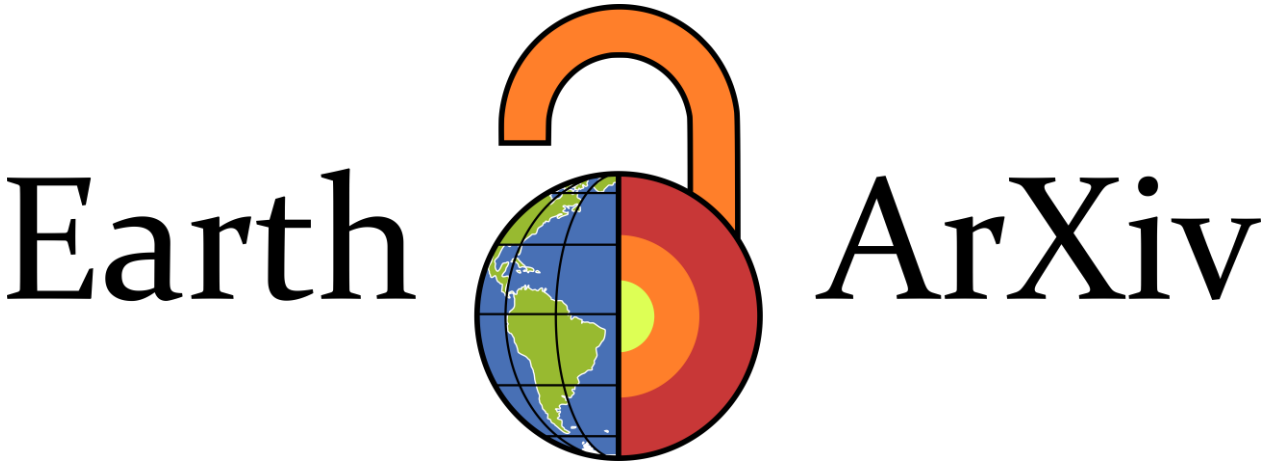

Denoising teleseismic data for deep Earth studies using a supervised deep-learning auto-encoder: a case study of diffracted waves from ULVZs

Stuart Russell[1,2], Katrin Hannemann[1], Carl Martin[3,4], Bennet Lindhorst[1], Waleed Esmail[5], Christine Thomas[1,6]

[1]Institut für Geophysik, Universität Münster, Correnstraße 24, 48149, Münster, Germany
[2]James Cook University, 1 James Cook Drive, Douglas, QLD, 4814, Australia
[3]Department of Geosciences, Utrecht University, 3584 CB, Netherlands
[4]Research School of Earth Sciences, The Australian National University, 142 Mills Road, Acton, ACT, 2601, Australia
[5]Institut für Kernphysik, Universität Münster, Wilhelm-Klemm-Straße 9, 48149, Münster
[6]Geological Survey of Denmark and Greenland, 1350 Copenhagen, Denmark

Key Points:

- A supervised auto-encoder effectively separates noise and signal in teleseismic traces, elucidating otherwise obscured signals.
- Models trained only on global noise and 1D synthetics perform well in a variety of applications including outside of the training paradigm.
- The denoiser can be readily incorporated in existing seismic workflows and has apparent potential in deep Earth seismic applications.

Plain language summary

Seismic waves generated by earthquakes are used to decipher the internal structure of the Earth. However, when these signals are recorded, they are often obscured or distorted by seismic noise – in the relevant frequency range, noise from the oceans is particularly pervasive. In this paper, we provide a demonstration that a deep-learning denoiser can effectively separate recorded teleseismic data into their signal and noise parts, allowing seismic phases to be more easily identified and studied. We present a computationally efficient workflow that uses synthetic seismograms combined with real seismic noise to train an algorithm, without the need for large datasets of real seismic phase observations or computationally costly simulations. Despite having never been exposed to real seismic signals, the resulting models prove effective at extracting real seismic signal and suppressing noise. We demonstrate a number of applications and provide recommendations for how to improve the performance of the model by applying a second training stage to the original model, known as transfer learning. Using the denoising workflow, we identify new short-period observations, as well as data sampling the Hawaiian ultra-low velocity zone from a new source-receiver geometry, demonstrating tangible benefits to deep Earth seismology. We hope to encourage the use of denoising in future studies, particularly for exploratory screening.

Corresponding author: Stuart Russell, stuart.russell@jcu.edu.au

Abstract

Seismic noise, particularly the microseism generated by the oceans, is a fundamental limitation on using short period (1 - 10 s) teleseismic data to study the deep Earth. Deep learning auto-encoders have proven effective at denoising seismic data in other applications. Here, we provide a demonstration of their potential for improving the quality of deep Earth seismic data, using S_{diff} post-cursors generated by ultra-low velocity zones as a case study. We present a computationally efficient and accessible workflow, utilising a supervised learning process with 1D synthetics and real, global seismic noise. This workflow does not require real teleseismic signals, which are limited in number and quality, and is computationally efficient, not requiring expensive 3D spectral-element simulations. We present a series of example applications, finding the denoiser to be effective at removing noise and extracting S_{diff} signal despite having not been trained with real teleseismic signals. Furthermore, we provide recommendations on how performance could be improved, including demonstrating transfer learning as an efficient and effective way to improve extraction of specific signals. Lastly we apply denoising to data sampling ULVZs, revealing observations of new short-period signals as well as sampling the Hawaiian ultra-low velocity zone using a novel source-receiver geometry from the East Pacific Rise to Japan. The potential impact of denoising deep Earth seismic data is apparent, and denoising can be implemented alongside traditional methods, such as stacking. We hope to encourage the use of deep learning denoising in future studies, particularly for exploratory screening.

1 Introduction

All seismic recordings contain noise. For seismic studies of the deep Earth, which typically use teleseismic arrivals from distant earthquakes to constrain the planet's internal structure, noise originates from a variety of different processes, both natural and anthropogenic. Examples of these processes include weather events such as storms, ocean waves, vehicle traffic, infrastructure such as wind farms, the seismic instrument itself, as well as many others (e.g. Bormann & Wielandt, 2013). Noise obscures and distorts the seismic signal and is a fundamental limitation on the use of seismic data. Typically, deep Earth seismologists use spectral filtering to isolate signals in a period band where noise is minimal, but this method is not effective when both signal and noise occur in the same period band. Where several seismic stations are available in a region, array methods such as stacking can be used to suppress incoherent noise (Rost & Thomas, 2002, 2009; Schweitzer et al., 2012), usually in conjunction with spectral filtering. Other methods for suppressing noise and extracting signal exist, including ensemble empirical mode decomposition (Han & van der Baan, 2015), wavelet based methods (Gaci, 2013; Mousavi et al., 2016; Langston & Mousavi, 2019), and the use of deep-learning, but these do not currently have common usage in deep Earth seismology.

Deep-learning denoisers in particular have received significant recent attention and have been shown to be effective in a wide variety of seismic applications, outperforming alternatives (e.g. Zhu et al., 2019; Heuel & Friederich, 2022; Tibi et al., 2023). Example uses to date include denoising regional-distance seismic data (Tibi et al., 2021), ocean bottom seismometers (Lin et al., 2024), offshore distributed acoustic sensing cables (Shi et al., 2025), ambient seismic field correlation functions (Viens & Van Houtte, 2020), microseismic event detection (e.g. Othman et al., 2021; Ma et al., 2023), observing deep Earthquake rupture (Shi & Denolle, 2023), downhole microseismic monitoring (Ge et al., 2024), seismic exploration (J. Li et al., 2023), nuclear monitoring (Steinberg et al., 2024), removal of anthropogenic noise (Heuel & Friederich, 2022; Yang et al., 2022), and for Martian seismology (Dahmen et al., 2022, 2024). Unsupervised deep-learning denoisers have been applied to improve receiver functions (Dalai et al., 2022; Feng et al., 2024), and recently to weak PcP signals (Oboué et al., 2026), but otherwise the use of deep-learning denoisers in deep Earth seismic studies is thus far limited, despite their potential to be highly impactful. In this study, we assess the potential of a supervised time-frequency domain denoising auto-encoder for denoising teleseismic data, using Ultra-low Velocity Zones (ULVZs) as a case study.

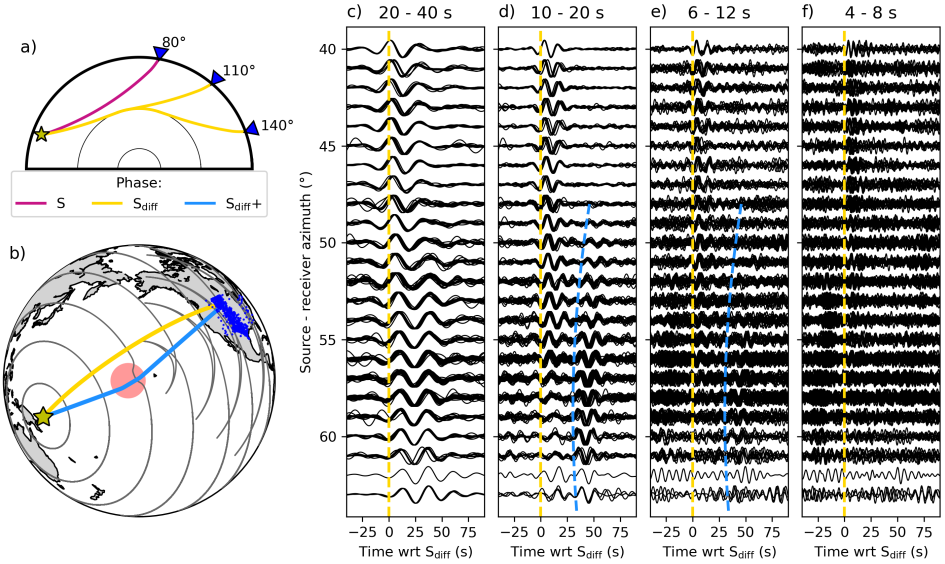

Figure 1. An introduction to S_{diff} and $S_{diff}+$. a) Example ray paths of S and S_{diff} in a cross-section of the Earth. b) Example ray paths and wavefronts of S_{diff} and $S_{diff}+$ in plan view for an earthquake in Papua New Guinea to stations the US, with the location of the Hawaiian ULVZ marked (Cottaar & Romanowicz, 2012). Wavefronts, curved by the ULVZ, are marked in grey. c) - f) Example waveforms for the source-receiver geometry shown in b) of Event 20100320. The same data are filtered in four different octave-width filter-bands as indicated. Arrival times of S_{diff} and $S_{diff}+$ are marked where visible. All traces are individually normalised and displayed rounded to the nearest degree of azimuth.

The core-mantle boundary (CMB) separates the solid silicate mantle from the liquid iron-alloy core at a depth of approximately 2890 km within the Earth (Oldham, 1906; Dziewonski & Anderson, 1981). The region approximately 200 - 300 km thick above the CMB is the lower thermal boundary layer of the convecting mantle and hosts significant seismic heterogeneity, of which ULVZs are the most extreme features. ULVZs are highly anomalous structures of severely reduced S-velocities – reported to be up to 50% slower than the background mantle – that sit atop the CMB (for overviews see e.g. Yu & Garnero, 2018; McNamara, 2019; Jackson & Thomas, 2021). They are typically only tens of kilometres thick but can be hundreds to thousands of kilometres in diameter (Cottaar & Romanowicz, 2012; Thorne et al., 2013; Yuan & Romanowicz, 2017; Cottaar et al., 2022). The formation mechanisms, composition, physical state and internal structure of ULVZs are all poorly understood; as is their relationship to other observed structures in the mantle, as well as to the dynamics of both the core and mantle. The key to better understanding ULVZs and their place in the Earth system is by better resolving their distribution, morphology and relationship to other structures.

One way to study ULVZs is with S_{diff} post-cursors. S_{diff} is a shear wave that diffracts along the CMB (Figure 1a). The vertically polarised S_V energy attenuates rapidly, while the horizontally polarised S_H is readily observed on the transverse (T) components of seismograms from approximately 100° to 130° epicentral distance. When S_{diff} encounters a ULVZ, the wave is lensed by the velocity contrast (Figure 1b), leading to a secondary arrival of energy. This secondary arrival is clear when S_{diff} arrivals are viewed as a function of source-to-receiver azimuth (Figure 1c-f), and is known as an S_{diff} post-cursor or S_{diff}+ (Cottaar & Romanowicz, 2012; Martin et al., 2023a).

Here, we consider some challenges faced by seismologists analysing S_{diff} and S_{diff}+ and highlight opportunities for inclusion of the DAE within existing workflows to potentially address these challenges. A major challenge is that the screening of record sections for S_{diff}+ remains a laborious process, typically involving expert examination of record sections (e.g. Martin et al., 2023a; Atkins et al., 2026), in which the visibility of S_{diff}+ depends on processing steps and signals can easily be overlooked – denoising provides a complementary perspective on this data. Beyond the screening stage, the behaviour of S_{diff}+ is frequency dependent. Observations of signals at shorter period can reveal the internal structure of ULVZs (e.g. Z. Li et al., 2022) and also facilitate the discovery of new ones (e.g. Martin et al., 2024), but observations at periods $<$10 s are challenging due to the high noise level (e.g. Figure 1). The dominant noise in this period range is the microseism, which is mid- to short-period noise generated by the coupling of the solid Earth with the oscillating oceans (e.g. Longuet-Higgins, 1950; Darbyshire, 1950; Iyer, 1958; Friedrich et al., 1998). There is a also a coverage challenge as there are fewer and more widely spaced seismometers deployed in the southern hemisphere. There is an over-reliance on stations in North America and Europe in particular, leading to over-sampling of certain regions of Earth's interior. Denoising can allow the use of lower quality data in regions with characteristically higher noise levels, such as ocean bottoms, ocean islands, or polar regions. In this way, it may be possible to sample more of the CMB in the search for ULVZs.

In this study, we use a deep-learning denoising auto-encoder (DAE) to denoise S_{diff} and S_{diff}+ waveforms that sample several different known ULVZs. We present a computationally efficient and accessible training workflow that uses Instaseis synthetics and real, global noise – thus not requiring computationally expensive 3D waveform simulations or data sets of real seismic phases. By using different variations of this training data, we attempt to ascertain the conditions in which the denoiser works well and how performance can be improved. We demonstrate the efficacy of denoising by applying to S_{diff} data, revealing observations of short-period signals from various ULVZs, and identifying data sampling the Hawaiian ULVZ using a novel source-receiver geometry. Finally, we discuss possible applications, and outline some recommendations for usage based on our experiences. We hope to encourage wider use of denoising in deep Earth seismology, particularly for screening, and, while this study uses S_{diff} and ULVZs as a case study, we envision denoising proving useful across a wide range of other deep Earth seismic applications.

2 Methods

2.1 Representation of signals

Our denoiser is based on the DAE model of Zhu et al. (2019), as implemented by Heuel & Friederich (2022) for removing noise generated by windfarms. Auto-encoders are a class of deep learning algorithms that aim to learn a sparse or latent representation of an input to extract specific features. The DAE comprises of two connected parts: the encoder compresses an input image through several convolutional layers into a latent space, while the decoder expands this latent representation back to its original dimensions through the same number of deconvolutional layers. This two-stage process allows the extraction of key features from the input image, and the accurate reconstruction of these key features in the output image. In the case of the DAE, the input is a time-frequency representation of a seismic trace and the aim is separate this into two independent parts: the signal and the noise.

In body wave studies of the deep Earth, teleseismic data are visualised foremost in the time domain, and a seismic trace, $y(t)$, can be viewed as the superposition of seismic signal time series, $s(t)$, containing the signal of interest, and a seismic noise time series, $n(t)$, that contains any number of other seismic oscillations that are not of interest,

$$y(t) = s(t) + n(t). \tag{1}$$

The terms signal and noise are relative and will vary between different seismologists, studies and applications. A seismic trace can be transformed into the time-frequency domain by application of some time-frequency transform, TFT,

$$X(t, f) = \mathrm{TFT}\{x(t)\}, \tag{2}$$

where x and X represent any seismic trace and the time-frequency representation of that trace. In the time-frequency domain the summation of signal and noise is equivalent to the time domain

$$Y(t, f) = S(t, f) + N(t, f). \tag{3}$$

The objective of the DAE is to decompose the input signal, Y, into approximations of S and N, denoted as $\hat{S}$ and $\hat{N}$, respectively. $\hat{S}$ is calculated from Y by

$$\hat{S}(t, f) = \hat{M}_S(t, f)\, Y(t, f), \tag{4}$$

and $\hat{N}$ in the same fashion

$$\hat{N}(t, f) = \hat{M}_N(t, f)\, Y(t, f), \tag{5}$$

where $\hat{M}_S$ and $\hat{M}_N$ are estimations of two mask functions, M_S and M_N, defined as

$$M_S(t, f) = \frac{1}{1 + \frac{|N(t,f)|}{|S(t,f)|}} \tag{6}$$

and

$$M_N(t,f) = \frac{\frac{|N(t,f)|}{|S(t,f)|}}{1+\frac{|N(t,f)|}{|S(t,f)|}} = 1 - M_S(t,f). \quad (7)$$

These mask functions have values between 0 and 1, and map each part of the time-frequency representation Y to the signal or noise. It is these mask functions that the algorithm seeks to learn. The training process is a supervised one where known signal and noise time series are summed, the time-frequency transforms calculated, and the algorithm optimises the process of separating these into the original parts.

It is worth noting two important implications of using mask functions rather than attempting to directly learn signal and noise. The first is that as the mask functions have a maximum value of 1 and are multiplied by Y, the outputs $\hat{S}$ and $\hat{N}$ cannot have amplitudes larger Y. This means that DAE cannot erroneously produce large amplitude signals where they are not already present. The second is that as the mask functions are applied equally to the real and imaginary parts of Y, the phase of Y is preserved in $\hat{S}$ and $\hat{N}$. This limits any time shifts or waveform distortion that can result from the application of the DAE.

The choice of time-frequency transform to use, as well as the length and sampling rate of the time series, are important parameters for the implementation of the DAE as they influence the size and resolution of the images used for learning. In this paper, the input time series all have a length of 600 s sampled at 5 Hz, therefore with a length of 3001 samples, which is appropriate for teleseismic data. For the time-frequency transform we follow Zhu et al. (2019) and Tibi et al. (2021) and use the short-time Fourier transform (STFT); this is as effective as and more efficient than the continuous wavelet transform (Heuel & Friederich, 2022). The STFT images used, and thus also the size of the mask functions, is 199×61, with 199 frequency points spanning 0 - 2.5 Hz with an interval of 0.013 Hz, and 61 time points spanning 0 - 600 s with an interval of 10 s.

2.2 Generation of training data

For the supervised training process we require high-quality examples of both signal and noise. For seismic noise there is no shortage of training data as most recorded seismograms are composed predominantly of noise. For teleseismic signal, however, there is a shortage of high-quality recordings. This is especially true when considering the often small (10^2 - 10^5 waveforms) datasets of specific deep Earth phases. Even where datasets are available, they are contaminated with noise, which will impede the learning process. In this paper we will circumvent this issue by using only synthetic signals to train the DAE (Mousavi & Beroza, 2022), an approach that has been successfully employed for other recent deep Earth machine learning applications (e.g. Zhou et al., 2024).

2.2.1 Noise datasets

We first construct a worldwide noise dataset using stations from the Global Seismographic Network (GSN; Romanowicz & Dziewonski, 1987; Ringler et al., 2020, 2022; Wilson et al., 2024). The GSN is a globally distributed seismic network currently comprising ∼140 stations, many of which have over 30 years near-continuous recording (see Supplementary Section S2 for more details). Using the ObsPy FDSN web-service (Beyreuther et al., 2010) we download seismic data for network codes CU, IC, II, and IU, selecting random 600 s second long time windows, starting from the deployment date of each station to the decommissioning date, or to September 2025 if the station is still active. As we are primarily interested in S_{diff}, we download only the horizontal components (north and east); across these two components, not accounting for outages, there would be >8000 years of possible seismic data.

While these seismograms are expected to be predominantly noise, some of the randomly selected time windows will contain seismic signal, which will contaminate the noise dataset. To minimise this contamination we use an STA/LTA discrimination process with an STA length of 100 samples, an LTA length of 1000 samples, and reject traces that have any STA/LTA values exceeding 5. Our final dataset

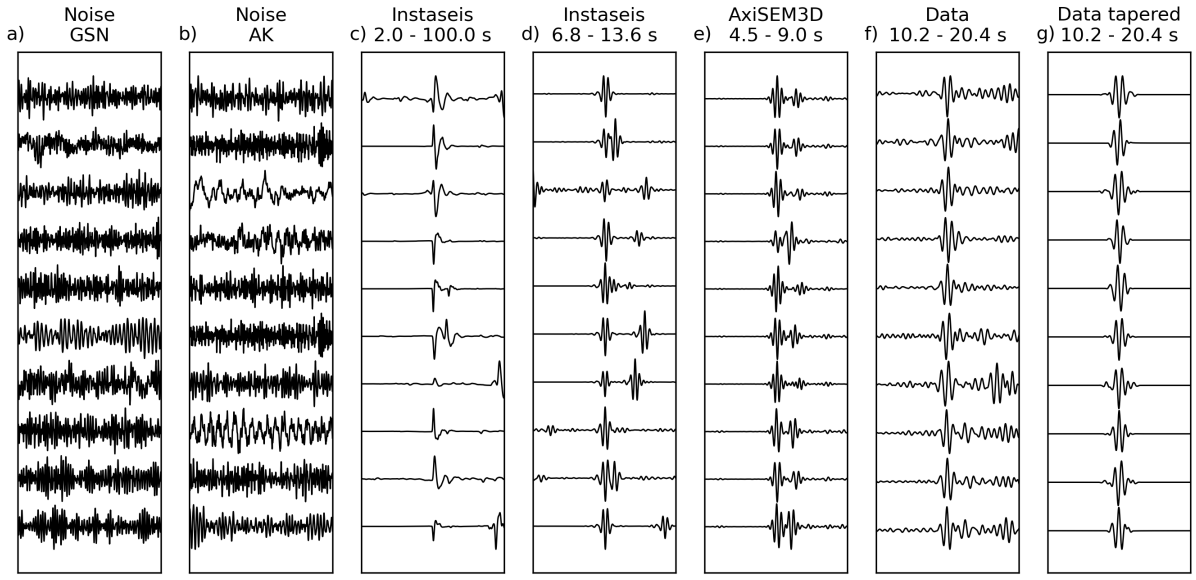

Figure 2. Example traces from various training datasets. a) Noise from the GSN; b) noise from the AK array in Alaska; c) Instaseis synthetics filtered 2 - 100 s; d) Instaseis synthetics filtered 6.8 - 13.6 s; e) AxiSEM3D synthetics with a ULVZ filtered 4.5 - 9.0 s; f) real data filtered 10.2 - 20.4 s; g) real data filtered 10.2 - 20.4 s with a taper applied around the S_{diff} signal. All traces are individually normalised. Only the central 300 s of each trace is shown.

comprises 512,000 individual traces, which is ∼0.12% of the possible available traces, assuming no overlaps or outages, for the GSN. Using this workflow we can also construct region- or network-specific noise datasets. Example noise traces are shown in Figure 2a and 2b.

2.2.2 Synthetics - Instaseis

Synthetic seismograms for a particular model of the Earth can be calculated using spectral-element simulations of seismic wave propagation (e.g. Komatitsch & Tromp, 1999; Chaljub et al., 2007). However, these are computationally extremely costly and computation time scales with frequency to the power of 3 - 4, depending on the parameterisation used (Leng et al., 2019). For global simulations at relevant periods, these simulations can take thousands or even millions of CPU hours per hour of seismogram (e.g. Komatitsch, 2011; Nissen-Meyer et al., 2014; Leng et al., 2019). Generating synthetic training data can therefore be very costly if done from scratch.

A large part of this computational cost can be avoided by using Instaseis (van Driel et al., 2015), a Python package that accesses pre-computed Green's function databases, calculated using AxiSEM (Nissen-Meyer et al., 2014), and generates synthetic seismograms for 1D Earth models near instantaneously by utilising the reciprocity of the Green's functions. We use the anisotropic version of the Preliminary Reference Earth Model (PREM; Dziewonski & Anderson, 1981), `PREM_a_2s`, hosted by Syngine (Hutko et al., 2017; Krischer et al., 2017). This database has a minimum valid period of 2 s, controlled by the mesh spacing during computation, and a maximum valid period of 100 s, due to the neglect of full-gravity by application of the Cowling approximation (Komatitsch & Tromp, 2002; van Driel et al., 2021); the synthetics must therefore be filtered from 2 - 100 s as a minimum requirement, but can be filtered to more specific period bands within these limits to encourage the extraction of signal of a specific frequency. We create datasets for six overlapping octave-width filter-bands, every half octave from 2 s. This covers the range of previous ULVZ observations with S_{diff}.

We generate traces such that each 600 s long trace has a synthetic T-component seismic arrival at the centre. To increase the diversity of the training data we do not limit ourselves to S_{diff}, but also include S and SS arrivals with an epicentral distance between 50° and 130°, source depths down to 600 km, and randomly generated but realistic moment tensors. Example Instaseis traces are shown in Figure 2c and 2d.

2.2.3 Synthetics - AxiSEM3D

By using Instaseis we save a significant computational cost, but are limited to 1D Earth models that do not contain any 3D structure, including ULVZs. The Instaseis synthetics thus only contain S_{diff} and not S_{diff}+, and are unlike real seismic data that have interacted with a highly heterogeneous mantle. We therefore also calculate 3D synthetics using AxiSEM3D (Leng et al., 2016, 2019), which has been used extensively to model S_{diff}+ and other ULVZ-associated seismic signals (e.g. Z. Li et al., 2022; Wolf & Long, 2023; Martin et al., 2024; Jagt et al., 2024; Wolf et al., 2024).

Due to the high computational cost, even with optimised inputs (following Russell et al., 2025), we only run a single AxiSEM3D simulation with a minimum period of 4 s. Our AxiSEM3D training dataset therefore only has a single source depth and a single set of ULVZ parameters, and does not contain 3D structure except for the ULVZ. We use a source depth of 500 km for a 3D model containing a cylindrical ULVZ (radius: 400 km; height: 20 km; S-velocity reduction: 20%) embedded on top of the core-mantle boundary within PREM. See Supplementary Section S3 for further details on the AxiSEM3D simulation. Example AxiSEM3D traces are shown in Figure 2e.

2.2.4 Real S_{diff} data

Our final training dataset utilises real S_{diff} traces recorded in the contiguous USA. From the Martin et al. (2023a) dataset of S_{diff} record sections that sample of Hawaiian ULVZ, we screen the data by eye and manually select traces that have high SNR S_{diff} arrivals in the 10 - 20 s period-band, giving 724 traces from 30 different earthquakes recorded in the contiguous US. Only the quality of S_{diff}, not S_{diff}+, is considered when selecting traces. We allow traces to be repeated and augment each trace in the frequency domain by implementing a random phase shift. We generate an augmented training dataset of 51,200 traces. These traces contain contaminating noise, and to assess the effect of this we create a further training dataset. For this dataset, a taper of variable length is applied centred on the S_{diff} peak which removes all noise along with any subsequent S_{diff}+ arrivals. Example traces from both of these datasets are shown in Figure 2f and 2g.

2.3 Model setup

The DAE is built in Python using the TensorFlow software library and our model setup follows that used by Heuel & Friederich (2022). The two input layers are the normalised real and imaginary parts of the STFT of a seismic trace. The two output layers are the estimates of the two mask functions, $\hat{M}_S$ and $\hat{M}_N$. The targets, M_S and M_N, are calculated directly from the training data using Equations (6) and (7), and the model is optimised by minimising a loss function, ℓ,

$$\ell = \frac{\ell_S + \ell_N}{2}, \tag{8}$$

where ℓ_S and ℓ_N are the mean binary cross-entropy loss for the signal and noise masks, respectively, and

$$\ell_X = \frac{1}{2N_0N_1} \sum_{i=0}^{N_0} \sum_{j=0}^{N_1} M_{X_{ij}} \log \hat{M}_{X_{ij}} + \left(1 - M_{X_{ij}}\right) \log \left(1 - \hat{M}_{X_{ij}}\right), \tag{9}$$

where X denotes whether the masks are for signal or noise, and the mask function images (and the input STFT images) have size $N_0 \times N_1$, where N_0 is the number of samples along the frequency axis, and N_1 is the number of samples along the time axis – in this application 199×61.

The encoder reduces the dimensions of the input layers and consists of six 2D convolutional layers with kernel size of 3×3 and stride length of 2×2, each followed by a leaky rectified linear unit (ReLU) activation function, batch normalisation (BN) (Ioffe & Szegedy, 2015), and a dropout layer (Srivastava et al., 2014) with a drop rate of 0.1. The setup of the decoder is the inverse of the encoder and increases the dimensions of the sparse representation to obtain the two output masks. 2D deconvolutional layers are followed by cropping, leaky ReLU, BN, and dropout. The final layer undergoes a further 2D convolution with a kernel size of 1×1 and a softmax activation function (Bridle, 1990) to generate the mask functions. As in previous implementations of this algorithm we also use skip connections to improve convergence (Ronneberger et al., 2015; Liu & Yang, 2018). With this configuration, the DAE has 1,183,002 trainable parameters.

2.4 Training process

During training, the data are randomly split into 80% for training and 20% for validation. We train our model on a single NVIDIA GeForce RTX 4090 GPU on the PALMA-II cluster of the University of Münster. We use a batch size of 32 traces; each batch is generated by selecting noise traces and signal traces at random from the pool of traces that have not been used yet that epoch. Traces are first demeaned and normalised and then random time shifts of up to ± 300 s are applied to the signal traces, which are then cut and zero-padded to the original length. The traces are scaled such that either signal or noise is up to 3 times larger amplitude than the other. The signal and noise are then summed to make a trace of 'noisy signal', which is again demeaned and normalised such that the absolute maximum amplitude has value 1 – this normalisation is essential for the convergence of the algorithm. In order to help the model identify pure noise we also sometimes set the entire signal trace amplitude to zero with an occurrence of 1%. Finally the input layers are calculated by taking the real and imaginary parts of the STFT of each trace. We use an Adam optimiser with a learning rate of 10^{-6}. We train our model for a maximum of 200 epochs and if the loss of the validation data does not improve for 10 successive epochs then training is stopped early. We have explored various hyper-parameter combinations including different batch sizes, learning rates and numbers of convolutional layers, settling on a configuration that gives stable learning curves, with sufficient validation loss in a reasonable training time-frame. Training typically takes 1 - 2 days.

As well as simply running a single training process, by which a model is optimised from a set of random initial weights, it is also possible to start from a previously trained model and update the weights for a new but related task – this is known as transfer learning (e.g. Weiss et al., 2016). Transfer learning allows computationally efficient training for tasks with relatively few training data, as the majority of the learning can be done with the first training step. The second training step is also typically much faster, as the learning distance from the starting model to the final model is shorter. In this way, a model trained on noise from the GSN could be retrained for a specific network with less available data. Alternatively, a model trained with Instaseis synthetics could be retrained with a smaller dataset of AxiSEM3D synthetics or even real data. During retraining we freeze the weights of the encoder, allowing only the decoder to retrain, which preserves the earlier layers typically responsible for feature extraction (Iman et al., 2023). To prevent over-fitting, the transfer learning stage has the same early stopping criterion as for the original training. A schematic demonstrating the two-step training process is shown in Figure 3.

2.5 Test S_{diff} data

Using the ObsPy FDSN web-service (Beyreuther et al., 2010) and HinetPy (Tian, 2024), we download a number of T-component seismic record sections from 2000 to 2024 that contain clear S_{diff} arrivals. Many of these also contain S_{diff}+ arrivals associated with the Hawaiian (Cottaar & Romanowicz, 2012; Z. Li et al., 2022; Martin et al., 2023a), Galápagos (Cottaar et al., 2022), Pitcairn (Z. Li

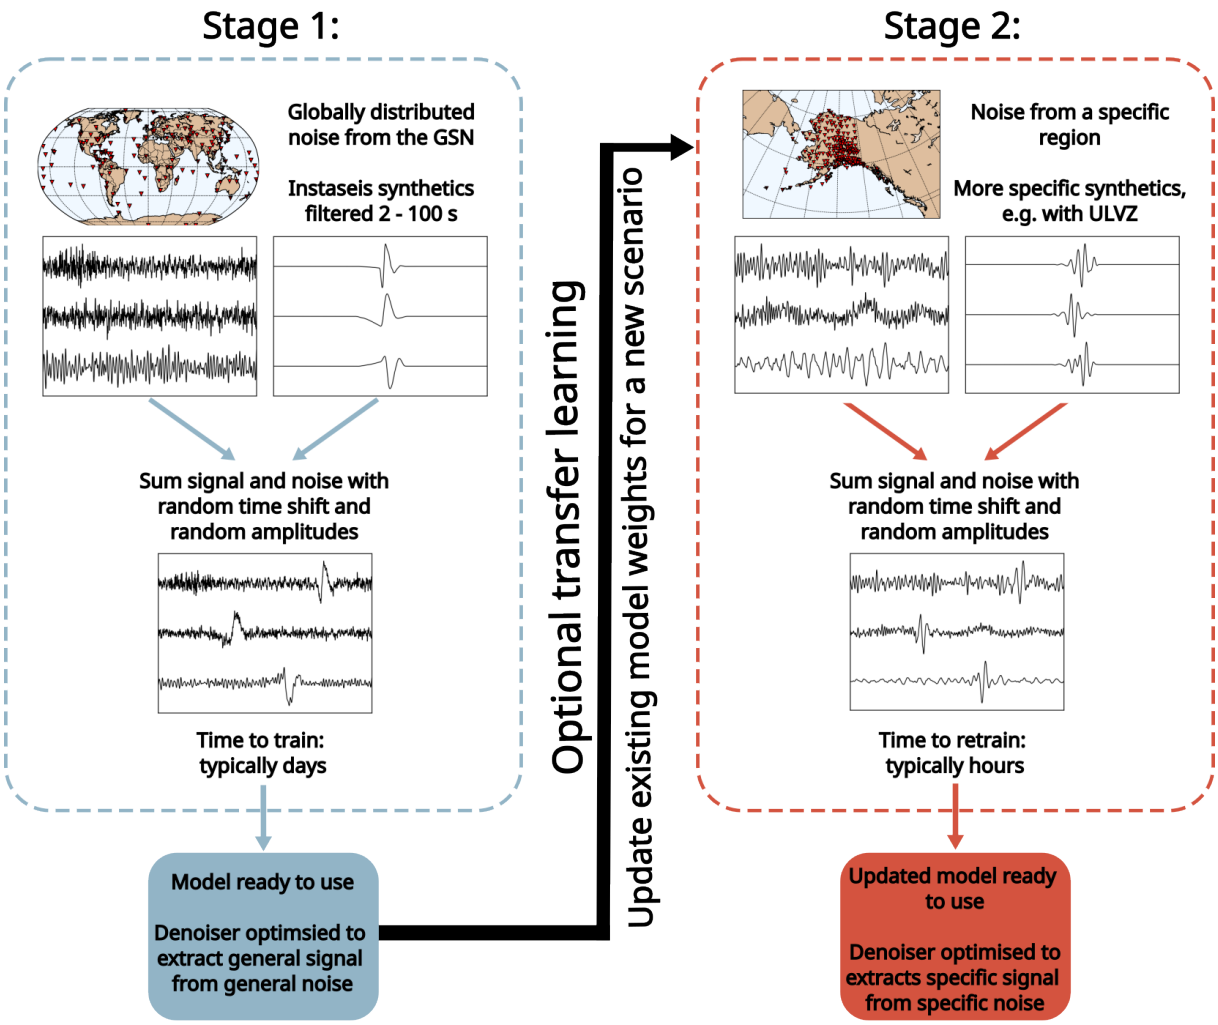

Figure 3. A flow chart demonstrating the transfer learning process. Training can be ended after either stage.

et al., 2024), Central Pacific (Martin et al., 2024) and Vanuatu (Martin et al., 2025) ULVZs. Station locations vary, but are mostly in the contiguous US, Alaska and Japan. As with the training data, these traces are 600 s long and sampled at 5 Hz. These record sections form the test dataset by which we will assess the performance of the trained DAE models on real data.

3 Verification and testing results

3.1 1D synthetic testing

Before application to real S_{diff} record sections, we first benchmark a trained DAE model against synthetic test datasets. The first tests are performed on an unseen dataset of 10,000 S_{diff} Instaseis synthetics and 10,000 noise traces from the US Transportable Array (FDSN network code TA), constructed in the same fashion as the original training data. The aim of these tests is to verify that the DAE has learned its task adequately, and to assess its performance over a range of noise levels. Noise level is defined by specifying a peak amplitude ratio, where during summation the signal and noise are normalised at a fixed ratio. An amplitude ratio of 1 is where the signal and noise have the same maximum amplitude, an amplitude ratio of 2 is where the signal has twice the amplitude of the noise, and so on. During summation, the signal is randomly shifted to test performance regardless of the location of signal within the trace.

A summary of the results are shown in Figure 4, with further detailed results presented in Supplementary Section S4. We quantify timing errors by cross-correlating the recovered and original traces, taking the maximum correlation to be the signal location, and we quantify amplitude errors by calculating the mean-squared-error (MSE) between the input and recovered signal traces,

$$\mathrm{MSE} = \frac{1}{n}\sum^{n}\left(s(t) - \hat{s}(t)\right)^2, \tag{10}$$

where $s(t)$ is the input signal, $\hat{s}(t)$ is the recovered signal, and n is the number of test traces, in this case 10,000. The MSE is calculated for a 20 s window around the phase arrival so that amplitude differences are explicitly quantified for the S_{diff} phase itself.

Across all three amplitude ratios shown, the signal is typically visually well-extracted. There is a visible reduction in performance at higher noise levels and there is also an associated broadening of the error distributions. The median absolute time shift is 0 or 1 samples (sampling rate is 5 Hz) for amplitude ratios down to 0.5, so where signal is half the amplitude of the noise. For smaller amplitude ratios the median error increases rapidly as the DAE fails to extract the signal at all. For all amplitude ratios tested, the most common absolute time shift is always 0 or 1 samples, indicating that even where errors are more likely in noisier data, the most common outcome is no significant time shift.

These results indicate that for reasonable noise levels, signal is accurately extracted, but errors are lower where the noise level is already low. This is unsurprising as signal that is already visible is clearly going to be easier for the DAE to identify. While these results suggest that the DAE is better at 'tidying up' noisy data rather than uncovering signal that is totally obscured by noise, the DAE certainly demonstrates remarkable aptitude for the latter. Signal with an amplitude ratio of 0.5, as shown in Figure 4a, would likely be considered devoid of signal with traditional visual screening, and those with an amplitude ratio of 1.0 would be considered ambiguous. That the DAE can elucidate these signals, indicating where hidden signals may be present even if extraction is not point-for-point exact, is remarkable and suggests that the DAE has learned well within its training paradigm.

3.2 3D synthetic testing

To test generalisation beyond the training distribution, we apply the Instaseis-trained DAE to AxiSEM3D synthetics that contain Sdiff+ arrivals not present in the training data. These are the same AxiSEM3D synthetics described in Section 2.2.3 and Supplementary Section S3 and these have not

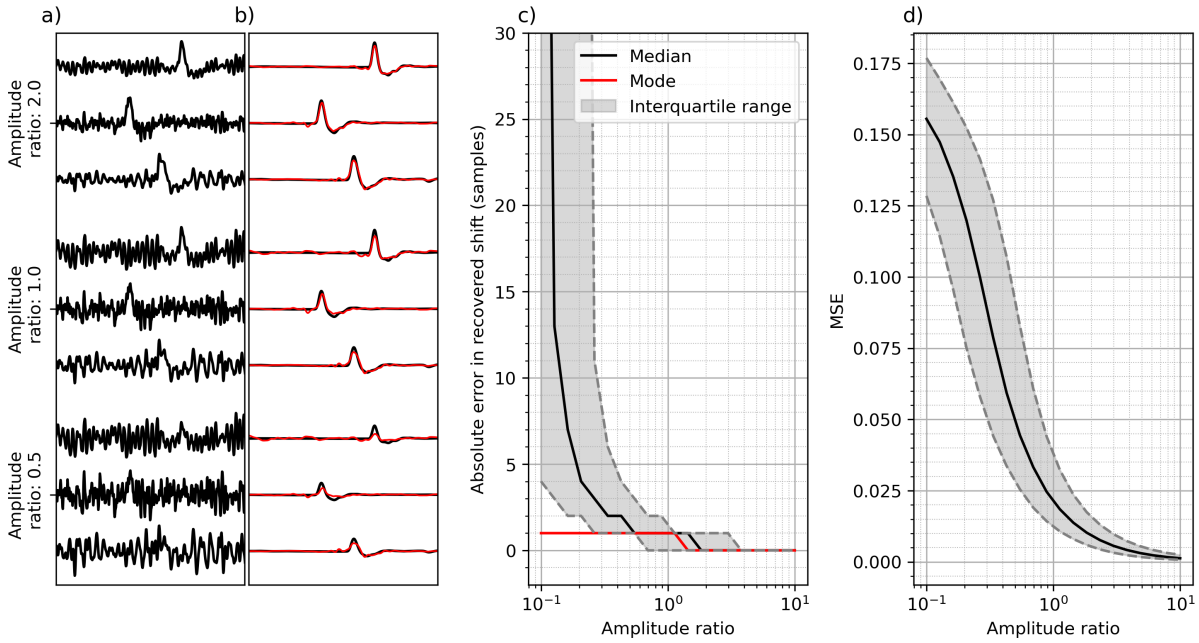

Figure 4. Synthetic testing results on unseen noise and Instaseis synthetics. a) The input signal that is the sum of an Instaseis S_{diff} arrival and noise from the US Transportable Array, for three different amplitude ratios. b) The input (black) and recovered (red) signal. c) The best-fitting shift in the recovered traces measured by cross-correlation, quantifying timing errors as a function of the imposed amplitude ratio. d) As for c), but showing the MSE between the input and recovered signal traces, quantifying amplitude errors.

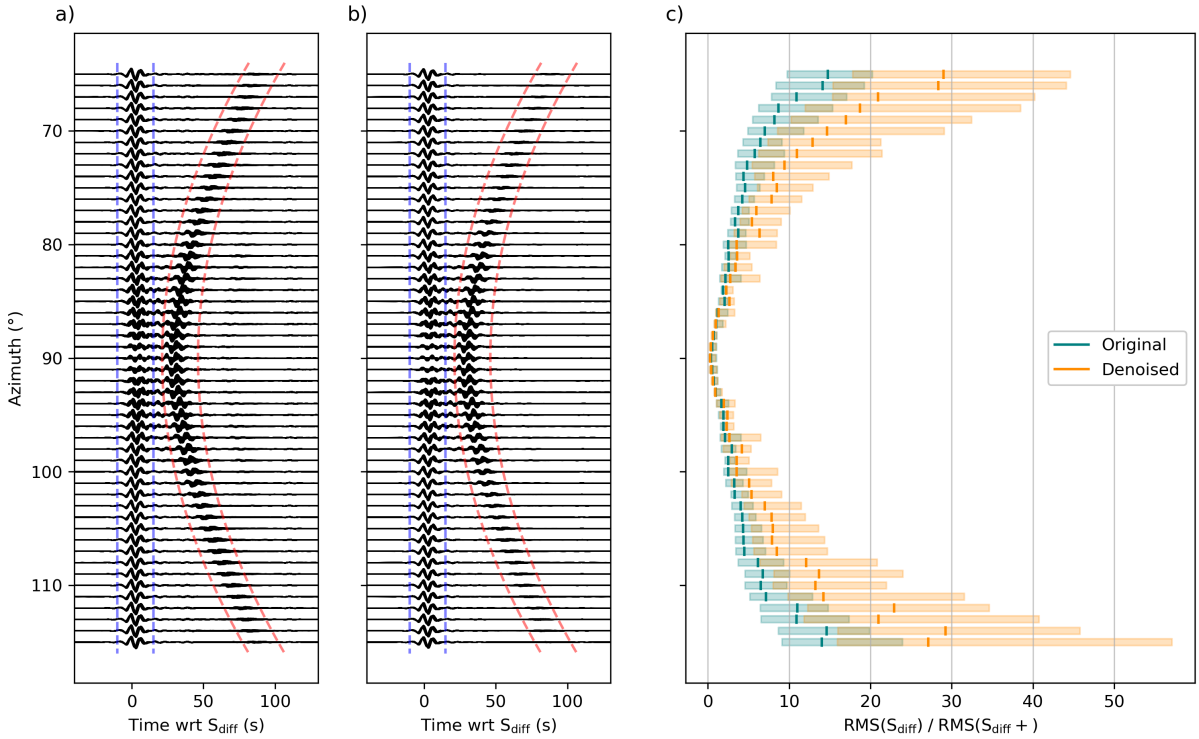

Figure 5. a) Input clean AxiSEM3D synthetics, b) the denoised synthetics. The blue and red dashed lines mark the 5 s long windows used for calculating the RMS of S_{diff} and $S_{diff}+$, respectively. c) The RMS ratio of S_{diff} to $S_{diff}+$ for the original synthetics (teal) and the denoised synthetics (orange). The lines show the median for each 1° wide azimuth bin, and the shaded boxes show the interquartile range.

been used to train this particular DAE. The aim of this test is to assess the performance of the DAE at extracting signals that it has not explicitly been trained to extract.

We use the DAE to denoise clean AxiSEM3D synthetics, so any discrepancies before and after denoising are due to the DAE itself. Results are shown in Figure 5, and further results with noise are presented in Supplementary Section S5. Qualitatively, we find that the DAE trained on Instaseis synthetics does satisfactorily extract the $S_{diff}+$, as shown by the likeness of the input and recovered synthetics shown in Figure 5a and 5b. We quantitively assess the relative performance on S_{diff} and $S_{diff}+$ by comparing the root-mean-square (RMS) amplitude of a 5 s long time window around each phase. As is expected, the off-axis $S_{diff}+$ is lower amplitude and the ratio $RMS_{S_{diff}} / RMS_{S_{diff}+}$ correspondingly has higher values off-axis. However, in the denoised synthetics this is even more exaggerated, indicating that off-axis $S_{diff}+$ is partially suppressed by the DAE.

This is unsurprising as off-axis $S_{diff}+$ is most distorted by the ULVZ and is farther from the training paradigm than S_{diff} or on-axis $S_{diff}+$. Whilst this suggests that caution must be used when interpreting denoised data that is outside of the training regime, that Figure 5a and 5b are so similar provides reassurance that the DAE can elucidate signals that it has not been specifically trained to extract, and motivates application to real $S_{diff}+$ data.

3.3 Application to real signals

Following the positive results of the synthetic tests, we apply the DAE to real $S_{diff}+$ waveforms. In the example shown in Figure 6, we apply a DAE to Event 20131123 (see Supplementary Section

S1 for details of all events referenced in this paper) recorded in the contiguous US that samples a ULVZ in the central Pacific (Martin et al., 2024). This DAE is trained on GSN noise, so has not seen the noise from any of the stations in the test data, and is trained on synthetic Instaseis signal, so has never seen real seismic waveforms. The synthetics used to train this particular DAE were filtered at 2 - 100 s, so the real traces are filtered to 2 - 100 s prior to denoising and are then filtered to the narrower period bands after denoising.

Figure 6 shows the original and denoised traces, and signal is extracted across all period bands. S_{diff}+ is clearly extracted in the 6.8 - 13.6 s and 4.5 - 9.0 s period bands but is weaker off-axis, as in the synthetic testing. S_{diff}+ is absent or weak in the extracted 3.0 - 6.0 s and 2.0 - 4.0 s period bands – this could also be mirroring the trend observed in the synthetic testing of poorer performance at higher noise levels, or could be a genuine lack of high-frequency S_{diff}+ from this ULVZ (Martin et al., 2024). Unlike in synthetic testing where a ground truth is known, it is not possible to know for real data. As is apparent from the difference between the top and bottom rows in Figure 6, there is always some amplitude loss due to non-perfect partitioning of signal and noise (Zhu et al., 2019) – this is consistent with the synthetics tests. These results show the promise of the DAE in extracting real teleseismic signal, including exposing signal that is barely visible in the original data, despite being trained on purely synthetic arrivals.

3.4 Specificity of the noise

Different locations on Earth's surface experience different seismic noise regimes. The aim of using the DAE in this study is primarily to suppress the microseism, which is globally pervasive in the period bands in which we wish to study ULVZs. In order to test how the DAE trained with noise from the GSN performs in different locations, we build three localised noise datasets: one each for the contiguous US (using network codes TA and US), Alaska (using network codes AK and AV), and Japan (using F-net stations). These are varied locations – the US is largely a continental environment, Alaska is a high-latitude environment, and Japan is an island. For each dataset we download 128,000 noise traces and we retrain a DAE that was originally trained with GSN noise.

We apply each retrained DAE to events recorded in each of these regions and assess their performance. The aim is to test how well local noise conditions are suppressed by a DAE trained on global noise, and to see if retraining to local noise conditions improves on this. We define SNR in decibels (e.g. Zhu et al., 2019; Heuel & Friederich, 2022) as

$$\mathrm{SNR} = 20 \log_{10} \left(\frac{\mathrm{RMS}_{\mathrm{signal}}}{\mathrm{RMS}_{\mathrm{noise}}} \right) \tag{11}$$

where $\mathrm{RMS}_{\mathrm{signal}}$ is the RMS amplitude of a 40 s long time window extending from the S_{diff} arrival time, t_0, to t_0+40 s, and $\mathrm{RMS}_{\mathrm{noise}}$ is the RMS amplitude of a 40 s long time window from t_0-70 s to $t_0 - 30$ s. We also define the improvement in SNR, ΔSNR, as

$$\Delta\mathrm{SNR} = \mathrm{SNR}_{\mathrm{denoised}} - \mathrm{SNR}_{\mathrm{original}}. \tag{12}$$

Figure 7 shows ΔSNR for each DAE for three events. The waveforms are filtered at 2.0 - 100 s to encompass noise suppression across the entire valid period range. In each region the DAE that achieves the best ΔSNR is typically the one retrained with noise from that region, but the improvements over the GSN DAE are marginal. In addition, retraining with noise from a different region usually reduces performance compared to the GSN DAE, although this is marginal. Furthermore, the waveforms themselves look highly similar (Supplementary Section S6).

These results suggest that the predominant local noise conditions in these three locations are well-represented in the GSN training data and correspondingly suppressed by the DAE; this is expected as the microseism is the main noise source affecting these data. We therefore suggest that the DAE

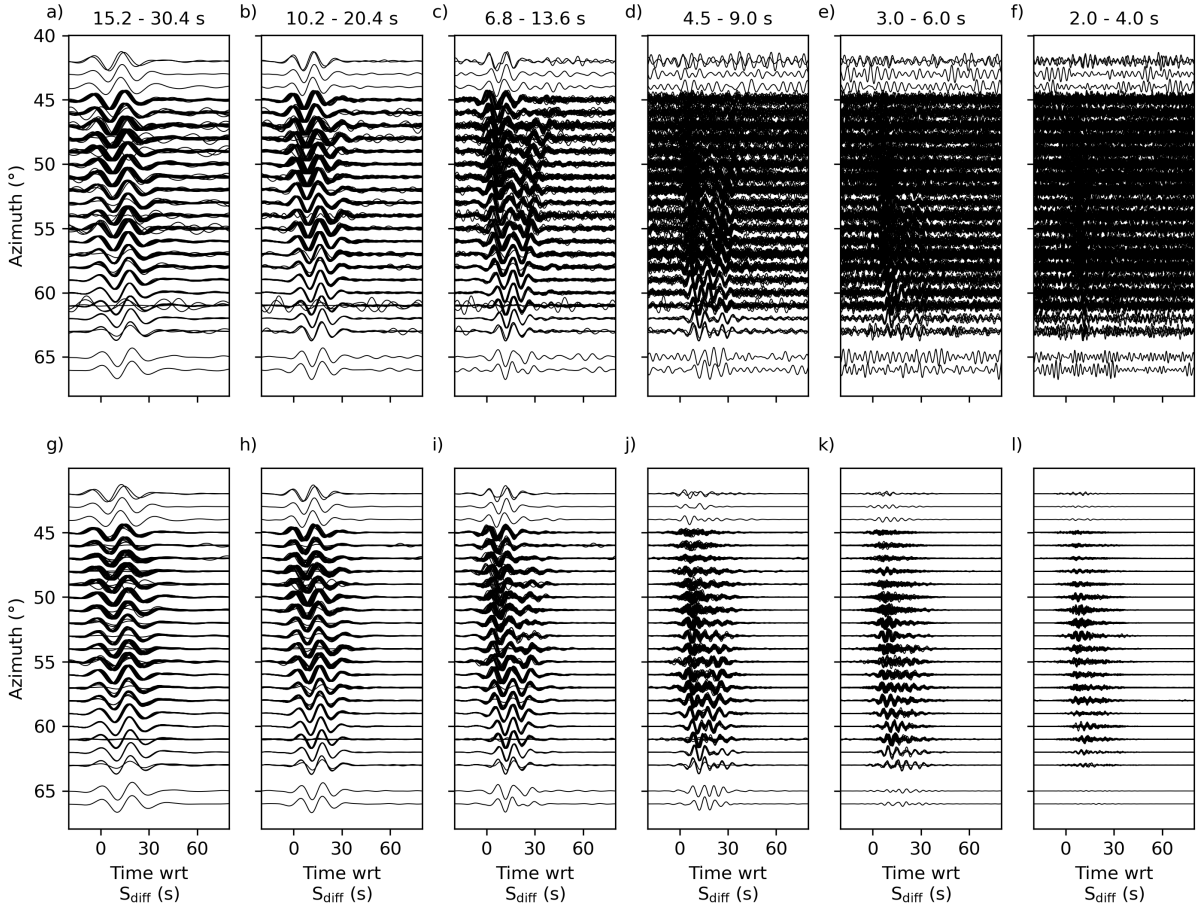

Figure 6. Example denoised waveforms for Event 20131123 recorded in the contiguous US in six different octave-width filter-bands. a) - e) Original waveforms in the period bands indicated. In the top row, traces are individually normalised such that the maximum amplitude corresponds to 0.75° azimuth, and traces are plotted rounded to the nearest degree of azimuth. f) - j) Denoised waveforms in the period bands indicated above each plot. The amplitudes in panels f) - j) are consistent with corresponding waveforms in panels a) - e).

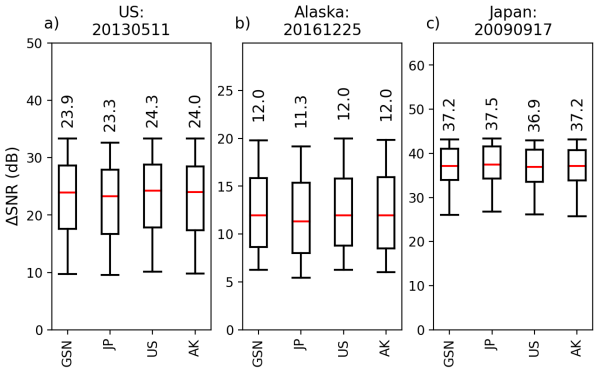

Figure 7. The performance of retrained DAE models on data recorded in different regions. a) The improvement in SNR in decibels for Event 20130511 recorded in the contiguous US, b) for Event 20161225 recorded in Alaska, and c) for Event 20240223 recorded in Japan. The median is shown in red, the edges of the boxes are the 25^{th} and 75^{th} percentile, and the ends of the arms are the 10^{th} and 90^{th} percentile. The value of the median is given above.

trained with GSN noise likely performs well globally. Nonetheless, for individual stations with highly specific noise regimes, for example those close to anthropogenic noise sources such as wind farms, highways, or deployed in urban environments, retraining with local noise may still be necessary to maximise noise suppression. In this fashion, a single DAE trained with GSN noise could be used for screening globally, but for more detailed down-stream analyses a retrained DAE could be implemented once a target set of signals is identified.

3.5 Specificity of the synthetics

As has been shown in the synthetic testing, $S_{diff}+$ can be suppressed by the Instaseis-trained DAE, even in the absence of noise. We therefore investigate whether retraining the DAE with AxiSEM3D synthetics that contain $S_{diff}+$ arrivals leads to better extraction of $S_{diff}+$ in real data. Figure 8 shows a DAE that has been retrained on AxiSEM3D synthetics applied to Event 20121207, recorded in the contiguous US, in the 6.8 - 13.6 s filter-band. This event has clear $S_{diff}+$ arrivals caused by the Pitcairn ULVZ (Z. Li et al., 2024). Whilst the denoised waveforms for all models are a clear improvement over the original data, all models show a decrease in the ratio of the RMS amplitude of $S_{diff}+$ to that of S_{diff} relative to the original data (panel e), measured in a 30 s window around each phase, from 65° to 75°. Consistent with expectations, this indicates that all of the DAEWs better extract S_{diff} than they do $S_{diff}+$.

The degradation of $S_{diff}+$ is somewhat remedied by retraining the DAE with AxiSEM3D synthetics that contain synthetic $S_{diff}+$ arrivals, which result in higher $S_{diff}+/S_{diff}$ ratios. $S_{diff}+$ extraction is further improved by retraining with synthetics that are filtered to the 6.8 - 13.6 s filter-band. It is important to note that both of these models also have a slightly lower improvement in SNR for S_{diff} (panel f) compared to the Instaseis-trained DAE, indicating that they are less effective at separating S_{diff} from noise, although due to the logarithmic definition of SNR the visual difference between the waveforms is not substantial.

That panels b, c and d are so visually similar indicates that the Instaseis-trained DAE already performs well. However, these results show that further specification during training can lead to better extraction. These results also reiterate the danger of interpreting relative amplitudes in denoised data, as not all phases may be extracted equally. The outcomes of retraining with 3D synthetics are variable and does not always improve the extraction of $S_{diff}+$ (see Supplementary Section S7), which

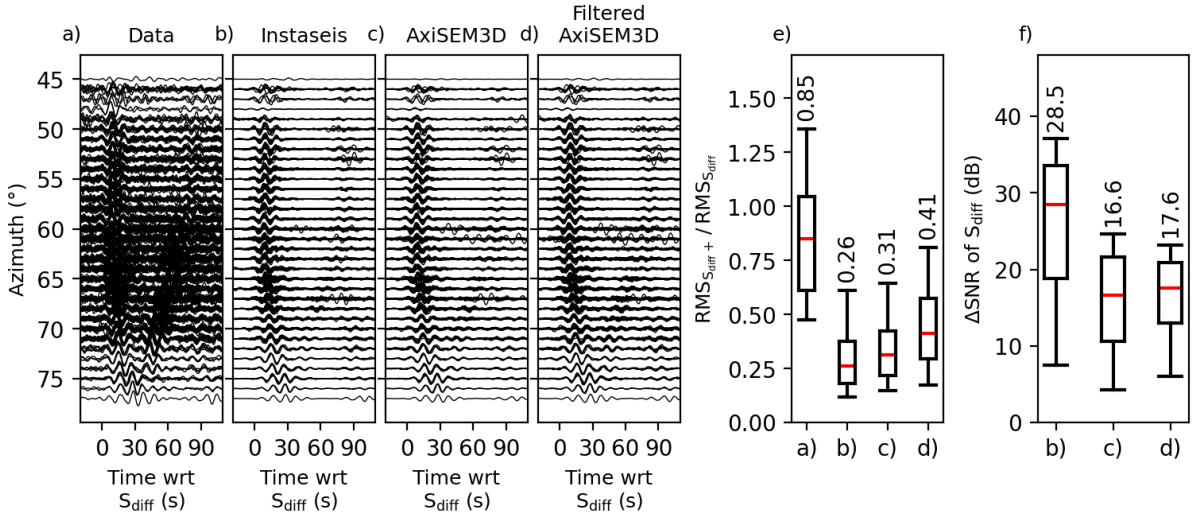

Figure 8. The performance of the DAE when retrained using AxiSEM3D synthetics applied to waveforms from Event 20121207 in the 6.8 - 13.6 s filter-band. a) The original data with each trace individually normalised such that the maximum amplitude corresponds to 0.75° azimuth, and traces are plotted rounded to the nearest degree of azimuth, b) denoised with the DAE trained with Instaseis synthetics, c) denoised with the DAE retrained with AxiSEM3D synthetics, and d) denoised with the DAE retrained with filtered AxiSEM3D synthetics. e) The ratios of the RMS of $S_{diff}+$ and S_{diff} for the waveforms in panels a) - d). The median is shown in red, the edges of the boxes are the 25^{th} and 75^{th} percentile, and the ends of the arms are the 10^{th} and 90^{th} percentile. The value of the median is given above. f) As for e) but showing the improvement in SNR for S_{diff} in decibels for panels b) - d) relative to a). Amplitudes in panels b) - d) are consistent with panel a).

is likely because the AxiSEM3D synthetics are too specific. Due to the computational cost involved, we only run one AxiSEM3D simulation with a single ULVZ that has only one combination of height, radius, S-velocity reduction, and a fixed source depth, parameters which may not be similar enough to an individual ULVZ to improve extraction. Generating a large dataset of ULVZ synthetics for a wide range of source and ULVZ parameters would be computationally extremely costly, but would be necessary for training a DAE that maximally extracts all $S_{diff}+$ signals.

3.6 Incorporating real training data

The final training datasets that we test incorporate real data. We retrain a DAE that was trained using the GSN and Instaseis synthetics using two different datasets of real data filtered to the 10.2 - 20.4 s filter band: both tapered and non-tapered. Figure 9 shows these models applied to Event 20150519 recorded in Japan, which has $S_{diff}+$ arrivals associated with the Vanuatu ULVZ (Martin et al., 2025). The real data used during retraining are recorded in the contiguous US, and therefore these data from Japan do not form part of the training dataset. All models have an improvement in SNR over the original data, but this is highest for the DAE retrained using tapered data. The DAE retrained with the untapered data sees only a minor improvement in SNR, although due to the logarithmic definition of SNR the visual improvement of the waveforms over the original is still significant. This demonstrates the insidious consequences of noise contamination in the training signals, which leads to the retrained DAE not performing as well as the original Instaseis DAE.

These results suggest that performance can be improved by incorporating real data even when applied to unseen data, but for best results one must be careful to isolate the signal as best as pos-

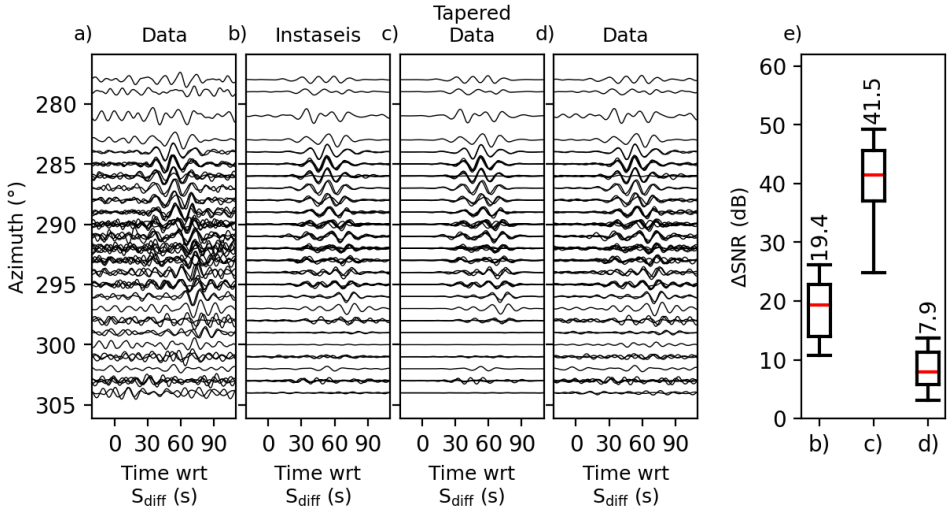

Figure 9. The performance of the DAE when retrained with real data, applied to Event 20150519 recorded in Japan in the filter-band 10.2 - 20.4 s. a) The original data with each trace individually normalised such that the maximum amplitude corresponds to 0.75° azimuth, and traces are plotted rounded to the nearest degree of azimuth, b) denoised with the DAE trained with Instaseis synthetics filtered 10.2 - 20.4 s, c) denoised with the DAE retrained with real data that has been tapered and filtered 10.2 - 20.4 s, and d) denoised with real data that has been filtered 10.2 - 20.4 s. e) The improvement in SNR for for the waveforms in panels b) - d) relative to a). The median is shown in red, the edges of the boxes are the 25^{th} and 75^{th} percentile, and the ends of the arms are the 10^{th} and 90^{th} percentile. The value of the median is given above. Amplitudes in panels b) - d) are consistent with panel a). Note that as the waveforms are severely delayed due to the Pacific LLSVP, the 40 s long time window used for the S_{diff} RMS starts at 30 s after the theoretical S_{diff} arrival time. The time window for the noise is unchanged.

sible to avoid noise contamination. Constructing training datasets of specific low-amplitude signals where noise levels are typically high may not be possible, and therefore this strategy may not help for extracting data that is most affected by noise. Real data are best incorporated during a second training stage due to the generally small number of waveforms available for high-quality training datasets, allowing the bulk of the learning to be done on synthetics prior to retraining.

3.7 Stacking as a companion to denoising

The most commonly used method in deep Earth seismology to improve SNR, after spectral filtering, is stacking (see Rost & Thomas, 2002, for a comprehensive overview of array methods including stacking). Where there are several seismic stations in a region, known as an array, it is possible to spatially average the signals from these stations (Rost & Thomas, 2002, 2009; Schweitzer et al., 2012). Arrays may be deployed specifically for this purpose, or be created ad-hoc by grouping nearby stations together. By making an assumption about the directionality of the incoming energy, usually that it is a plane wave with specified slowness and back-azimuth, one can correct for the travel time difference between stations, a process known as beamforming, and then average the shifted seismic traces in the time domain. This leads to the suppression of incoherent energy (i.e. noise) and the amplification of coherent energy (i.e. the incoming teleseismic plane wave). Stacking functions as a slowness-back-azimuth filter, allowing the extraction of signal in the same period band as pervasive noise, provided the directional properties of the signal and noise differ.

There are a plethora of different averaging methods of varying complexity that can be implemented following beamforming. Examples include linear stacking (Davies et al., 1971), n^{th}-root stacking (Kanasewich et al., 1973), phase-weighted stacking (Schimmel & Paulssen, 1997), F-trace stacking (Blandford, 1974), and iterative stacking schemes (Hansen et al., 2021). Each method of stacking has its own advantages and disadvantages, including varying efficacy of noise suppression, waveform distortion and amplitude fidelity. Linear stacking is conceptually the easiest to implement, most commonly used, and has the advantage that it preserves waveform timing, shape and amplitude (Rost & Thomas, 2002).

The slowness and back-azimuth of the target signal must be known in order to stack optimally, otherwise coherent signals with differing directional properties will be suppressed; for example in our application, stacking for S_{diff} would suppress S_{diff}+ as it arrives from a different direction. For anomalous signals, the directional information may not be known in advance, and so one can iterate through different slowness and back-azimuth values, taking the highest amplitude stack as the truth. This is a slow iterative process but can identify the direction and slowness of incoming signal. It is by this method that ULVZs can be located from the direction of the post-cursory energy (e.g. Cottaar & Romanowicz, 2012; Martin et al., 2023b; Jagt et al., 2024).

Here, we test how the DAE performs relative to linear stacking, and also demonstrate stacking as a useful companion to denoising. We also verify that the directions recovered from beamforming S_{diff}+ waveforms are consistent in both original and denoised waveforms. We use waveforms from Event 20100320 recorded in the contiguous US, which have high-amplitude post-cursors generated by the Hawaiian ULVZ; the post-cursors of this event have been well-examined by previous studies (Cottaar & Romanowicz, 2012; Lai et al., 2022; Z. Li et al., 2022; J. Li et al., 2022; Martin et al., 2023a; Jagt et al., 2024). For every station we linearly stack all stations within a 2° (∼220 km) radius. As both S_{diff} and S_{diff}+ diffract along the CMB, we fix the stacking slowness to 8.36°/s, which is the critical slowness for the CMB in PREM. As S_{diff} and S_{diff}+ arrive from different directions, we create stacks separately for each phase; for S_{diff} we fix the back-azimuth to the great circle path, and for S_{diff}+ we allow back-azimuth to vary from $-25°$ to 0° relative to the great circle path in 1° intervals, taking the stack with the highest amplitude.

We use a DAE that has been retrained to Instaseis synthetics filtered to the specific period band, and results for the 10.2 - 20.4 s and 3.0 - 6.0 s filter-bands are shown in Figure 10, the former filter-band being close to that used by Cottaar & Romanowicz (2012) and the latter to that used by Z. Li et al. (2022). In panel a, the standard post-cursor delayed by 20 - 40 s is visible, and in panel g the short-period post-cursor delayed by 40 - 80 s is visible. In both filter-bands, both stacking and denoising increase the SNR and improve the visibility of the S_{diff}+, but the best results are achieved when both techniques are used together by stacking the denoised data.

The back-azimuths measured from the original and denoised data are consistent, demonstrating that array methods can be applied effectively to denoised data while preserving the directional information. Furthermore, as is apparent in the 3.0 - 6.0 s filter-band, comparing panels i and j, stacking can help to suppress ringing and erroneous signal present in the denoised data. These results suggest that combining these two complimentary techniques can lead to more effective noise suppression than each technique used in isolation.

4 Example applications to known ULVZs

4.1 Short-period signals

The short-period post-cursor used by Z. Li et al. (2022) to infer a greater velocity reduction at the base of the Hawaiian ULVZ is a consistent and coherent feature in the short-period denoised data shown in Figure 10 panels g - j, giving confidence that this is a robust signal. We highlight two further signals that are especially apparent in the denoised waveforms in panel j: there is a high frequency pre-cursor, marked in purple, and a second signal between S_{diff} and S_{diff}+, marked in red. Both of these signals have similar move-out to the post-cursors, i.e. arriving out of plane, indicating that they orig-

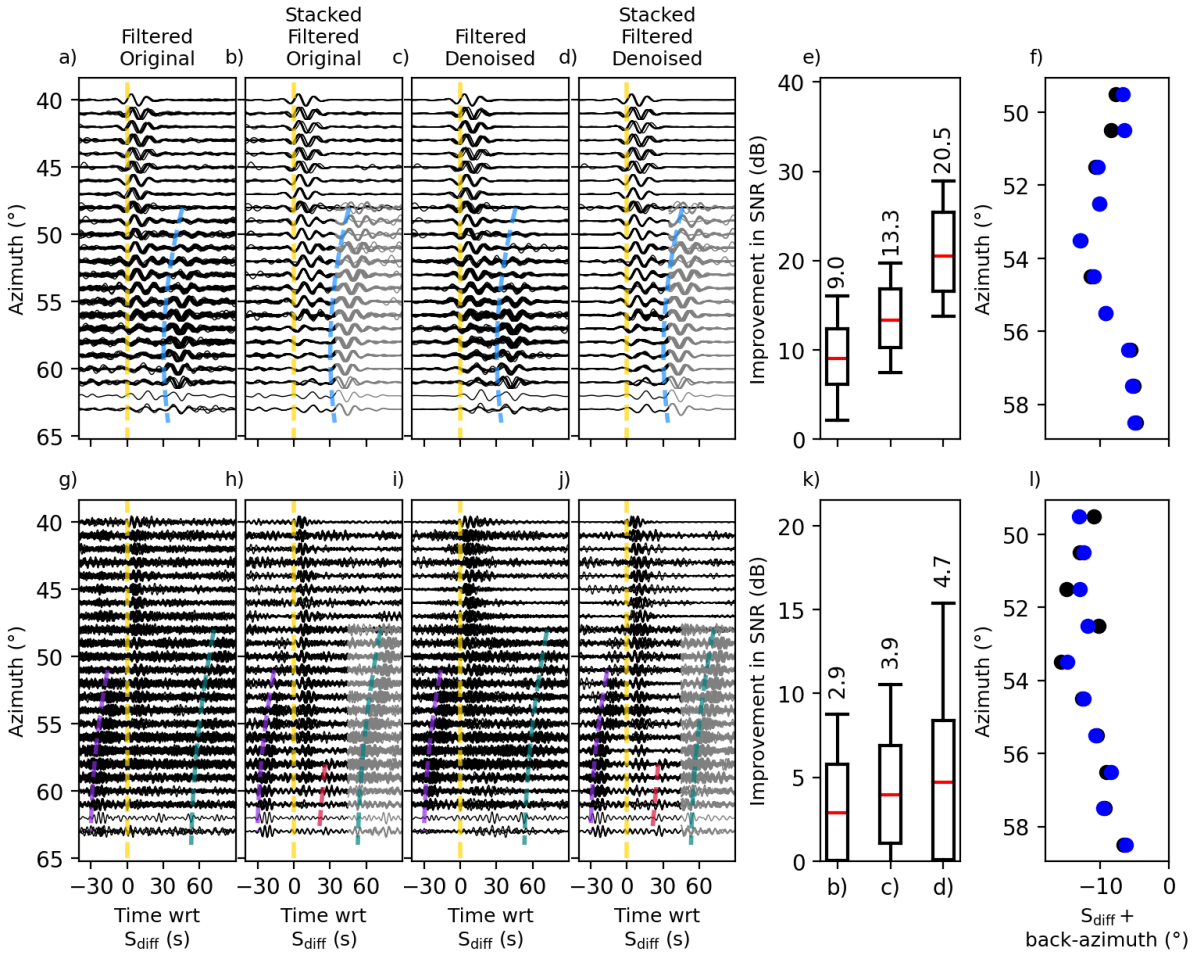

Figure 10. Comparison of the DAE with stacking for Event 20100320 recorded in the contiguous US filtered 10.2 - 20.4 s (top row) and 3.0 - 6.0 s (bottom row). a) The original data, b) the stacked data – the separate stacks for $S_{diff}+$ are shown in grey where calculated, c) the denoised data, d) the stacked denoised data. Traces and stacks in panels a) - d) are individually normalised such that the maximum amplitude corresponds to 0.5° azimuth, and are plotted rounded to the nearest degree of azimuth. Coherent signals are marked as dashed coloured lines to guide the eye, including the 10 - 20 s $S_{diff}+$ (Cottaar & Romanowicz, 2012) in blue and the short-period $S_{diff}+$ (Z. Li et al., 2022) in teal. e) The improvement in SNR for the waveforms in panels b) - d) relative to a). The median is shown in red, the edges of the boxes are the 25th and 75th percentile, and the ends of the arms are the 10th and 90th percentile. The value of the median is given above. f) The mean recovered back-azimuths of the greatest amplitude stacks for $S_{diff}+$ in 1° azimuth bins, for the original (black) and denoised (blue) data. Panels g) - l) are as for a) - f) but filtered at 3.0 - 6.0 s.

inate from 3D structure, likely the ULVZ. To the best of our knowledge, neither of these signals has been identified by previous studies of ULVZs, but are clear and coherent when denoising and stacking are applied.

We apply this method of denoising and stacking to several events that sample known ULVZs. For two events, 2010080407 and 20100417, that sample the Hawaiian ULVZ and for two further events, 20161225 and 20121207, that sample the Galápagos and Pitcairn ULVZs, respectively. The resulting resulting waveforms are shown in Figure 11. Both events that sample the Hawaiian ULVZ show clear pre-cursors, while those that sample the Galápagos and Pitcairn ULVZs also show possible pre-cursors. Some of these signals are still poor quality and correspondingly are marked with question marks to denote the uncertainty.

In many of these examples, the short-period signals are not visible in the original or stacked data, showing that denoising can make a material difference for identifying low-amplitude high-frequency signals from ULVZs. The cause of these out-of-plane signals – which have not been identified in the literature – is still unknown, but we suggest the the precursor, as a feature across several events, is likely scattered energy related to the ULVZ or nearby structure. These signals warrant further investigation.

4.2 Signals of the Hawaiian ULVZ recorded in Japan

The Hawaiian ULVZ is the most studied ULVZ, as well as the largest, but there is still much that is unknown. In particular, the shape of the ULVZ, which is usually assumed to be cylindrical for modelling purposes, remains poorly constrained. In large part this is due to the over-reliance on source-receiver geometries from the south-western Pacific, where there are many earthquakes, to North America, where there are many seismometers. Due to trade-offs along the S_{diff} diffracted path, it is not possible to uniquely constrain the size and shape of a ULVZ from unidirectional S_{diff} observations, and almost all previous literature using S_{diff} has focussed on this geometry (Cottaar & Romanowicz, 2012; Lai et al., 2022; Z. Li et al., 2022; J. Li et al., 2022; Martin et al., 2023a).

Martin et al. (2023a) compiled a dataset of almost 100 events that sample the Hawaiian ULVZ, providing unprecedented data coverage. They also noted the possible source-receiver geometry from events on the southern East Pacific Rise to Japan – where there is a large network of broadband seismometers, F-net (Okada et al., 2004) – but did not unambiguously identify $S_{diff}+$ signals from these events. A single event, 20090917, was marked as possibly containing $S_{diff}+$, but the low data quality precluded certainty. No data was found from stations in Japan that could be used to provide a complimentary dataset to $S_{diff}+$ recorded in North America. A recent global compilation of $S_{diff}+$ also contained no observations of the Hawaiian ULVZ from stations in Japan (Atkins et al., 2026).

As a final application, we apply a DAE to S_{diff} record sections recorded on F-net seismometers in Japan originating from earthquakes on the East Pacific Rise. To maximise possible detection we use a DAE that is originally trained on GSN noise and Instaseis synthetics, that has been retrained with real tapered S_{diff} recorded in North America. We examine the signals a function effective azimuth, defined as the difference between the source-to-ULVZ azimuth, and the ULVZ-to-station azimuth; effective azimuth is therefore a measure of expected refraction angle with respect to the ULVZ (Martin et al., 2023a). If the location of the ULVZ is correct, the $S_{diff}+$ signals are symmetric about 0° effective azimuth; we use the preferred centre location of Z. Li et al. (2022), 15.4°N and 172.3°W.

Using our DAE-assisted workflow, we identify four events containing $S_{diff}+$ signals (Figure 12). These post-cursors are different in character to those caused by the Hawaiian ULVZ as recorded in North America. The move-out parabolas are shallower, and $S_{diff}+$ is not well separated from the main S_{diff}. These $S_{diff}+$ arrivals are, however, visible over a relatively wide azimuthal range. We note that, judging by eye, the centre of the parabola typically centres between $-5°$ and $-10°$ effective azimuth, indicating that the best-fitting ULVZ location for the data in North America does not well-fit the data in Japan. Our new observations appear to be consistent with a more south-westerly location and elliptical structure, as suggested by Martin et al. (2023a).

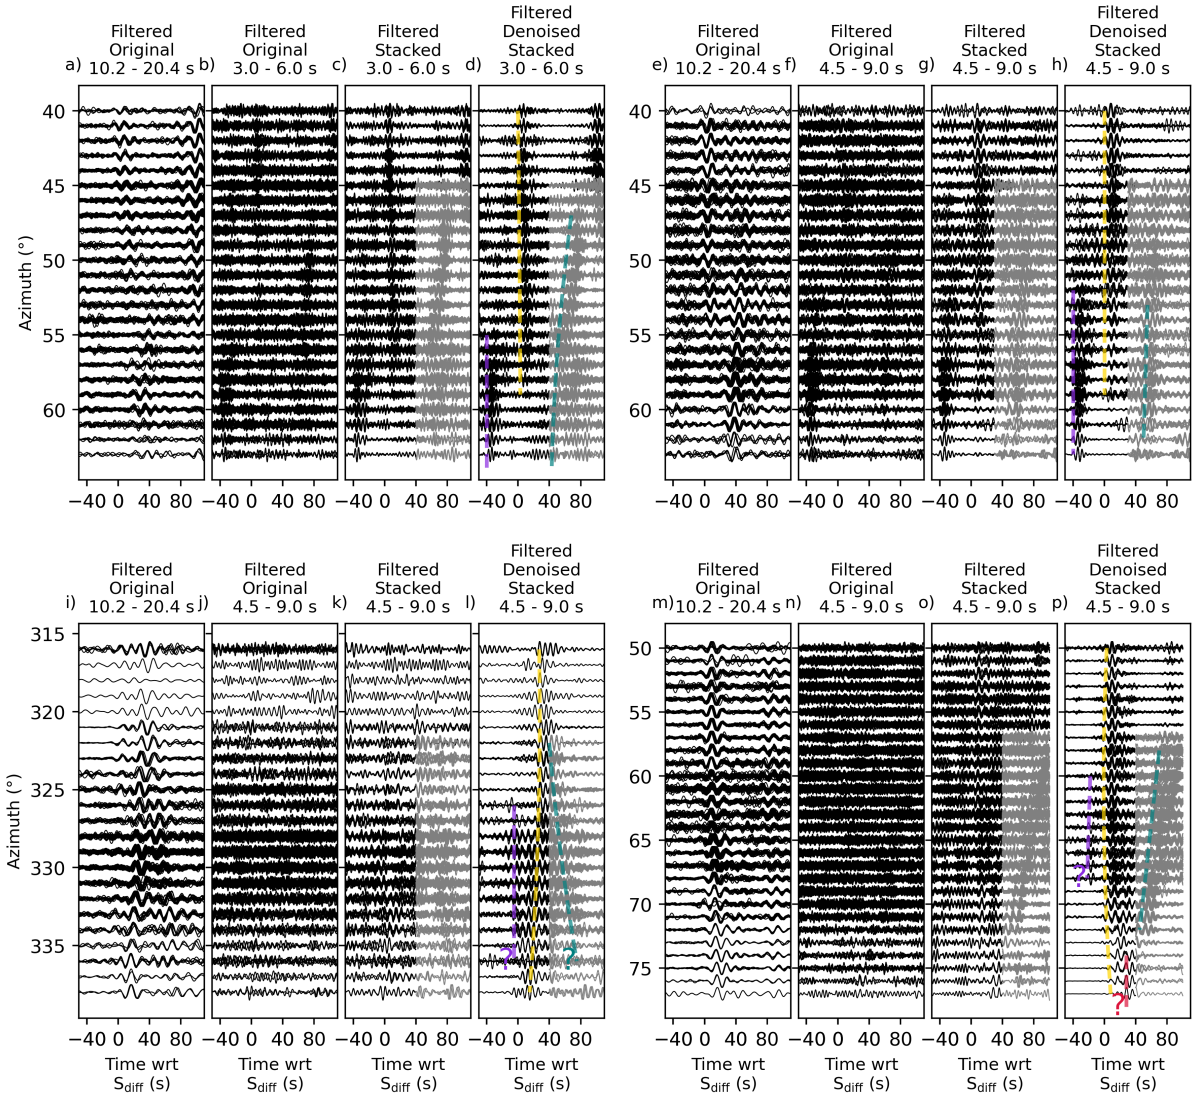

Figure 11. Examples of short-period signals elucidated by the DAE. Event 20100804 sampling the Hawaiian ULVZ filtered 3.0 - 6.0 s (panels a - d), Event 20100417 sampling the Hawaiian ULVZ filtered 4.5 - 9.0 s (panels e - h), Event 20161225 sampling the Galápagos ULVZ filtered 3.0 - 6.0 s (panels i - l), and Event 20121207 sampling the Pitcairn ULVZ filtered 2.0 - 4.0 s (panels m - p). For all events the first column is the original data filtered at longer period, the second is the original data filtered at short period, the third is short-period stacked data, and the fourth is short-period stacked denoised data. Traces where $S_{diff}+$ is expected is stacked separately for this time window and plotted in grey. Coherent signals are marked in the third column where visible; tentative signals are marked with a question mark.

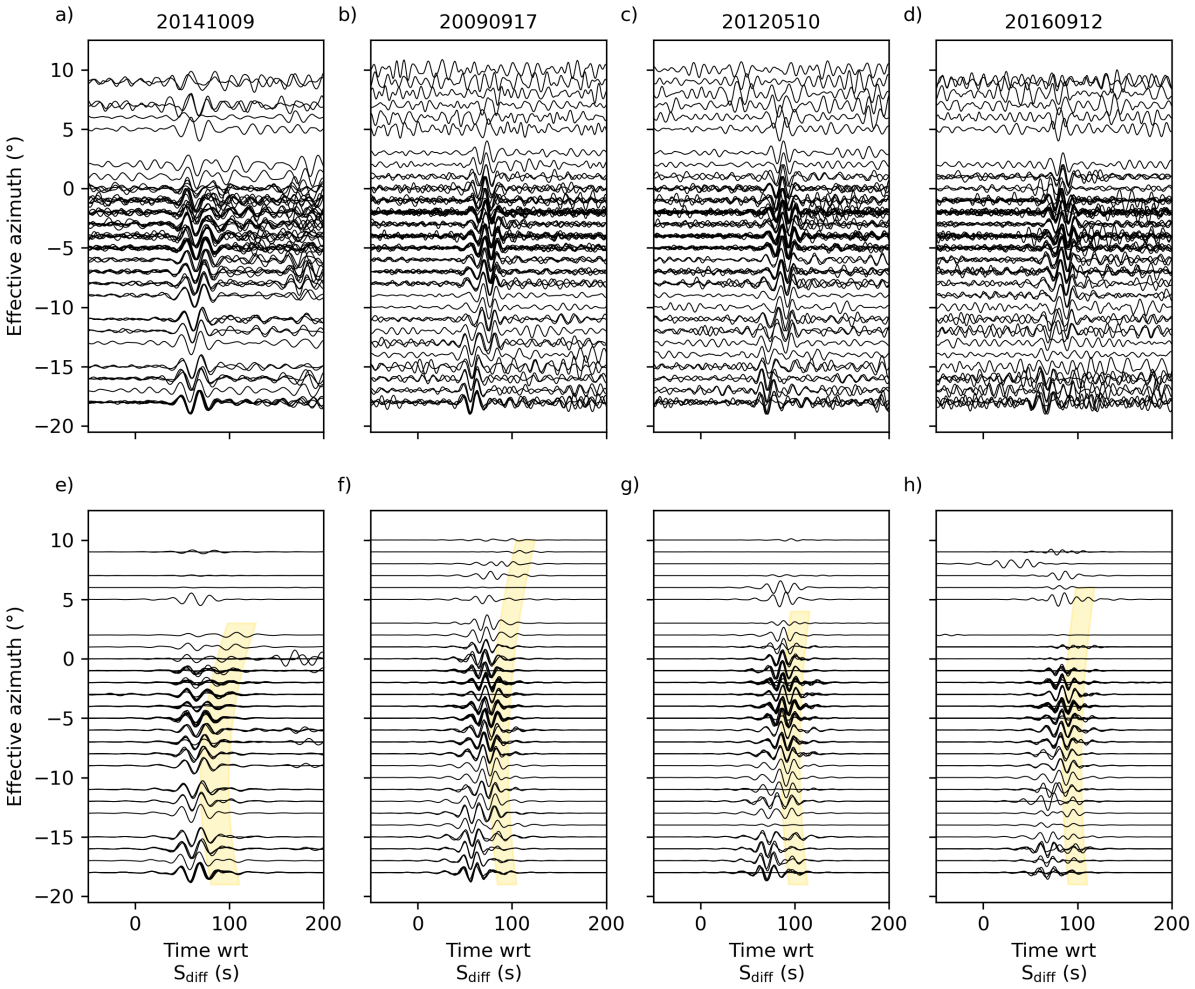

Figure 12. Examples of S_{diff}+ resulting from the Hawaiian ULVZ recorded by F-net station in Japan. a) - d) shows the filtered data normalised such that the maximum amplitude corresponds to 0.75° azimuth, and traces are plotted rounded to the nearest degree of azimuth. e) - h) shows the denoised data with amplitudes consistent with the top row. Potential post-cursors are highlighted in yellow to guide the eye. All panels are filtered at 10.2 - 20.4 s, except for panels a) and e), which are filtered at 15.2 - 30.5 s. The significant time delay of these arrivals relative to PREM is due to the Pacific LLVP.

Some of these S_{diff}+ signals are visible in the original data, but are difficult to identify due to the high noise level in Japan. Denoising highlights these signals and demonstrates its utility to deep Earth studies. These signals require further investigation to ascertain exactly why the manifestation of the Hawaiian ULVZ in Japan is so different to in North America.

5 Discussion

5.1 Reliability of the DAE

There are several points about the workflow that give us high confidence in the signals extracted by the DAE: (1) Traces are denoised one-at-a-time and the DAE has no spatial or temporal knowledge of the stations or traces. This means that it is not possible for the DAE to introduce coherent signal between traces unless there is coherent energy already present in those traces. While there may be uncertainty in any particular signal extracted in a single trace, there is high confidence in signals that are coherent across stations as a function of distance and/or azimuth, even when they are low amplitude, as they have been extracted independently. (2) The mask functions have values between 0 and 1, and so any particular point in the resulting time-frequency images can only be as high in amplitude as the same point in the original time-frequency representation. In other words, seismic energy cannot be created where it does not already exist. This does not mean that all extracted signal is correct. For instance, erroneous extracted signals are apparent in some of the figures shown in this paper, but these are typically low amplitude and are always incoherent across the record sections. Stacking the denoised traces provides an effective and complementary way of suppressing these incoherent erroneous signals. (3) As the same mask functions are applied to both the real and imaginary parts of the STFT, there is no phase shift introduced between the input and output traces. The potential for waveform distortion is therefore minimal and important time-axis observables, such as the arrival times of seismic phases, are well-preserved.

There are, of course, still warnings and caveats that must be considered when using the DAE. Most importantly, as we have noted, amplitude fidelity is a major limitation and amplitude loss may not be consistent across a trace or set of arrivals. Making amplitude measurements, even relative ones, on denoised traces is therefore to be discouraged. Despite the significant advantages of using mask functions outlined above, the mask functions themselves are a source of imperfect partitioning of signal and noise, as noted by previous studies (Zhu et al., 2019; Heuel & Friederich, 2022). It is possible for seismic energy to be incorrectly partitioned and for signals to be missed entirely, especially where the application data are unlike the training data. A lack of extracted signal should not be taken as a true lack of signal, simply a lack of signal that is similar enough to the training data to be extracted. We suggest retraining with more representative signals if this is suspected.

Finally, we note that the DAE is best at identifying signals when the noise level is already low – this is unsurprising as clearly clean signals are more trivial to identify, even for a deep learning algorithm. In cases where the initial noise level is low, all metrics including amplitudes and times are better preserved and these degrade as the initial noise level increases. Whilst the DAE is therefore expected to perform better at 'tidying up' signal that is only partly obscured by noise, our results clearly demonstrate that it is also successful at revealing hidden signals that are completely obscured by noise. For screening purposes, where the aim to to elucidate the possible presence of signal, the DAE would be highly appropriate in spite of sub-optimal extraction at high noise levels.

5.2 Potential deep Earth use cases

We have used S_{diff} and its ULVZ-generated post-cursor, S_{diff}+, as a case study to demonstrate the effectiveness and potential of the DAE in improving data quality for deep Earth body wave studies. The benefits to denoising these phases are readily apparent in helping to better understand ULVZs, for which many important properties remain poorly constrained. There are a wide array of other possible applications of deep-learning based denoising in deep Earth seismology, although due to the loss of amplitudes, which may not be consistent between phases, the DAE would not be appropriate

for studies that intend to use amplitudes. We envision the DAE being most useful for the initial screening and identification of seismic phases that can then be the subject of further detailed investigation.

The DAE has proven particularly effective at removing ocean-generated microseism and so would be expected to be useful for stations with high oceanic noise levels such as the Caribbean, ocean islands, or OBS deployments, where teleseismic observations are challenging – as demonstrated by strong performance on data from Japan. Furthermore, by transfer learning, the DAE can be efficiently optimised for use in places with a unique noise environment such as Antarctica (Anthony et al., 2015), a region that is increasingly being used for deep Earth seismic research (e.g. Brett & Deuss, 2020; Hansen et al., 2023).

Increasing global data coverage in this manner is highly desirable. For the ULVZs that have been the focus of this study, better global sampling could allow imaging of ULVZs from new directions, as we have demonstrated using Japanese F-net data, allowing better constraints on size and morphology, as well as the discovery of ULVZs in new locations. This would help to better map global distributions (e.g. Yu & Garnero, 2018; Thorne et al., 2021; Atkins et al., 2026). Many other deep Earth applications suffer from geographic bias and, most notably, this has a major impact on tomographic models (e.g. Foulger et al., 2013; Ruan et al., 2019). As the DAE preserves travel time information, there is clear potential for the DAE to be used in the construction of tomographic models using phase-based metrics.

The study of Earth's inner core is another area where denoising may prove useful; since the inner core is highly attenuating and body wave observations are often low amplitude (e.g. Cormier et al., 2021), and adequate and representative data distribution is key (e.g. Frost et al., 2021; Waszek et al., 2023). Constraining the magnitude and directionality of P-wave anisotropy is essential; in particular, further robust observations of polar geometries are required (e.g. Brett & Deuss, 2020), but such measurements are often made from stations in noisy high-latitude environments with low station density. The innermost inner core, a potentially distinct region at the very centre of the planet, is the subject of intense research to unambiguously identify and constrain its properties (e.g. Lythgoe et al., 2014; Stephenson et al., 2021; Brett et al., 2022; Costa de Lima et al., 2022). This task would be helped by observations of deeper-bottoming phases, which has been the subject of recent efforts with machine learning (Zhou et al., 2024).

5.3 Possible future developments

Based on the criteria that we have set for downloading GSN noise, we have only used $\sim$0.12% of the $>$70 million hours of available training noise. This is only for the GSN – incorporating other networks could massively increase the available training noise again. The training synthetics are also limited; for example, the Instaseis synthetics are an idealised and unrealistic signal prior. However, creating a diverse training dataset of 3D synthetics, by incorporating ULVZs, tomographic or crustal models, would involve significant computational cost. An alternative to running costly 3D simulations is to convolve 1D synthetics with crustal responses (Zhou et al., 2024). Other S phases or P waves could also be incorporated to diversify the training data. The potential to develop better performing models, based on more numerous and varied training data, is apparent, but performance scaling has not been quantified.

The learning process could also be modified. At present, the classification of each point of the input STFT into signal and noise is treated as a binary classification task and performance is measured by binary cross-entropy on the masks themselves. Implementing multi-resolution STFT inputs or performance metrics may better characterise frequency content, allowing more effective discrimination (e.g. Engel et al., 2020; Yamamoto et al., 2020). Reformulation into a regression task on the masks or measuring performance on the traces in the time domain are alternative options. Outputting signal and noise directly is another possibility, but this would forfeit the benefits of using mask functions.

One could also modify the architecture of the DAE to incorporate array information. By integrating temporo-spatial information, as stacking does, we hypothesise that even greater noise suppression and signal extraction could be achieved. But this potential advantage would come with a trade-off: the present trace-by-trace denoising implementation provides an important safety check, which would be forfeited if the DAE was able to extract spurious signal coherently.

While the denoisers already perform well in our examples, the models used in this study represent only an initial demonstration of the potential of denoising teleseismic data with supervised deep-learning. Alternative learning objectives and model architectures remain to be investigated, and the performance demonstrated in this paper represents a lower bound on the ability of the DAE for this task.

6 Conclusions and recommendations

In this study, we have demonstrated the potential of the DAE for improving teleseismic data quality for studying the deep Earth. We have presented a computationally efficient workflow that utilises only Instaseis synthetics and global seismic noise for training, that manages to extract ULVZ-associated signals from real seismic data. We have shown that denoising outperforms traditional spectral filtering and can be used successfully alongside stacking to elucidate signals from ULVZs. We have also established that the noise from the GSN provides sufficient global noise exposure for the DAE to work in the regions that we have tested without further optimisation, but that specification to particular types of seismic signal can enhance extraction. Therefore, when seeking the extraction of highly specific seismic signals, retraining via transfer learning can boost performance, including utilising small datasets of real data.

We have focussed on denoising record sections that have visibly coherent seismic signals between them, which provides a practical check on the outputs of the DAE: extracted signals that are coherent between stations can be more confidently interpreted as genuine signals. While denoising and interpreting signals in single stations is possible, it forfeits this simple and effective check and we advise caution when doing so. Amplitude fidelity is a major limitation, and amplitude measurements, even relative ones, should not be made on the outputs of the DAE, particularly where the application is dissimilar to the training data. Conversely the extraction of travel times is typically robust. Performance by all metrics is reduced at high noise levels, but the DAE shows remarkable aptitude for elucidating signals that are completely obscured by noise.

We have applied denoising in two applications, revealing new short-period observations likely related to ULVZs, and identifying the signal of the Hawaiian ULVZ using a novel source-receiver geometry from the East Pacific Rise to Japan. The results and applications suggest that denoising is a powerful tool that can compliment existing methodologies for processing and visualising deep Earth seismic data. The potential for denoising to impact deep Earth seismology is apparent, and we hope that by highlighting its capability and providing clear steps and recommendations for implementation, we will encourage the use of denoising in future studies.

Acknowledgements

SR acknowledges funding from the Deutsche Forschungsgemeinschaft (DFG), grant number TH1530/16-3. WE acknowledges funding from the German Federal Ministry of Research, Technology and Space (BMFTR), ErUM-WAVE project 05D2022. This study has made use of the computing cluster PALMA II of Universität Münster, subsidised by the DFG (INST 211/667-1). We would like to thank Maggie Chen at the University of Oxford and Sanne Cottaar at the University of Cambridge for valuable conversations relevant to this work.

Data Availability

All seismic data used in this study is freely available online, either through the National Research Institute for Earth Science and Disaster Resilience (NIED; https://www.hinet.bosai.go.jp/) for Japanese F-net waveform data, or the U.S. National Science Foundation's Seismological Facility for the Advancement of Geoscience (NSF SAGE; https://www.iris.edu/hq/) for all other waveform data. Data are available for convenient download using Python by using the ObsPy https://pypi.org/project/obspy/ and HinetPy https://pypi.org/project/HinetPy/ packages. The codes for training data compilation, training, retraining and denoising used in this study are freely available in a public GitHub repository https://github.com/StuartJRussell/DeepSeisDAE/ under a GNU General Public License. A DAE is also implemented in SeisBench https://seisbench.readthedocs.io/en/stable/.

Competing interests

The authors declare that there are no conflicts of interest for this manuscript.

References

Anthony, R. E., Aster, R. C., Wiens, D., Nyblade, A., Anandakrishnan, S., Huerta, A., ... Rowe, C. (2015). The seismic noise environment of Antarctica. *Seismological Research Letters*, *86*(1), 89–100. doi: 10.1785/0220140109

Atkins, J. R., Martin, C., & Cottaar, S. (2026). Global presence and absence of ultra-low velocity zones as seen by Sdiff postcursors. *Journal of Geophysical Research: Solid Earth*, *131*(4), e2025JB033171. doi: 10.1029/2025JB033171

Beyreuther, M., Barsch, R., Krischer, L., Megies, T., Behr, Y., & Wassermann, J. (2010). ObsPy: A Python toolbox for seismology. *Seismological Research Letters*, *81*(3), 530–533. doi: 10.1785/gssrl.81.3.530

Blandford, R. (1974). An automatic event detector at the Tonto Forest seismic observatory. *Geophysics*, *39*(5), 633–643. doi: 10.1190/1.1440453

Bormann, P., & Wielandt, E. (2013). Seismic signals and noise. In *New manual of seismological observatory practice 2 (nmsop2)* (pp. 1–62). Deutsches GeoForschungsZentrum GFZ.

Brett, H., & Deuss, A. (2020). Inner core anisotropy measured using new ultra-polar PKIKP paths. *Geophysical Journal International*, *223*(2), 1230–1246. doi: 10.1093/gji/ggaa348

Brett, H., Hawkins, R., Waszek, L., Lythgoe, K., & Deuss, A. (2022). 3D transdimensional seismic tomography of the inner core. *Earth and Planetary Science Letters*, *593*, 117688. doi: 10.1016/j.epsl.2022.117688

Bridle, J. S. (1990). Probabilistic interpretation of feedforward classification network outputs, with relationships to statistical pattern recognition. In *Neurocomputing: Algorithms, architectures and applications* (pp. 227–236). Springer. doi: 10.1007/978-3-642-76153-9_28

Chaljub, E., Komatitsch, D., Vilotte, J.-P., Capdeville, Y., Valette, B., & Festa, G. (2007). Spectral-element analysis in seismology. *Advances in geophysics*, *48*, 365–419. doi: 10.1016/S0065-2687(06)48007-9

Cormier, V. F., Bergman, M. I., & Olson, P. L. (2021). *Earth's core: Geophysics of a planet's deepest interior*. Elsevier.

Costa de Lima, T., Tkalčić, H., & Waszek, L. (2022). A new probe into the innermost inner core anisotropy via the global coda-correlation wavefield. *Journal of Geophysical Research: Solid Earth*, *127*(4), e2021JB023540. doi: 10.1029/2021JB023540

Cottaar, S., Martin, C., Li, Z., & Parai, R. (2022). The root to the Galápagos mantle plume on the core-mantle boundary. *Seismica*, *1*(1), 197. doi: 10.26443/seismica.v1i1.197

Cottaar, S., & Romanowicz, B. (2012). An unsually large ULVZ at the base of the mantle near Hawaii. *Earth and Planetary Science Letters*, *355*, 213–222. doi: 10.1016/j.epsl.2012.09.005

Dahmen, N., Clinton, J., Stähler, S., Meier, M.-A., Ceylan, S., Euchner, F., ... others (2024). Re-

visiting Martian seismicity with deep learning-based denoising. *Geophysical Journal International*, *239*(1), 434–454. doi: 10.1093/gji/ggae279

Dahmen, N., Clinton, J. F., Meier, M.-A., Stähler, S. C., Ceylan, S., Kim, D., ... Giardini, D. (2022). MarsQuakeNet: A more complete marsquake catalog obtained by deep learning techniques. *Journal of Geophysical Research: Planets*, *127*(11), e2022JE007503. doi: 10.1029/2022JE007503

Dalai, B., Kumar, P., Srinu, U., & Sen, M. K. (2022). De-noising receiver function data using the unsupervised deep learning approach. *Geophysical Journal International*, *229*(2), 737–749. doi: 10.1093/gji/ggab494

Darbyshire, J. (1950). Identification of microseismic activity with sea waves. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, *202*(1070), 439–448. doi: 10.1098/rspa.1950.0111

Davies, D., Kelly, E., & Filson, J. (1971). Vespa process for analysis of seismic signals. *Nature Physical Science*, *232*(27), 8–13. doi: 10.1038/physci232008a0

Dziewonski, A. M., & Anderson, D. L. (1981). Preliminary reference earth model. *Physics of the earth and planetary interiors*, *25*(4), 297–356. doi: 10.1016/0031-9201(81)90046-7

Engel, J., Gu, C., & Roberts, A. (2020). DDSP: Differentiable digital signal processing. In *International conference on learning representations (iclr)*.

Feng, M., Chen, L., Wei, S., Muksin, U., Simanjuntak, A. V., Chen, Y., & Gong, C. (2024). Deep Learning–Based Denoising Improves Receiver Function Imaging Using Dense Short-Period Teleseismic Data. *Seismological Research Letters*, *95*(6), 3696–3708. doi: 10.1785/0220240017

Foulger, G. R., Panza, G. F., Artemieva, I. M., Bastow, I. D., Cammarano, F., Evans, J. R., ... others (2013). Caveats on tomographic images. *Terra Nova*, *25*(4), 259–281. doi: 10.1111/ter.12041

Friedrich, A., Krüger, F., & Klinge, K. (1998). Ocean-generated microseismic noise located with the Gräfenberg array. *Journal of seismology*, *2*(1), 47–64. doi: 10.1023/A:1009788904007

Frost, D. A., Lasbleis, M., Chandler, B., & Romanowicz, B. (2021). Dynamic history of the inner core constrained by seismic anisotropy. *Nature Geoscience*, *14*(7), 531–535. doi: 10.1038/s41561-021-00761-w

Gaci, S. (2013). The use of wavelet-based denoising techniques to enhance the first-arrival picking on seismic traces. *IEEE Transactions on Geoscience and Remote Sensing*, *52*(8), 4558–4563. doi: 10.1109/TGRS.2013.2282422

Ge, W., Mao, Q., Zhou, W., Gui, Z., & Wang, P. (2024). Signal enhancement for downhole microseismic data using improved attention mechanism based on autoencoder network. *IEEE Access*. doi: 10.1109/ACCESS.2024.3483196

Han, J., & van der Baan, M. (2015). Microseismic and seismic denoising via ensemble empirical mode decomposition and adaptive thresholding. *Geophysics*, *80*(6), KS69–KS80. doi: 10.1190/geo2014-0423.1

Hansen, S. E., Garnero, E. J., Li, M., Shim, S.-H., & Rost, S. (2023). Globally distributed subducted materials along the Earth's core-mantle boundary: Implications for ultralow velocity zones. *Science Advances*, *9*(14), eadd4838. doi: 10.1126/sciadv.add4838

Hansen, S. E., Garnero, E. J., & Rost, S. (2021). Historical interstation pattern referencing (HIPR): An application to PcP waves recorded in the Antarctic for ULVZ imaging. *Journal of Geophysical Research: Solid Earth*, *126*(10), e2021JB022741. doi: 10.1029/2021JB022741

Heuel, J., & Friederich, W. (2022). Suppression of wind turbine noise from seismological data using nonlinear thresholding and denoising autoencoder. *Journal of Seismology*, *26*(5), 913–934. doi: 10.1007/s10950-022-10097-6

Hutko, A. R., Bahavar, M., Trabant, C., Weekly, R. T., Fossen, M. V., & Ahern, T. (2017). Data products at the IRIS-DMC: Growth and usage. *Seismological Research Letters*, *88*(3), 892–903. doi: 10.1785/0220160190

Iman, M., Arabnia, H. R., & Rasheed, K. (2023). A review of deep transfer learning and recent advancements. *Technologies*, *11*(2), 40. doi: 10.3390/technologies11020040

Ioffe, S., & Szegedy, C. (2015). Batch normalization: Accelerating deep network training by re-

ducing internal covariate shift. In *International conference on machine learning* (pp. 448–456).
Iyer, H. (1958). A study on the direction of arrival of microseisms at kew observatory. *Geophysical Journal International*, *1*(1), 32–43. doi: 10.1111/j.1365-246X.1958.tb00032.x
Jackson, J. M., & Thomas, C. (2021). Seismic and mineral physics constraints on the d″ layer. *Mantle convection and surface expressions*, 193–227. doi: 10.1002/9781119528609.ch8
Jagt, L., Martin, C., Millet, F., Russell, S., & Cottaar, S. (2024). Pdiff Postcursors from the Base of the Hawaiian ULVZ. *The Seismic Record*, *4*(3), 204–213. doi: 10.1785/0320240016
Kanasewich, E., Alpaslan, T., & Hemmings, C. (1973). Nth-root stack nonlinear multichannel filter. *Geophysics*, *38*(2), 327–338. doi: 10.1190/1.1440343
Komatitsch, D. (2011). Fluid–solid coupling on a cluster of gpu graphics cards for seismic wave propagation. *Comptes Rendus Mecanique*, *339*(2-3), 125–135. doi: 10.1016/j.crme.2010.11.007
Komatitsch, D., & Tromp, J. (1999). Introduction to the spectral element method for three-dimensional seismic wave propagation. *Geophysical journal international*, *139*(3), 806–822. doi: 10.1046/j.1365-246x.1999.00967.x
Komatitsch, D., & Tromp, J. (2002). Spectral-element simulations of global seismic wave propagation—II. Three-dimensional models, oceans, rotation and self-gravitation. *Geophysical journal international*, *150*(1), 303–318. doi: 10.1046/j.1365-246X.2002.01716.x
Krischer, L., Hutko, A. R., van Driel, M., Stähler, S., Bahavar, M., Trabant, C., & Nissen-Meyer, T. (2017, 04). On-Demand Custom Broadband Synthetic Seismograms. *Seismological Research Letters*, *88*(4), 1127-1140. doi: 10.1785/0220160210
Lai, V. H., Helmberger, D. V., Dobrosavljevic, V. V., Wu, W., Sun, D., Jackson, J. M., & Gurnis, M. (2022). Strong ULVZ and slab interaction at the northeastern edge of the Pacific LLSVP favors plume generation. *Geochemistry, Geophysics, Geosystems*, *23*(2), e2021GC010020. doi: 10.1029/2021GC010020
Langston, C. A., & Mousavi, S. M. (2019). Separating signal from noise and from other signal using nonlinear thresholding and scale-time windowing of continuous wavelet transforms. *Bulletin of the Seismological Society of America*, *109*(5), 1691–1700. doi: 10.1785/0120190073
Leng, K., Nissen-Meyer, T., & van Driel, M. (2016). Efficient global wave propagation adapted to 3-D structural complexity: a pseudospectral/spectral-element approach. *Geophysical Supplements to the Monthly Notices of the Royal Astronomical Society*, *207*(3), 1700–1721. doi: 10.1093/gji/ggw363
Leng, K., Nissen-Meyer, T., Van Driel, M., Hosseini, K., & Al-Attar, D. (2019). AxiSEM3D: broad-band seismic wavefields in 3-D global earth models with undulating discontinuities. *Geophysical Journal International*, *217*(3), 2125–2146. doi: 10.1093/gji/ggz092
Li, J., Qu, R., & Lu, C. (2023). Multiple attention mechanisms-based convolutional neural network for desert seismic denoising. *Pure and Applied Geophysics*, *180*(6), 2135–2155. doi: 10.1007/s00024-023-03255-5
Li, J., Sun, D., & Bower, D. J. (2022). Slab control on the mega-sized north pacific ultra-low velocity zone. *Nature Communications*, *13*(1), 1042. doi: 10.1038/s41467-022-28708-8
Li, Z., Leng, K., Jenkins, J., & Cottaar, S. (2022). Kilometer-scale structure on the core–mantle boundary near Hawaii. *Nature Communications*, *13*(1), 2787. doi: 10.1038/s41467-022-30502-5
Li, Z., Martin, C., & Cottaar, S. (2024). Seismic observation of a new ULVZ beneath the Southern Pacific. *Journal of Geophysical Research: Solid Earth*, *129*(4), e2023JB026941. doi: 10.1029/2023JB026941
Lin, H., Xu, J., Xing, L., Li, Q., & Liu, H. (2024). Random noise attenuation of ocean bottom seismometers based on a substep deep denoising autoencoder. *Geophysical Prospecting*, *72*(Special Issue: Seabed Prospecting Technology), 1428–1441. doi: 10.1111/1365-2478.13302
Liu, J.-Y., & Yang, Y.-H. (2018). Denoising auto-encoder with recurrent skip connections and residual regression for music source separation. In *2018 17th ieee international conference on machine learning and applications (icmla)* (pp. 773–778). doi: 10.1109/ICMLA.2018.00123
Longuet-Higgins, M. S. (1950). A theory of the origin of microseisms. *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences*, *243*(857), 1–35.

doi: 10.1098/rsta.1950.0012
Lythgoe, K., Deuss, A., Rudge, J., & Neufeld, J. (2014). Earth's inner core: Innermost inner core or hemispherical variations? *Earth and Planetary Science Letters*, *385*, 181–189. doi: 10.1016/j.epsl.2013.10.049
Ma, Z., Li, H., Liu, M., Huang, Y., Zhang, S., & Lv, Y. (2023). Micro-seismic events detection and its tectonic implications in Northeastern Hainan Province. *Frontiers in Earth Science*, *11*, 1169877. doi: 10.3389/feart.2023.1169877
Martin, C., Bodin, T., & Cottaar, S. (2023a). Mapping structures on the core–mantle boundary using Sdiff postcursors: Part II. Application to the Hawaiian ULVZ. *Geophysical Journal International*, *235*(3), 2399–2409. doi: 10.1093/gji/ggad345
Martin, C., Bodin, T., & Cottaar, S. (2023b). Mapping structures on the core–mantle boundary using Sdiff postcursors: Part I. Method and Validation. *Geophysical Journal International*, *235*(3), 2385–2398. doi: 10.1093/gji/ggad340
Martin, C., Harmsma, L., Atkins, J., Deuss, A., & Cottaar, S. (2025). Evidence of a ULVZ near Vanuatu from Sdiff postcursors. *Journal of Studies of Earth's Deep Interior*, *1*(1). doi: 10.46298/jsedi.15883
Martin, C., Russell, S., & Cottaar, S. (2024). Detection of a ULVZ in the Central Pacific using high frequency Sdiff postcursors. *Earth and Planetary Science Letters*, *648*, 119028. doi: 10.1016/j.epsl.2024.119028
McNamara, A. K. (2019). A review of large low shear velocity provinces and ultra low velocity zones. *Tectonophysics*, *760*, 199–220. doi: 10.1016/j.tecto.2018.04.015
Mousavi, S. M., & Beroza, G. C. (2022). Deep-learning seismology. *Science*, *377*(6607), eabm4470. doi: 10.1126/science.abm4470
Mousavi, S. M., Langston, C. A., & Horton, S. P. (2016). Automatic microseismic denoising and onset detection using the synchrosqueezed continuous wavelet transform. *Geophysics*, *81*(4), V341–V355. doi: 10.1190/geo2015-0598.1
Nissen-Meyer, T., van Driel, M., Stähler, S. C., Hosseini, K., Hempel, S., Auer, L., ... Fournier, A. (2014). AxiSEM: broadband 3-D seismic wavefields in axisymmetric media. *Solid Earth*, *5*(1), 425–445. doi: 10.5194/se-5-425-2014
Oboué, Y. A. S. I., Ying, H., Ma, S., Zuo, P., & Chen, Y. (2026). Deep Learning-Based Recovery of Weak PcP Phases from Noisy Seismic Records. *Big Data and Earth System*, 100041. doi: 10.1016/j.bdes.2026.100041
Okada, Y., Kasahara, K., Hori, S., Obara, K., Sekiguchi, S., Fujiwara, H., & Yamamoto, A. (2004). Recent progress of seismic observation networks in japan—hi-net, f-net, k-net and kik-net—. *Earth, Planets and Space*, *56*(8), xv–xxviii. doi: 10.1186/BF03353076
Oldham, R. D. (1906). The constitution of the interior of the Earth, as revealed by earthquakes. *Quarterly Journal of the Geological Society*, *62*(1-4), 456–475.
Othman, A., Iqbal, N., Hanafy, S. M., & Waheed, U. B. (2021). Automated event detection and denoising method for passive seismic data using residual deep convolutional neural networks. *IEEE Transactions on Geoscience and Remote Sensing*, *60*, 1–11. doi: 10.1109/TGRS.2021.3054071
Ringler, A. T., Anthony, R. E., Aster, R., Ammon, C., Arrowsmith, S., Benz, H., ... others (2022). Achievements and prospects of global broadband seismographic networks after 30 years of continuous geophysical observations. *Reviews of Geophysics*, *60*(3), e2021RG000749. doi: 10.1029/2021RG000749
Ringler, A. T., Steim, J., Wilson, D. C., Widmer-Schnidrig, R., & Anthony, R. E. (2020). Improvements in seismic resolution and current limitations in the Global Seismographic Network. *Geophysical Journal International*, *220*(1), 508–521. doi: 10.1093/gji/ggz473
Romanowicz, B. A., & Dziewonski, A. M. (1987). Global digital seismographic network: Research opportunities and recent initiatives. *Composition, structure and dynamics of the lithosphere-asthenosphere system*, *16*, 99–110. doi: 10.1029/GD016p0099
Ronneberger, O., Fischer, P., & Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention–miccai*

2015: 18th international conference, munich, germany, october 5-9, 2015, proceedings, part iii 18 (pp. 234–241). doi: 10.1007/978-3-319-24574-4_28

Rost, S., & Thomas, C. (2002). Array seismology: Methods and applications. *Reviews of geophysics*, *40*(3), 2–1. doi: 10.1029/2000RG000100

Rost, S., & Thomas, C. (2009). Improving seismic resolution through array processing techniques. *Surveys in geophysics*, *30*, 271–299. doi: 10.1007/s10712-009-9070-6

Ruan, Y., Lei, W., Modrak, R., Örsvuran, R., Bozdağ, E., & Tromp, J. (2019). Balancing unevenly distributed data in seismic tomography: a global adjoint tomography example. *Geophysical Journal International*, *219*(2), 1225–1236. doi: 10.1093/gji/ggz356

Russell, S., Magali, J. K., Vallenton, K., & Thomas, C. (2025). The effect of source-side subduction on PKP differential times and implications for inner core anisotropy. *Physics of the Earth and Planetary Interiors*, 107382. doi: 10.1016/j.pepi.2025.107382

Schimmel, M., & Paulssen, H. (1997). Noise reduction and detection of weak, coherent signals through phase-weighted stacks. *Geophysical Journal International*, *130*(2), 497–505. doi: 10.1111/j.1365-246X.1997.tb05664.x

Schweitzer, J., Fyen, J., Mykkeltveit, S., Gibbons, S. J., Pirli, M., Kühn, D., & Kværna, T. (2012). Seismic arrays. In *New manual of seismological observatory practice 2 (nmsop-2)* (pp. 1–80). Deutsches GeoForschungsZentrum GFZ.

Shi, Q., & Denolle, M. A. (2023). Improved observations of deep earthquake ruptures using machine learning. *Journal of Geophysical Research: Solid Earth*, *128*(12), e2023JB027334. doi: 10.1029/2023JB027334

Shi, Q., Denolle, M. A., Ni, Y., Williams, E. F., & You, N. (2025). Denoising offshore distributed acoustic sensing using masked auto-encoders to enhance earthquake detection. *Journal of Geophysical Research: Solid Earth*, *130*(2), e2024JB029728. doi: 10.1029/2024JB029728

Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. (2014). Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, *15*(1), 1929–1958.

Steinberg, A., Gaebler, P., Hartmann, G., Lehr, J., & Pilger, C. (2024). Deep Neural Networks Based Denoising of Regional Seismic Waveforms and Impact on Analysis of North Korean Nuclear Tests. *Pure and Applied Geophysics*, 1–21. doi: 10.1007/s00024-024-03491-3

Stephenson, J., Tkalčić, H., & Sambridge, M. (2021). Evidence for the innermost inner core: robust parameter search for radially varying anisotropy using the neighborhood algorithm. *Journal of Geophysical Research: Solid Earth*, *126*(1), e2020JB020545. doi: 10.1029/2020JB020545

Thorne, M. S., Garnero, E. J., Jahnke, G., Igel, H., & McNamara, A. K. (2013). Mega ultra low velocity zone and mantle flow. *Earth and Planetary Science Letters*, *364*, 59–67. doi: 10.1016/j.epsl.2012.12.034

Thorne, M. S., Leng, K., Pachhai, S., Rost, S., Wicks, J., & Nissen-Meyer, T. (2021). The most parsimonious ultralow-velocity zone distribution from highly anomalous SPdKS waveforms. *Geochemistry, Geophysics, Geosystems*, *22*(1), e2020GC009467. doi: 10.1029/2020GC009467

Tian, D. (2024). HinetPy: A Python package for accessing and processing NIED Hi-net seismic data. *Journal of Open Source Software*, *9*(98), 6840. doi: 10.21105/joss.06840

Tibi, R., Hammond, P., Brogan, R., Young, C. J., & Koper, K. (2021). Deep learning denoising applied to regional distance seismic data in Utah. *Bulletin of the Seismological Society of America*, *111*(2), 775–790. doi: 10.1785/0120200292

Tibi, R., Young, C. J., & Porritt, R. W. (2023). Comparative study of the performance of seismic waveform denoising methods using local and near-regional data. *Bulletin of the Seismological Society of America*, *113*(2), 548–561. doi: 10.1785/0120220105

van Driel, M., Kemper, J., & Boehm, C. (2021). On the modelling of self-gravitation for full 3-d global seismic wave propagation. *Geophysical Journal International*, *227*(1), 632–643. doi: 10.1093/gji/ggab237

van Driel, M., Krischer, L., Stähler, S. C., Hosseini, K., & Nissen-Meyer, T. (2015). Instaseis: instant global seismograms based on a broadband waveform database. *Solid Earth*, *6*(2), 701–717. doi: 10.5194/se-6-701-2015,2015.

Viens, L., & Van Houtte, C. (2020). Denoising ambient seismic field correlation functions with convolutional autoencoders. *Geophysical Journal International*, *220*(3), 1521–1535. doi: 10.1093/gji/ggz509

Waszek, L., Irving, J., Phạm, T.-S., & Tkalčić, H. (2023). Seismic insights into Earth's core. *nature communications*, *14*(1), 6029. doi: 10.1038/s41467-023-41725-5

Weiss, K., Khoshgoftaar, T. M., & Wang, D. (2016). A survey of transfer learning. *Journal of Big data*, *3*(1), 9. doi: 10.1186/s40537-016-0043-6

Wilson, D. C., Hutt, C. R., Gee, L. S., Ringler, A. T., & Anthony, R. E. (2024). Global seismic networks operated by the US Geological Survey. *Seismological Research Letters*, *95*(3), 1578–1590. doi: 10.1785/0220230178

Wolf, J., & Long, M. D. (2023). Lowermost mantle structure beneath the central Pacific Ocean: Ultralow velocity zones and seismic anisotropy. *Geochemistry, Geophysics, Geosystems*, *24*(6), e2022GC010853. doi: 10.1029/2022GC010853

Wolf, J., Long, M. D., & Frost, D. A. (2024). Ultralow velocity zone and deep mantle flow beneath the Himalayas linked to subducted slab. *Nature Geoscience*, *17*(4), 302–308. doi: 10.1038/s41561-024-01386-5

Yamamoto, R., Song, E., & Kim, J.-M. (2020). Parallel wavegan: A fast waveform generation model based on generative adversarial networks with multi-resolution spectrogram. In *Icassp 2020-2020 ieee international conference on acoustics, speech and signal processing (icassp)* (pp. 6199–6203).

Yang, L., Liu, X., Zhu, W., Zhao, L., & Beroza, G. C. (2022). Toward improved urban earthquake monitoring through deep-learning-based noise suppression. *Science advances*, *8*(15), eabl3564. doi: 10.1126/sciadv.abl3564

Yu, S., & Garnero, E. J. (2018). Ultralow velocity zone locations: A global assessment. *Geochemistry, Geophysics, Geosystems*, *19*(2), 396–414. doi: 10.1002/2017GC007281

Yuan, K., & Romanowicz, B. (2017). Seismic evidence for partial melting at the root of major hot spot plumes. *Science*, *357*(6349), 393–397. doi: 10.1126/science.aan076

Zhou, J., Phạm, T.-S., & Tkalčić, H. (2024). Deep-learning phase-onset picker for deep Earth seismology: PKIKP waves. *Journal of Geophysical Research: Solid Earth*, *129*(9), e2024JB029360. doi: 10.1029/2024JB029360

Zhu, W., Mousavi, S. M., & Beroza, G. C. (2019). Seismic signal denoising and decomposition using deep neural networks. *IEEE Transactions on Geoscience and Remote Sensing*, *57*(11), 9476–9488. doi: 10.1109/TGRS.2019.2926772

Supplementary Materials for:

Denoising teleseismic data for deep Earth studies using a supervised deep-learning auto-encoder: a case study of diffracted waves from ULVZs

Stuart Russell[*,1,2], Katrin Hannemann[1], Carl Martin[3,4], Bennet Lindhorst[1], Waleed Esmail[5], Christine Thomas[1,6]

[1] Institut für Geophysik, Universität Münster, Correnstraße 24, 48149, Münster, Germany

[2] James Cook University, 1 James Cook Drive, Douglas, QLD, 4814, Australia

[3] Department of Geosciences, Utrecht University, 3584 CB, Netherlands

[4] Research School of Earth Science, The Australian National University, Canberra, ACT, Australia

[5] Institut für Kernphysik, Universität Münster, Wilhelm-Klemm-Straße 9, 48149, Münster

[6] Geological Survey of Denmark and Greenland, 1350 Copenhagen, Denmark

[*] Corresponding author, email: stuart.russell@jcu.edu.au

Section S1: Event details

Event name	**Event time** (UTC)	**Latitude** (°N)	**Longitude** (°E)	**Depth** (km)	**ULVZ sampled**	**Predominant station location**
20090917	2009-09-17T23:21	-29.14	-112.27	10	Hawaii	Japan
20100320	2010-03-20T11:55	-3.38	152.28	418	Hawaii	Contiguous USA
20100417	2010-04-17T23:15	-6.70	147.30	418	Hawaii	Contiguous USA
20100804	2010-08-04T07:08	-5.52	146.86	220	Hawaii	Contiguous USA
20120510	2012-05-10T02:13	-28.60	-112.45	10	Hawaii	Japan
20121207	2012-12-07T18:19	-38.36	176.07	171	Pitcairn	Contiguous USA
20130511	2013-05-11T20:47	-17.94	-175.12	210	Hawaii	Contiguous USA
20131123	2013-11-23T07:48	-17.16	-176.52	370	Central Pacific	Contiguous USA
20141009	2014-10-09T02:14	-32.30	-110.92	4	Hawaii	Japan
20150519	2015-05-19T15:25	-15.36	-132.27	18	Vanuatu	Japan
20160831	2016-08-31T03:08	-3.69	152.79	499	Hawaii	Contiguous USA
20160912	2016-09-12T06:51	-28.76	-112.01	10	Galápagos	Alaska
20161225	2016-12-25T14:22	-43.41	-73.94	38	Galápagos	Alaska
20240223	2024-02-23T10:51	-35.09	-110.77	10	Hawaii	Japan

Table S1: The event details of all earthquakes referenced in this study. Events are referred to by YYYYMMDD format.

Section S2: GSN station locations

Figure S1 shows a map of the Global Seismograph Network stations that were used when training the DAE models. In total there are 199 stations for which data is available via IRIS, distributed over four network codes as shown in Table S2. The station distribution has good global coverage at all latitudes, including in the oceans as well as on the coasts and interiors of all of the continents. The GSN is expected to give a representative sample of globally distributed noise affecting teleseismic observations.

Network code	DOI	Total stations	Stations active as of Jan. 2026
CU	10.7914/SN/CU	9	9
IC	10.7914/SN/IC	10	10
II	10.7914/SN/II	59	45
IU	10.7914/SN/IU	121	76
Total:		199	140

Table S2: The details of the networks used for the GSN training dataset.

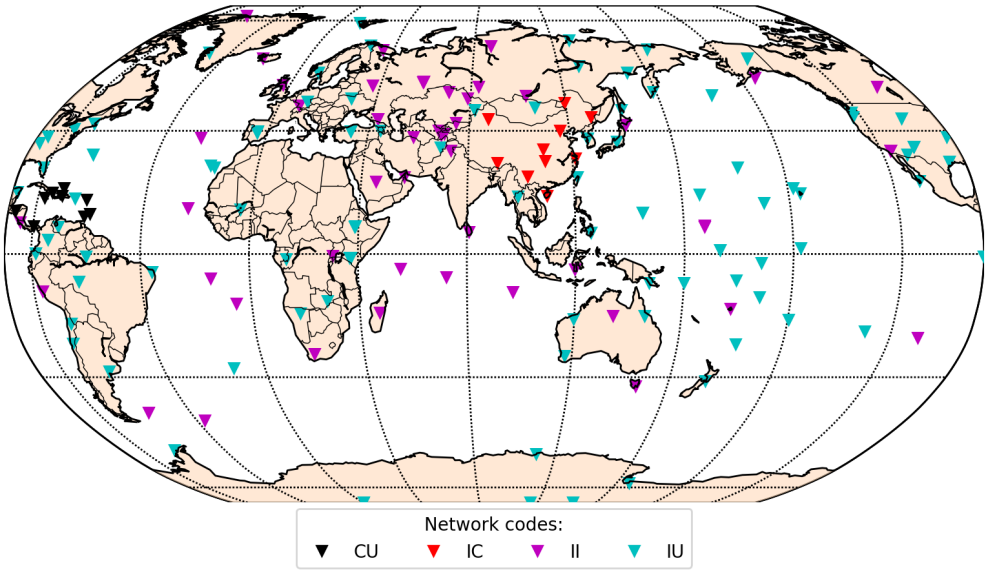

Figure S1: A map showing the locations of the GSN stations used for the GSN training dataset.

Section S3: Generation of AxiSEM3D synthetics

AxiSEM3D (*Leng et al.*, 2016, 2019) is a spectral element method (SEM) code for calculating synthetic seismograms in 3D Earth models. AxiSEM3D differs from other SEM codes by solving in a 2D computational domain and parameterising the azimuthal complexity of the wavefield as a Fourier expansion

$$u\left(r,\theta,\phi;t\right) \simeq \sum_{\alpha=0}^{N_r} u^{\alpha}\left(r,\theta;t\right) e^{\mathrm{i}\alpha\phi}, \tag{1}$$

where r and θ are the in-plane directions, ϕ is the azimuthal direction, and t is the time dimension. u is the displacement field and u^{α} is the αth order complex Fourier coefficient field, and i is $\sqrt{-1}$. The total number of terms to be summed is N_r, which is a user-chosen parameter that can vary throughout the mesh and depends on both the complexity of the model and the desired accuracy of the simulation. For global simulations AxiSEM3D can be incredibly efficient compared to traditional SEM codes with a fully-3D computational domain, provided that a simulation is properly optimised by minimising N_r.

We minimise N_r while maximising the accuracy of S_{diff} and $S_{diff}+$ by following a similar procedure to *Russell et al.* (2025). We set N_r below a certain depth (3745 km, i.e. below the CMB), before a certain distance (60°), and beyond a certain distance (130°) to a value of 5, which is the minimum value required represent a moment tensor. The accuracy of the expansion in these regions does not affect the S_{diff} phases that we are interested in. The value elsewhere is set to 500, and we verify that this setup produces accurate waveforms. Even once optimised these simulations are still extremely costly, and computation time scales with frequency to the power 3. We run a single simulation at 4 s minimum period, which on 432 CPUs on the PALMA-II cluster of the University of Münster took 15 hours. The parameters of the simulation are summarised in Table S3 and a map of the simulation geometry, as well as synthetic waveforms are shown in Figure S2.

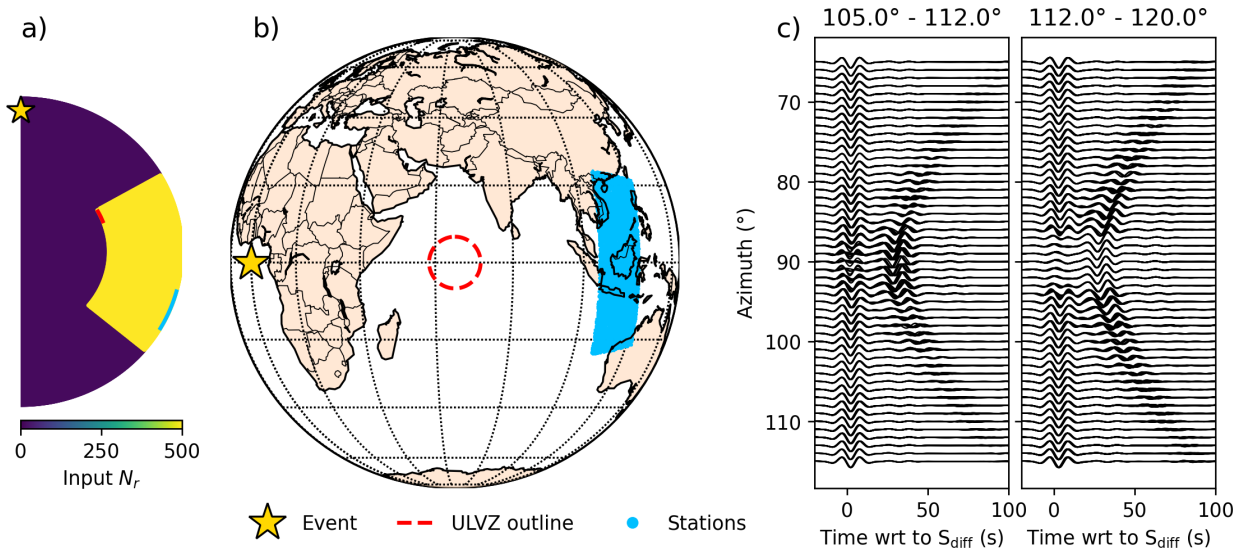

Figure S2: a) The N_r field used for the simulation shown on the 2D mesh. The locations of the event, ULVZ and stations are marked. b) Overview map of the simulation showing the position of the event, ULVZ and stations. The continents are shown to give a sense of scale, but have no meaning in the simulation. c) Example waveforms showing S_{diff} and $S_{diff}+$ for two different distance ranges indicated in the title of each plot. Traces are individually normalised such that the maximum amplitude corresponds to 0.75° azimuth, and traces are plotted rounded to the nearest degree of azimuth.

Parameter description	Value	Further remarks
General parameters		
Background model	PREM	
Minimum period	4 s	
Courant number	0.6	
Elements per wavelength	2	
Length of simulation	1,800 s	
Source parameters		
Source location	0°N, 0 °E	
Source depth	500 km	
Moment tensor	$[0, 0, 0, 10^{13}, 0, 10^{13}]$	CMTSOLUTION format, units Nm
ULVZ parameters		
Centre of the ULVZ	0°N, 65 °E	ULVZ is cylindrical in shape
Radius of ULVZ	400 km	Approximately 6.5° at the CMB
Height of ULVZ	30 km	
S-velocity reduction inside the ULVZ	-20 %	Velocity in the ULVZ is uniform
Station parameters		
Parameter recorded	Displacement	
Distance range for stations	$[105°, 120°]$	
Azimuth range for stations	$[60°, 120°]$	
Total number of stations	10,000	Stations are distributed randomly
N_r **parameters**		
Radius below which N_r is set to 5	3475 km	
Distance before which N_r is set to 5	60°	
Distance beyond which N_r is set to 5	130°	
N_r in all other regions	500	

Table S3: The parameters used for the AxiSEM3D simulation used to calculate the training synthetics.

Section S4: Synthetic testing - 1D synthetics

In this section we expand on the synthetic testing shown in the main paper. We download 10,000 noise traces from the US Transportable Array (FDSN network code TA) and 10,000 S_{diff} traces via Instaseis filtered from 2 - 100 s. Here we show the results for a model trained using the GSN, with Instaseis synthetics filtered 2 - 100 s. Both the noise and synthetics in the test dataset are unseen by this DAE during training.

Figures S3, S4 and S5 are an expansion of Figure 4a and 4b in the main paper. Figure S3 shows the inputs and outputs of the denoiser for 30 randomly chosen combinations of signal and noise, where the peak amplitude ratio of signal to noise is fixed at 2 (signal has twice the amplitude of the noise). Figures S4 and S5 show cases where the peak amplitude ratio is fixed at 1 (signal and noise are the same amplitude) and 0.5 (signal has half the amplitude of the noise), respectively. In all three of these cases, the signal output is fairly good as most signals are identified and extracted, although there are always some differences. The results are visibly poorer where the peak amplitude ratio is smaller, and in some cases the DAE can fail to extract signal where it is completely obscured by noise. However, these results are extremely promising: whereas the data shown in Figure S5 would be considered unusable after manual inspection, denoising reveals the presence of this signal, showing clear potential for elucidating hidden signals.

Figure S6 shows a test case where the DAE is given pure signal, and in this case it identifies the signal well. In Figure S7 we give the DAE pure noise, and the DAE identifies this almost perfectly and does not create any erroneous signal of significant amplitude. In none of our tests do we observe a DAE 'inventing' realistic looking signals. In Figure S7 we include a further panel, demonstrating the danger of re-normalising recovered signal after denoising. Low amplitude and spurious energy included in the signal output can be massively amplified by the normalisation, where in reality this signal is very low amplitude relative to the target phase. Despite being a common pre-plotting step to visualising record sections, re-normalising voids using the absolute amplitude as a control, and therefore should be done with caution so as not to interpret false signals.

Where signal is extracted, the phase and timing of the signal does not usually visibly change, in contrast to amplitudes which can vary significantly. Nonetheless time shifts and waveform distortion are possible, particularly where the noise level is high, and so we quantify timing and amplitude fidelity across the test dataset of 10,000 traces. The offset of the recovered arrival relative to the input is measured by cross-correlating the denoised signal trace with the clean input signal trace, and the amplitude difference is calculated using the mean-squared error between the denoised signal trace with the clean input signal trace. Differences are quantified across several imposed noise levels and results are shown in Figures S8 for amplitudes, and Figures S9 for timings; these show the distributions for five different amplitude ratios from Figure 4c and 4d in the main paper.

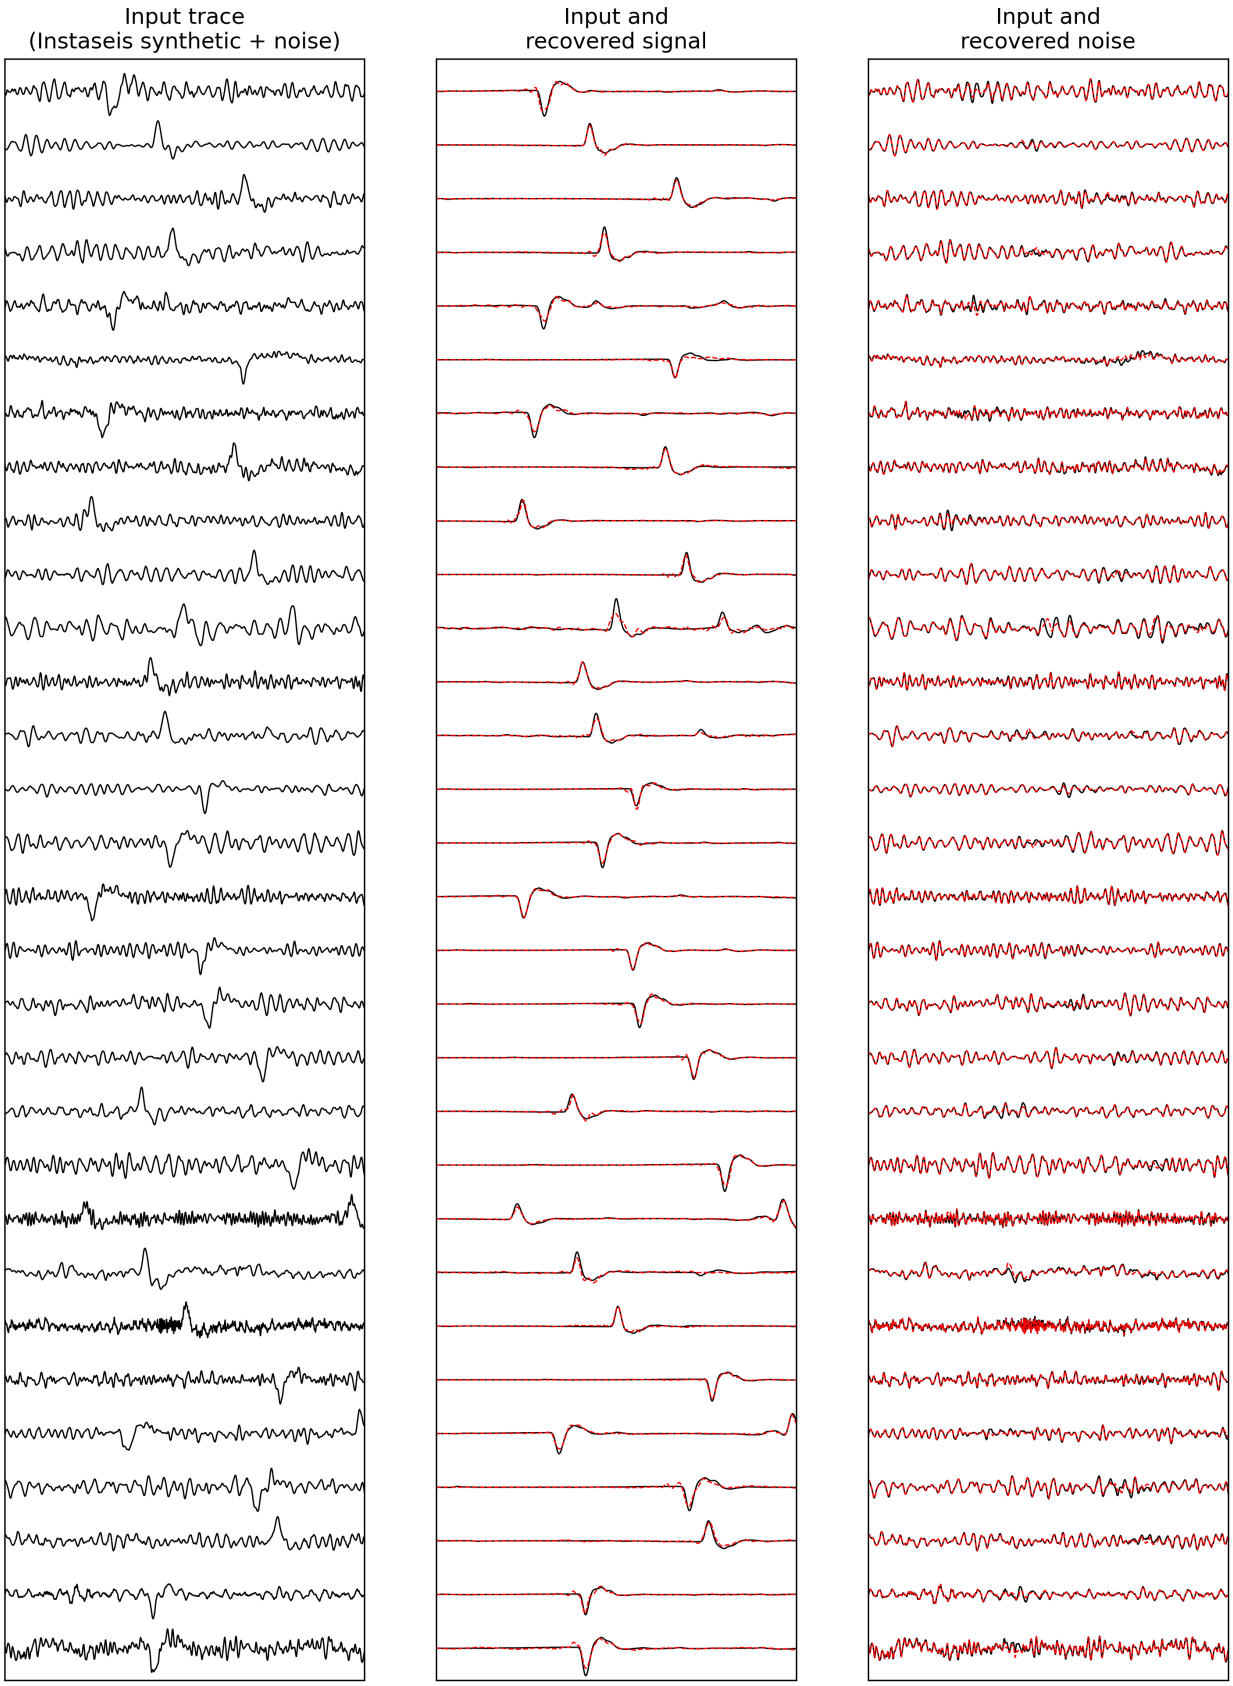

Figure S3: Application of the DAE to an unseen Instaseis synthetic test dataset where the peak amplitude ratio of signal to noise is fixed at 2. Input traces are plotted in solid black, and recovered traces are plotted on top in dashed red, with consistent amplitudes. 30 randomly selected traces are shown.

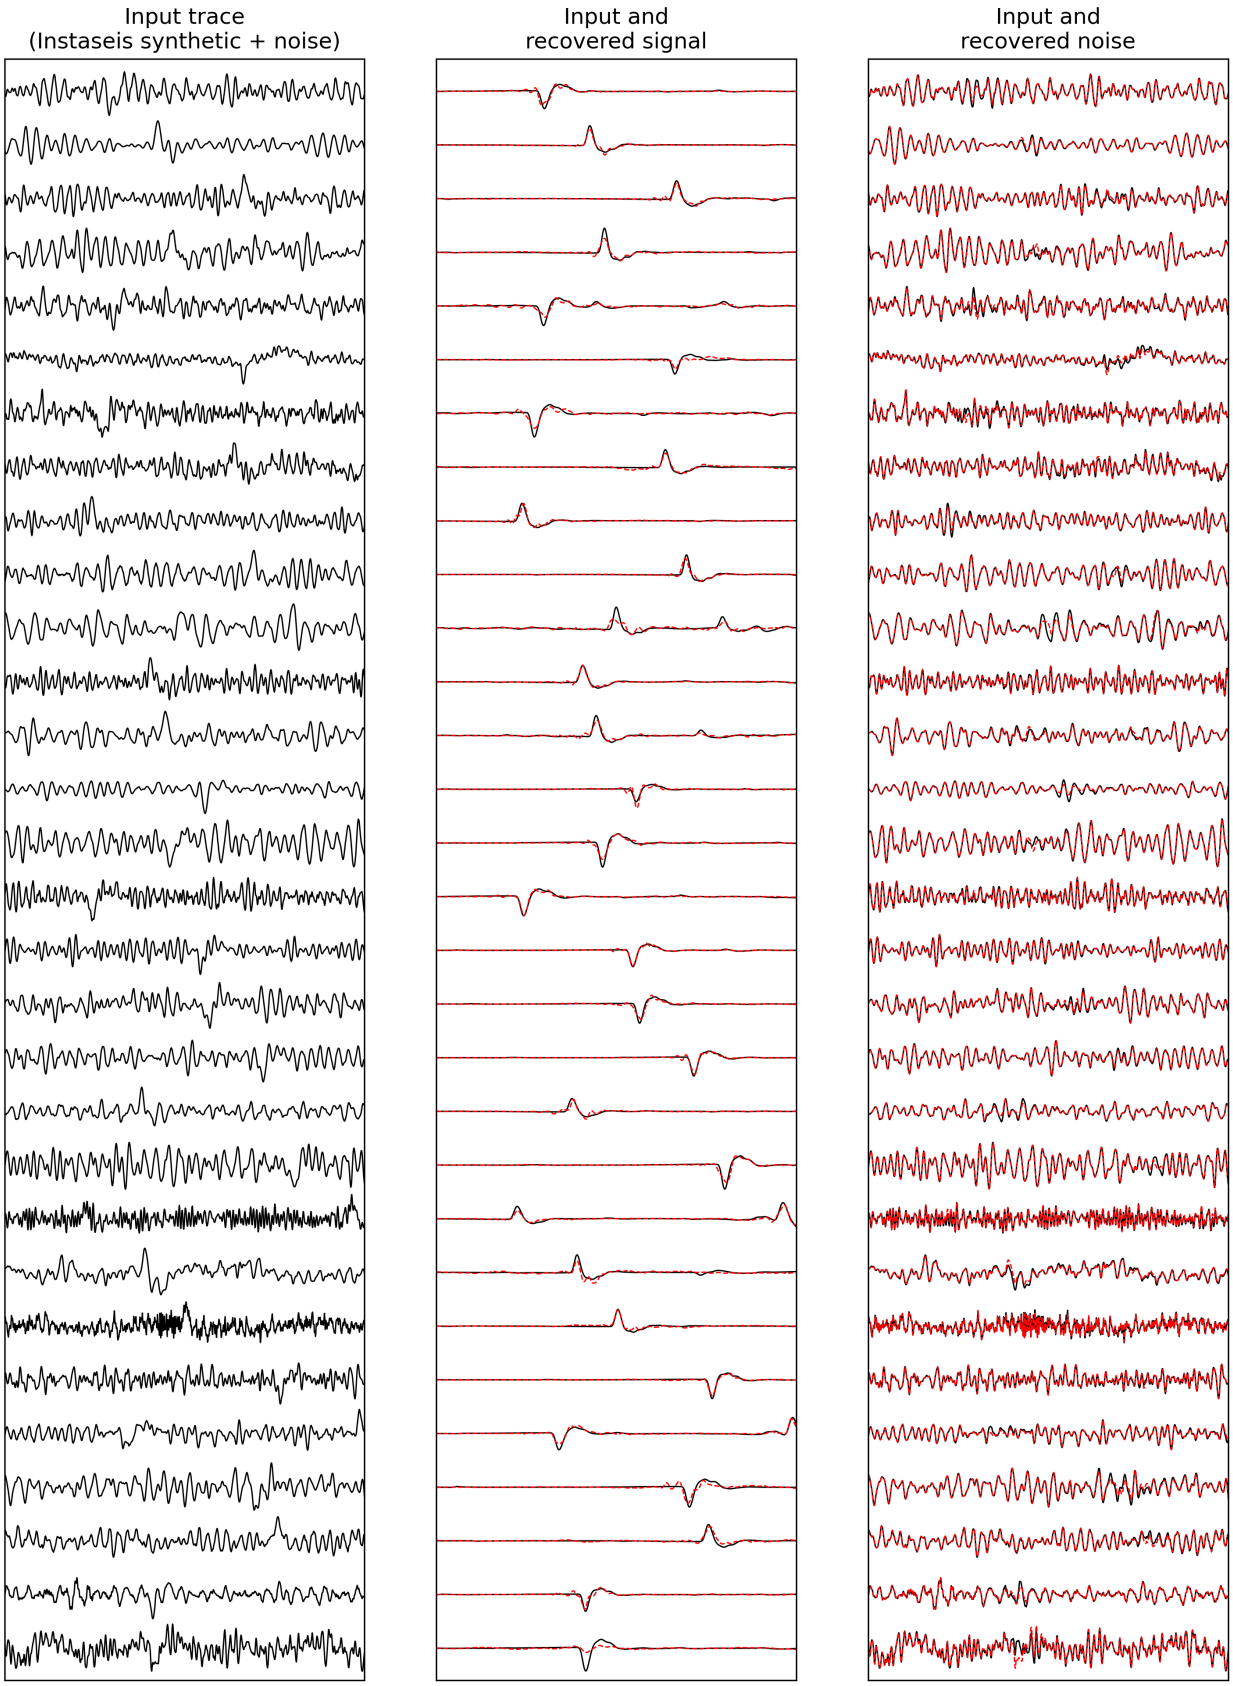

Figure S4: As for Figure S3, but where the peak amplitude ratio of signal to noise is fixed at 1.

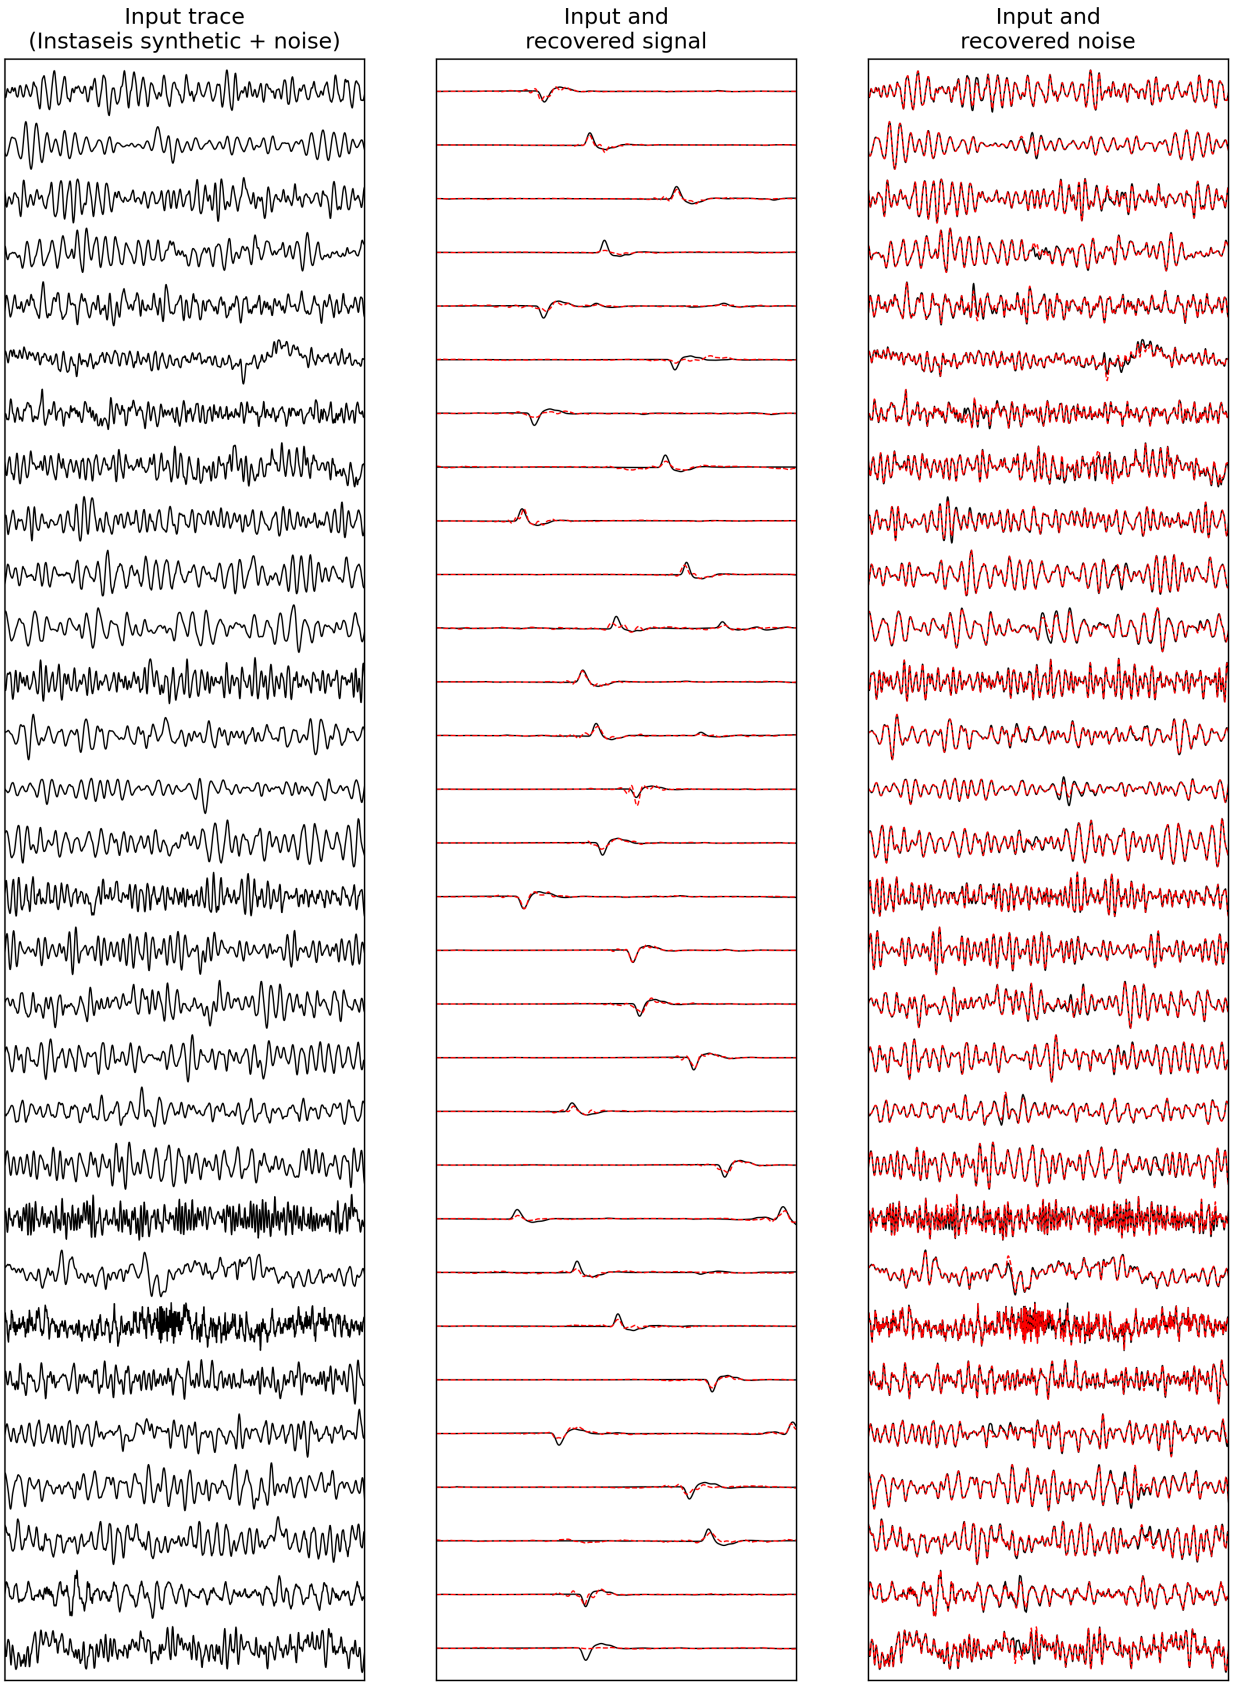

Figure S5: As for Figure S3, but where the peak amplitude ratio of signal to noise is fixed at 0.5.

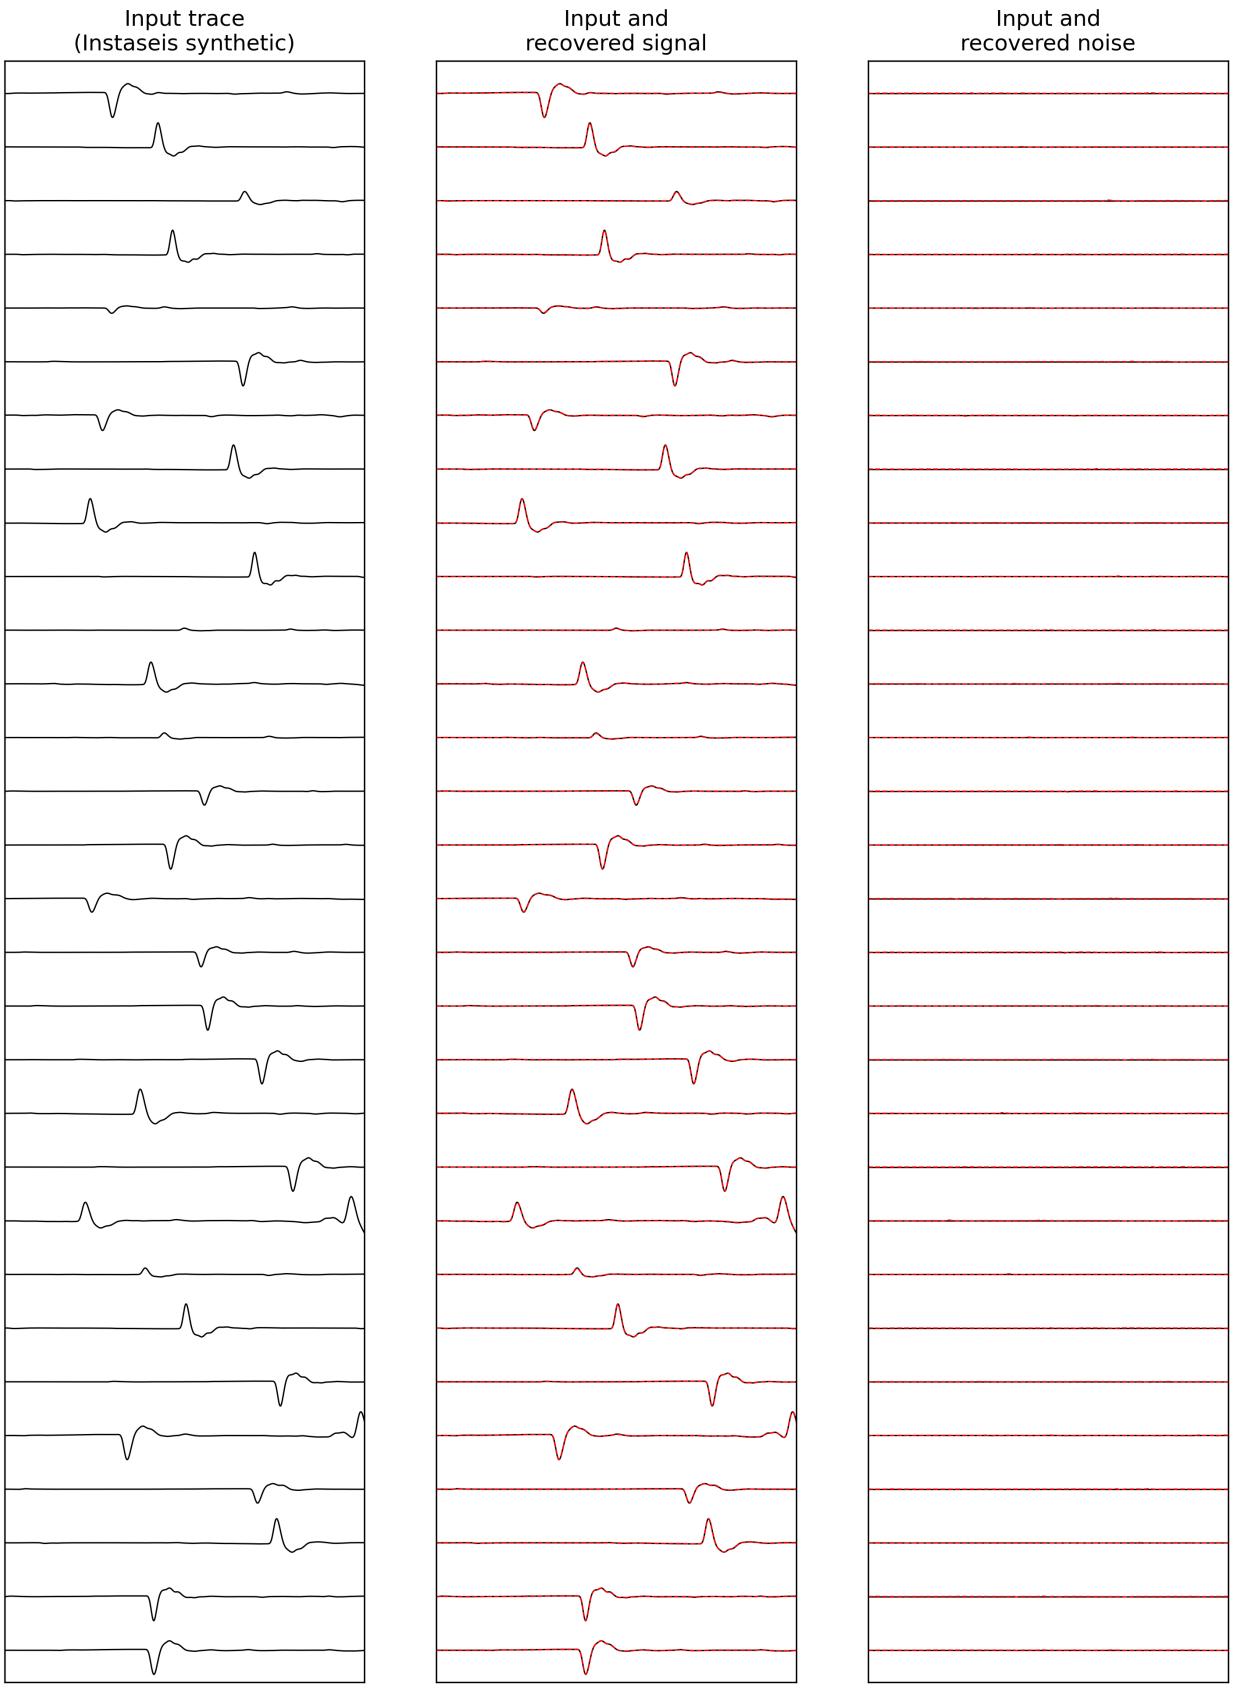

Figure S6: As for Figure S3, but where the DAE is given pure signal.

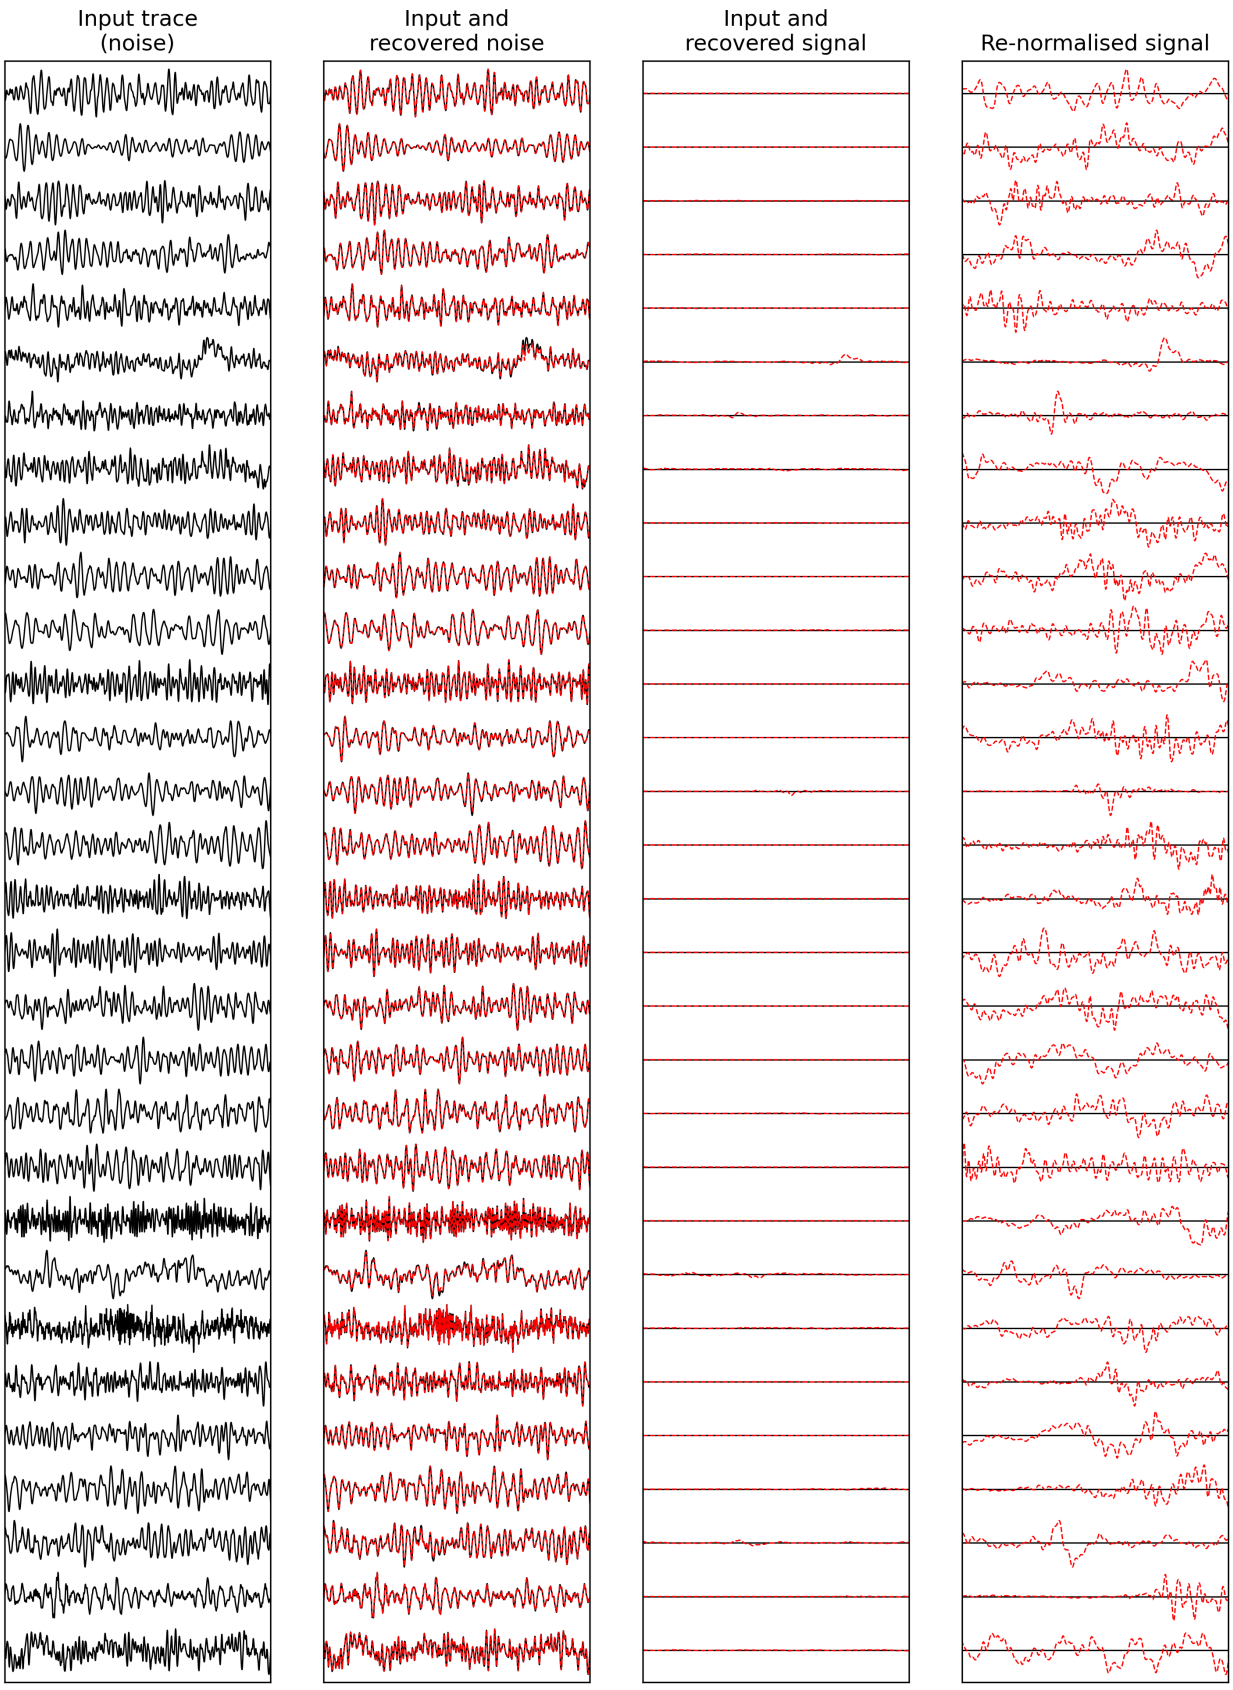

Figure S7: As for Figure S3, but where the DAE is given pure noise. An additional panel is shown where the recovered signal is re-normalised, showing the dangerous amplification of erroneous signal.

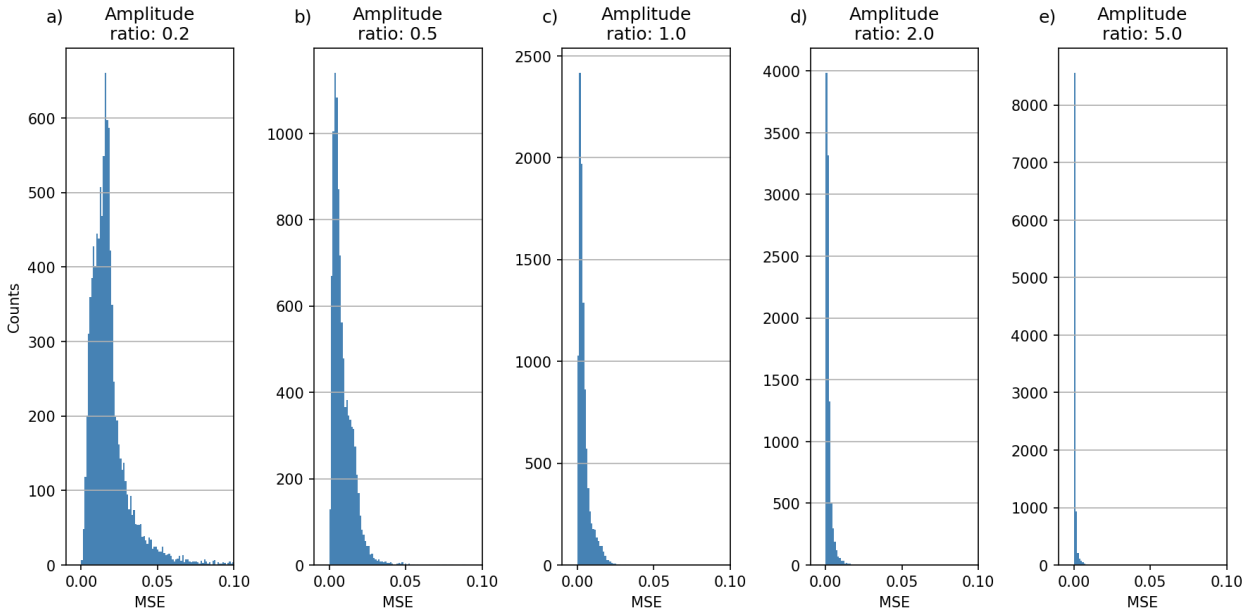

Figure S8: The mean-squared-error (MSE) of the recovered signal trace for denoising 1D synthetics for five different imposed amplitude ratios. The distribution is wider for noisier data (lower amplitude ratios).

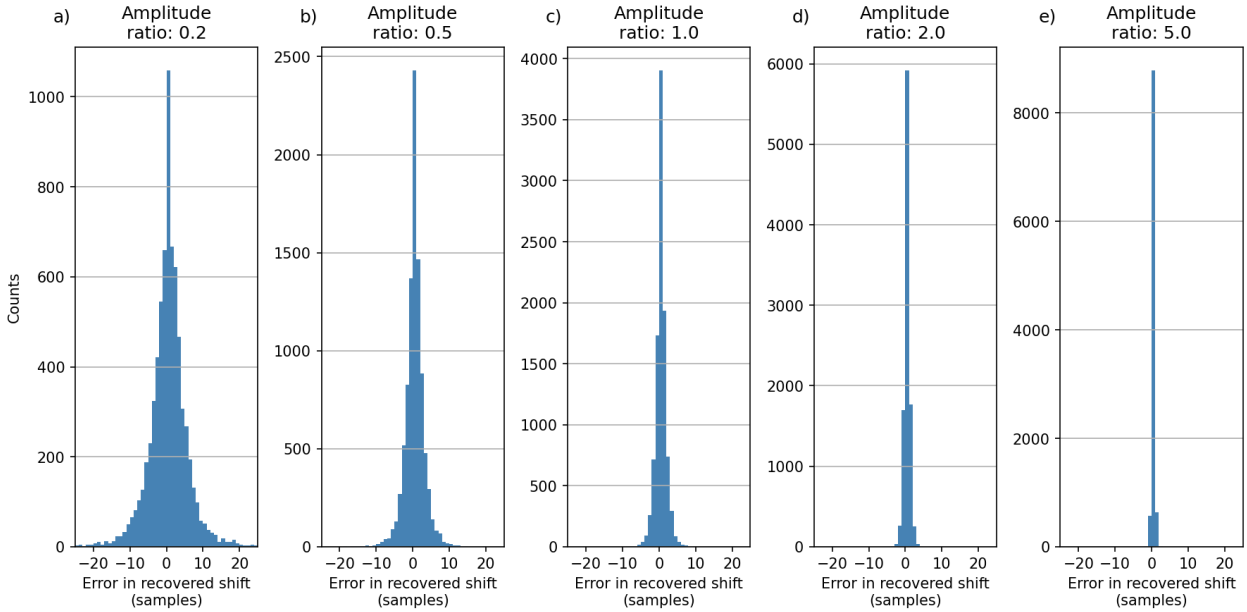

Figure S9: The error in the recovered arrival time in samples for denoising 1D synthetics for five different imposed amplitude ratios. The distribution is wider for noisier data (lower amplitude ratios).

Section S5: Synthetic testing - 3D synthetics

Beyond what is presented in the main paper, and building on what is presented in the previous section, Figures S10, S11 and S12 show the denoising of AxiSEM3D synthetics with different amplitude ratios. The results mirror the conclusions of the previous section, where extraction is better at lower noise levels, but even at high noise levels the DAE can elucidate low-amplitude signal that is obscured by the noise. In such cases, the presence of erroneous signals complicates interpretation, but these are typically low amplitude and incoherent as a function of azimuth. The DAE therefore shows great promise at screening for hidden S_{diff}+ – this is particularly exemplified in Figure S12 where the signals are barely visible in the input traces.

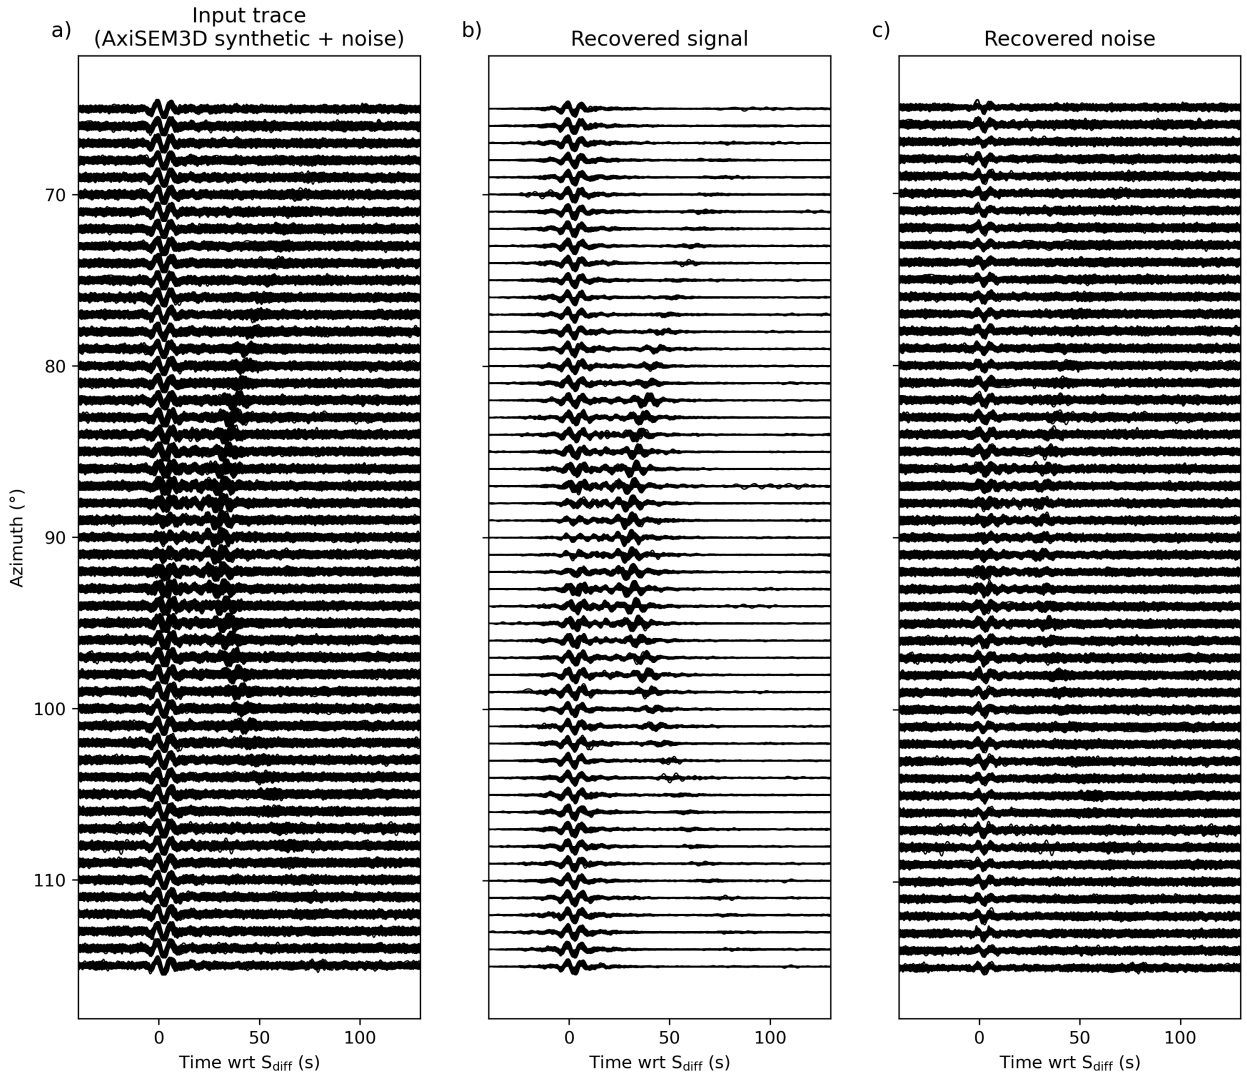

Figure S10: Application of the DAE to an unseen AxiSEM3D synthetics where the peak amplitude ratio of signal to noise is fixed at 2.0.

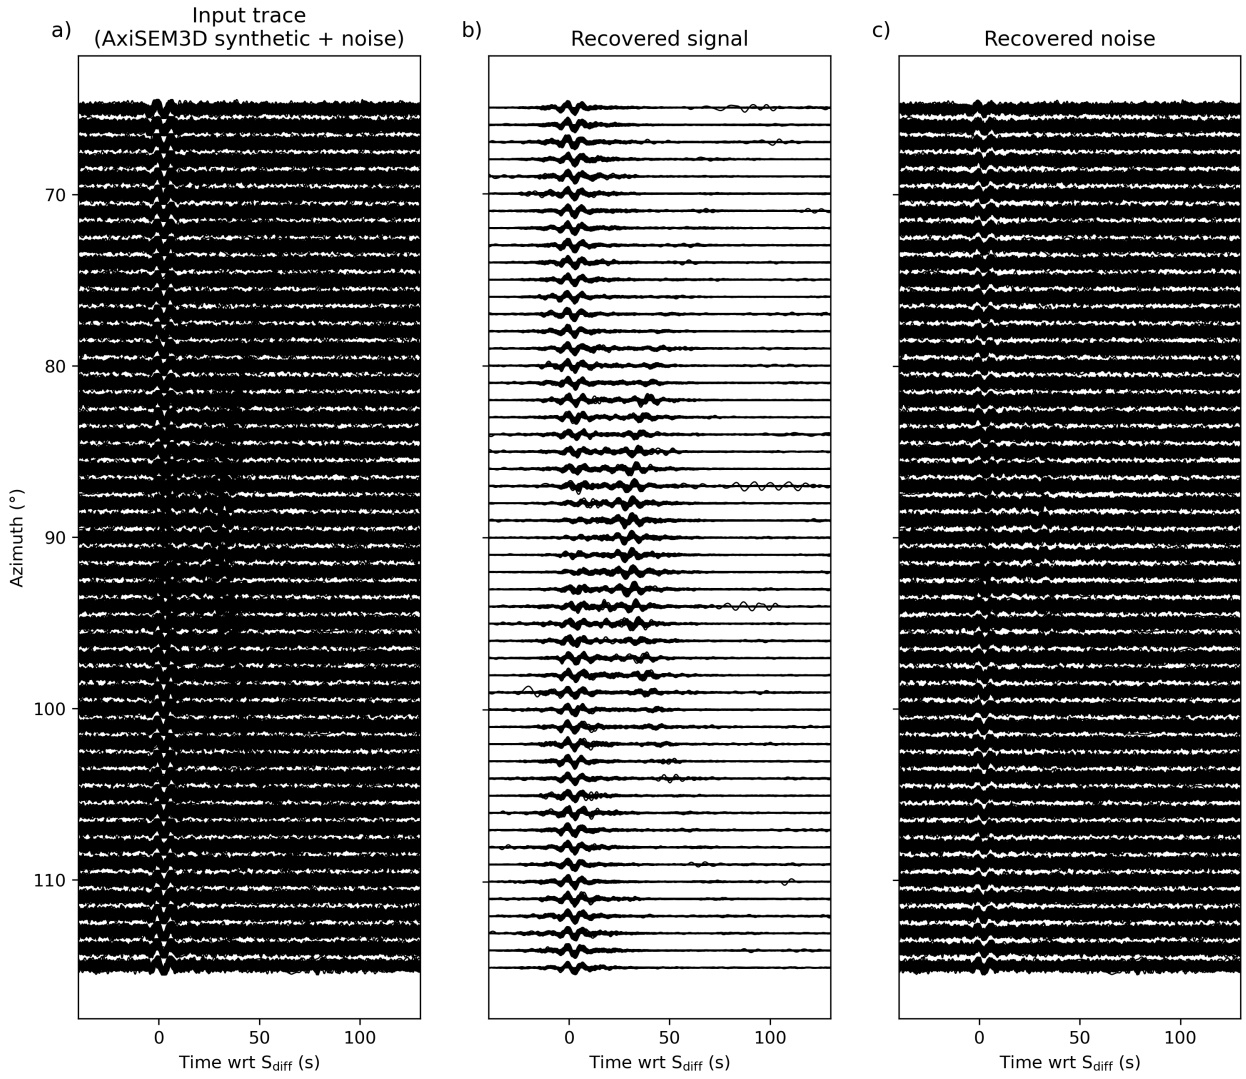

Figure S11: As for Figure S10, but where the peak amplitude ratio is fixed at 1.0.

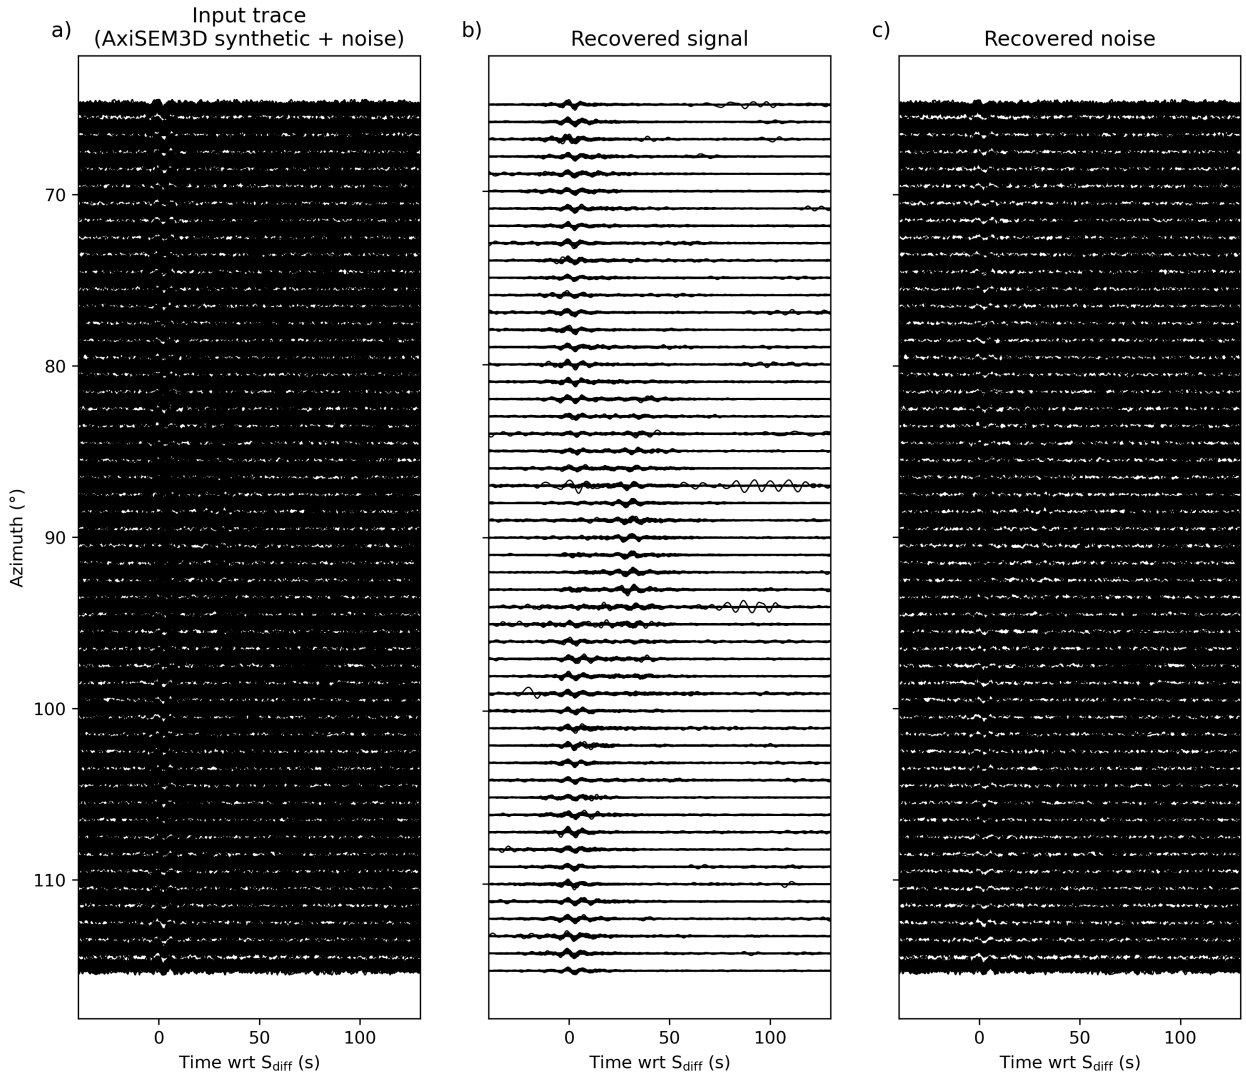

Figure S12: As for Figure S10, but where the peak amplitude ratio is fixed at 0.5.

Section S6: Noise specificity waveforms

Figures S13, S14 and S15 show the waveforms that correspond to the bars in the three panels of Figure 7 in the main paper. These data are filtered from 2.0 - 100.0 s, which is the same as the DAEs that are used for denoising.

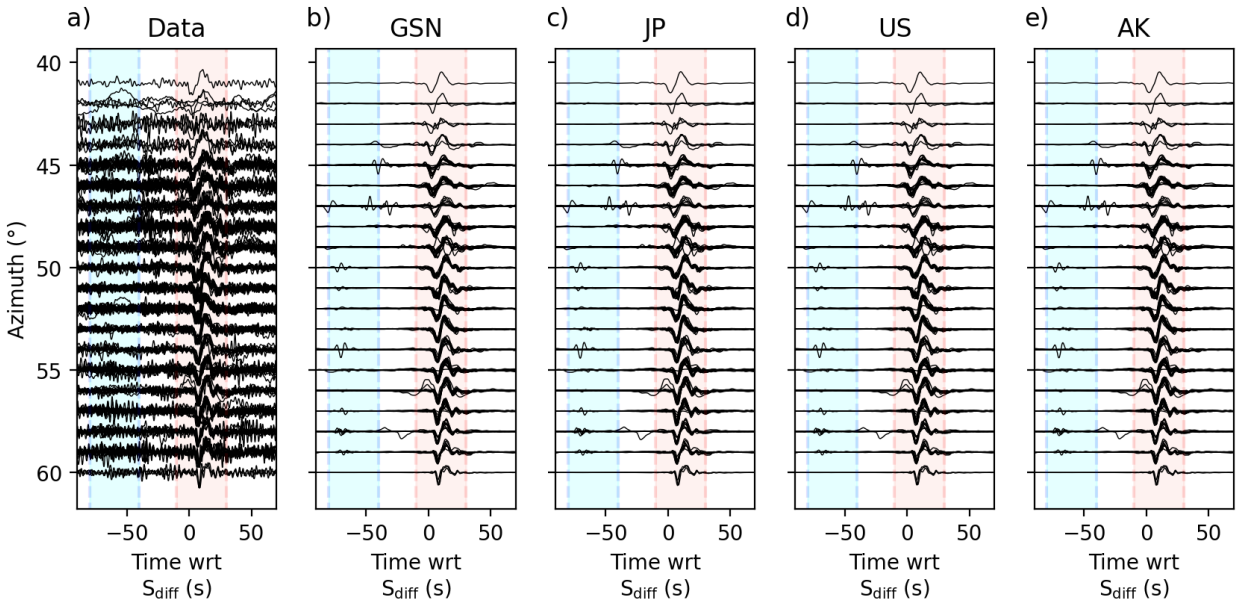

Figure S13: The waveforms corresponding to the SNR improvements shown in Figure 7 in the main paper, for event 20130511 recorded in the contiguous US. These waveforms sample the Hawaiian ULVZ. The data used to retrain the DAE that is used is indicated by the titles. In a) traces are individually normalised such that the maximum amplitude corresponds to 0.75° azimuth, and traces are plotted rounded to the nearest degree of azimuth. The amplitudes in panels b) - d) are consistent with corresponding waveforms in panel a). The shaded blue and red regions are used for calculating the RMS amplitude of the noise and the signal, respectively.

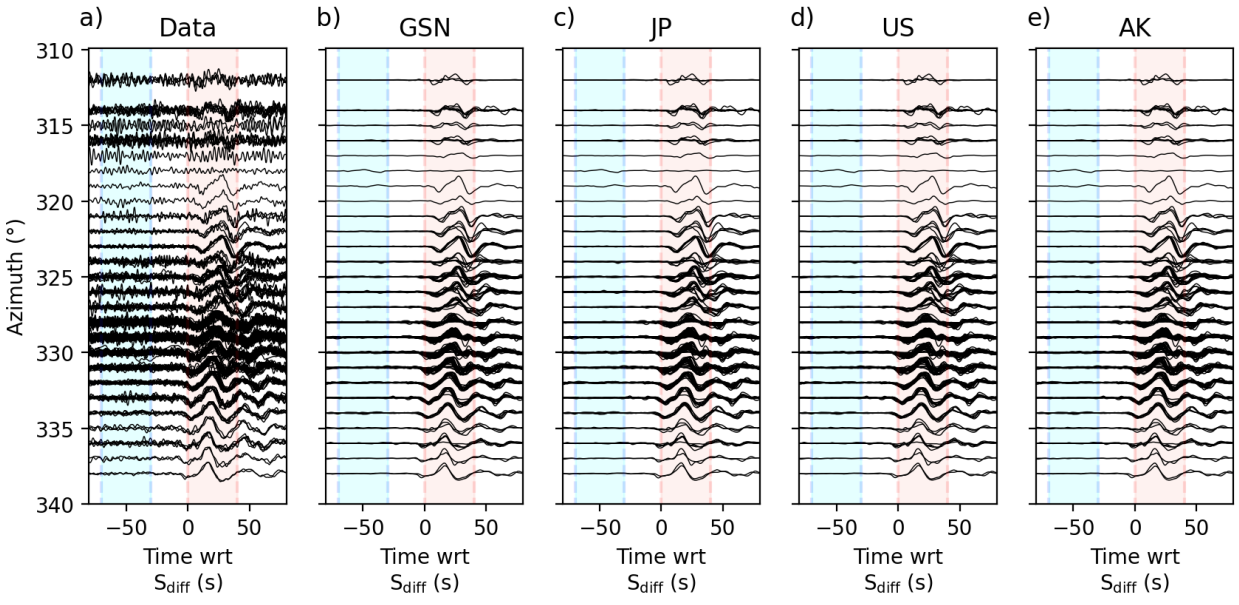

Figure S14: As for Figure S13 but for event 20161225 recorded in Alaska. These waveforms sample the Galápagos ULVZ.

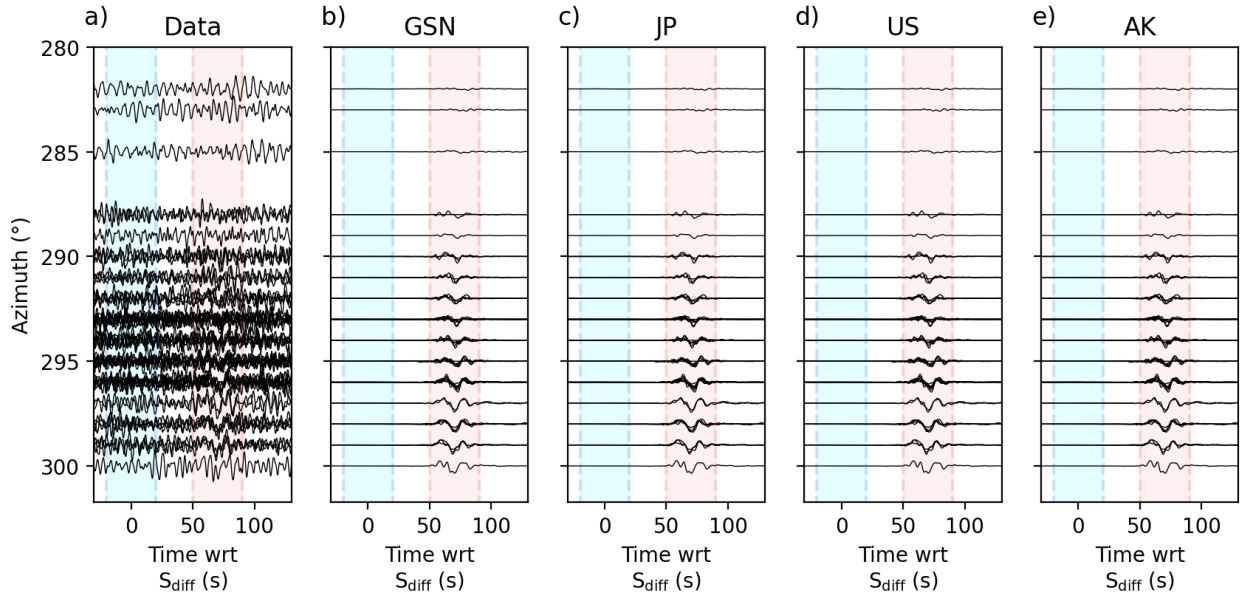

Figure S15: As for Figure S13 but for event 20090917 recorded in Japan. These waveforms sample the Hawaiian ULVZ. Note that as the waveforms are severely delayed due to the Pacific LLVP.

Section S7: Further synthetic specificity examples

In this section we show more examples of how retraining with AxiSEM3D synthetics affects the results. Specifically we show two examples where the results are not so starkly positive as in Figure 8 of the main paper. Figure S16 shows the results for Event 20160831 filtered 6.8 - 13.6 s – this event samples the Hawaiian ULVZ. The results for this event broadly mirror what is shown in the main paper, but the increase in performance when using AxiSEM3D synthetics is less. Figure S17 shows the results for Event 20161225 filtered 10.2 - 20.4 s – this event samples the Galápagos ULVZ. For this event, using AxiSEM3D synthetics has little effect on the ratio of $S_{diff}+$ to S_{diff}, but does marginally increase the SNR of S_{diff} to noise. We speculate that the ULVZ synthetics are too specific to ensure better extraction across all ULVZs, particularly those that are unlike the ULVZ in the single simulation that we ran.

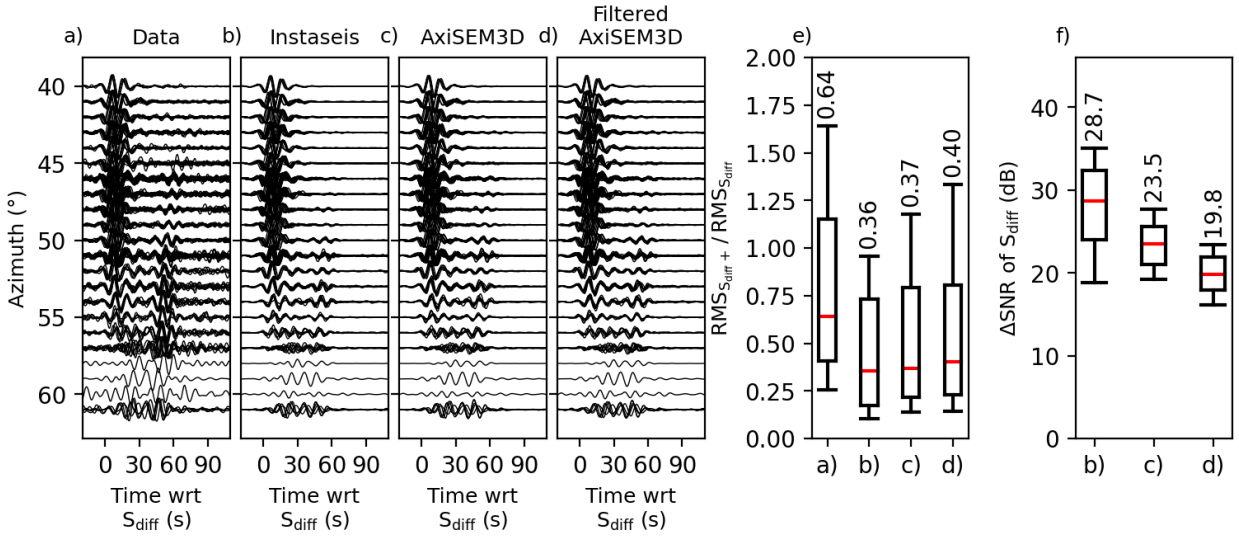

Figure S16: As for Figure 8 in the main paper, but for Event 20160831 filtered 6.8 - 13.6 s. a) The original data with each trace individually normalised such that the maximum amplitude corresponds to 0.75° azimuth, and traces are plotted rounded to the nearest degree of azimuth, b) denoised with the model trained with Instaseis synthetics, c) denoised with the model retrained with AxiSEM3D synthetics, and d) denoised with the model retrained with filtered AxiSEM3D synthetics. e) The ratios of the RMS of $S_{diff}+$ and S_{diff} for the waveforms in panels a) - d). The median is shown in red, the edges of the boxes are the 25^{th} and 75^{th} percentile, and the ends of the arms are the 10^{th} and 90^{th} percentile. The value of the median is shown above. f) As for e) but showing the improvement in SNR for S_{diff} in decibels for panels b) - d) relative to a). Absolute amplitudes in panels b) - d) are consistent with panel a).

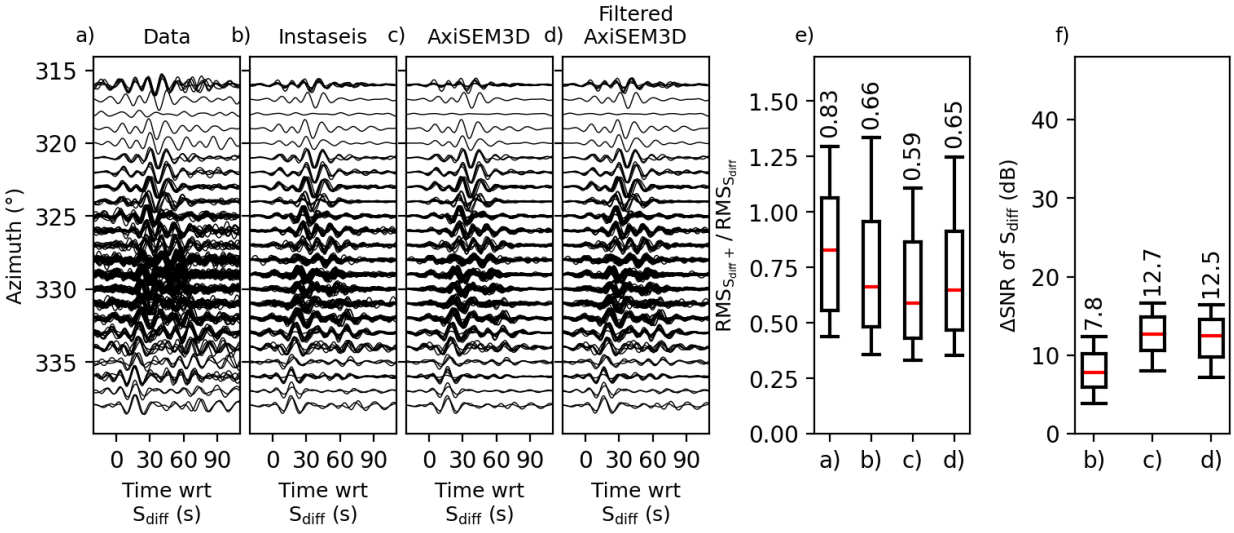

Figure S17: As for Figure S16 but for Event 20161225 filtered 10.2 - 20.4 s.

References

Leng, K., T. Nissen-Meyer, and M. van Driel, Efficient global wave propagation adapted to 3-D structural complexity: a pseudospectral/spectral-element approach, *Geophysical Supplements to the Monthly Notices of the Royal Astronomical Society*, *207*(3), 1700–1721, doi:10.1093/gji/ggw363, 2016.

Leng, K., T. Nissen-Meyer, M. Van Driel, K. Hosseini, and D. Al-Attar, AxiSEM3D: broad-band seismic wavefields in 3-D global earth models with undulating discontinuities, *Geophysical Journal International*, *217*(3), 2125–2146, doi:10.1093/gji/ggz092, 2019.

Russell, S., J. K. Magali, K. Vallenton, and C. Thomas, The effect of source-side subduction on PKP differential times and implications for inner core anisotropy, *Physics of the Earth and Planetary Interiors*, p. 107382, doi:10.1016/j.pepi.2025.107382, 2025.