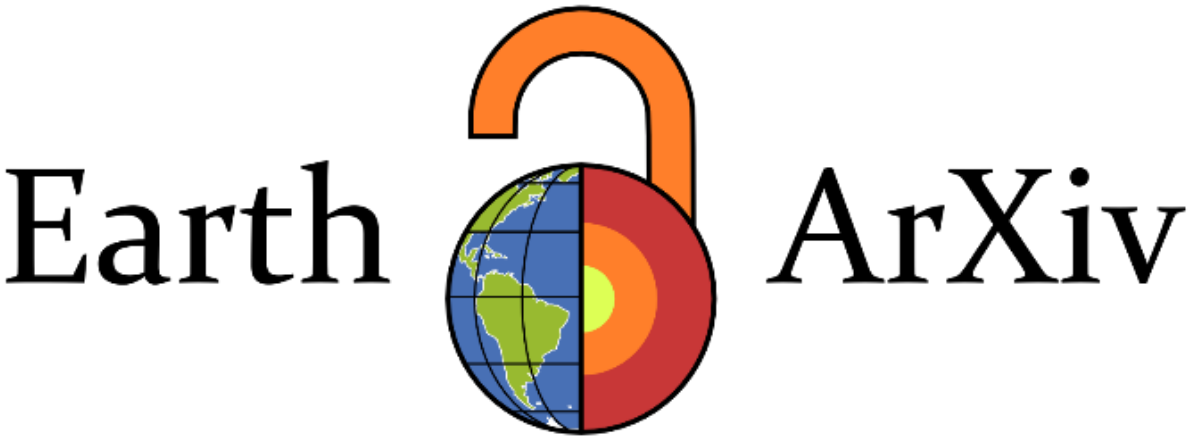

submitted to *Geophys. J. Int.*

Towards Entire Wavefield Inversion in Highly Scattering Volcanic Environments using Fourier Neural Operators

E. J. Totten[1], C. J. Bean[1] and G. S. O'Brien[2]

[1] *Geophysics Section, School of Cosmic Physics, Dublin Institute for Advanced Studies (DIAS), Dublin,* D02Y006*, Ireland. Email: etotten@cp.dias.ie*

[2] *Microsoft Ireland, Dublin, Ireland*

Received 2026 March 9; in original form 2025 December 1

SUMMARY

Seismic imaging in volcanic environments is highly challenging due to significant scattering of seismic waves on multiple spatial scales. When these wavefields are recorded on the surface by seismic arrays, seismograms generally contain information-rich codas in addition to ballistic first arrivals. Later reflected and refracted arrivals are often completely masked by the scattered coda waves. Despite the potential information contained in the coda waves their incoherency usually forces imaging to rely solely on first arrival travel time tomography, which smooths out fine scale detail, such as important individual feeder dykes. Volcanic regions such as Iceland may provide the opportunity to directly use the scattered wavefield to produce improved velocity models from seismic data, using machine learning approaches. Machine learning-driven inversion also has the potential to offer marked speed-up over conventional seismic inversion once a neural network is trained to acceptable performance. Rapid velocity model estimation is of particular importance in volcanic settings, where magmatic intrusions during periods of increased volcanic activity can pose a significant risk to both society and industry. In this study we show how machine learning could play a role in real-time seismic imaging of the upper crust beyond the current resolution of seismic tomography. We also show how machine learning could play a role in producing higher resolution images of the upper volcanic crust. We

train a Fourier Neural Operator-based neural network to invert for 2D P-wave velocity models from a single earthquake gather. We first create a core dataset of 30,000 2D velocity models with depthwise gradients representative of Icelandic crust, perturbed by up to 25% with anti-persistent Von Kármán series and 1000m correlation lengths. Secondly, we forward model the wavefield gather through each velocity model using SPECFEM2D, accounting for attenuation and broadband source properties. Thirdly we train a Fourier Neural Operator (FNO) to predict 2D P-wave velocity models from single earthquake gathers. From here we show that FNO performance generalises to unseen earthquake gathers not included during training, recovering the broadscale 2D velocity structure. Finally, we repeat FNO training with 4,000 additional models and earthquake gathers with random walk-based, low-velocity dyke intrusions down to the order of 200m scale-widths. We show that the trained FNO model can successfully recover these fine-scale dykes using strongly scattered seismic data but that recovery fidelity is influenced by earthquake source location. We discuss how FNO-based seismic inversion can complement cognate methodologies such as emerging fibre-optic geodesy and seismicity approaches to monitor for dyke and sill intrusions in real-time.

Key words: Seismic imaging, Volcanoes, Earthquakes, Seismic waves, Seismic inversion, Machine Learning, Neural Network, Fourier Neural Operators (FNO), Dykes, Magmatic Intrusion, Time lapse imaging

1 INTRODUCTION

Seismic imaging is a crucial geophysical method for understanding volcanic environments by identifying the seismic velocity structure in the upper crust. When repeated over adequate timescales and compared with seismicity monitoring, it can facilitate an improved understanding of dynamic properties in volcanoes related to heat and mass transfer processes. This is important in volcanic regions such as Iceland, where elevated crustal temperatures are both a critical energy heat source and a potential indicator of volcanic hazard. For example, ray-based seismic travel-time tomography has been applied in the Reykjanes geothermal region of SW Iceland, providing a robust 3D velocity model (Jousset et al., 2016). Glück et al. (2025) similarly performed local-scale tomography of the Krafla volcanic field down to 3km depth. Rahimi Dalkhani et al. (2024) used the

ambient noise field for transdimensional ambient-noise surface wave tomography to 20km depth beneath Reykjanes geothermal fields at resolutions sufficient for geothermal exploration.

However, periods of sustained volcanism where magma enters the upper few kilometres of the crust pose a significant hazard to society and industry. Dyke intrusions facilitate magmatic transport to the surface and can range from 0.1m to 100m in outcrop scale (Gudmundsson 2025). In order to image such fine-scale structure, both high frequencies and seismic reflections and diffractions are required. Travel-time tomography generally 'smoothes out' structures at this important scale. Due to their physical heterogeneity, volcanoes are highly scattering to seismic waves (De Siena et. al. 2016; Bracale et. al. 2025), producing substantial high amplitude seismic coda that prevents identification of coherent seismic reflections. This makes high resolution reflection imagery on volcanoes challenging. The seismic coda that typically follows the main arrivals in a given seismogram is produced by multiple scattering of the wavefield throughout various points in the velocity model domain. It is important to stress that, although complex, these coda waves should not be considered as noise. When collectively considered in an earthquake gather from a heterogeneous medium, they have the potential to 'fingerprint' the velocity model's spatial distribution of heterogeneity. However, inverting for all orders of seismic wavefield scattering is a significant challenge. Seismic imaging is also limited by the seismic coverage available and can result in an under-recovery of fine-scale features such as low velocity zones (Paulatto et al. 2022). Li et. al. (2025) mitigated for this by leveraging low-frequency distributed acoustic sensing on fibre-optic cables to track dyke intrusions on a timescale of minutes near Grindavík, Iceland. Glastonbury-Southern et al. (2025) used an earthquake source relocation approach on 12,000 earthquakes to constrain the spatial extent of the 2021 Fagradalsfjall dyke intrusion. We therefore identify a need for improved seismic imaging methods that utilise as much of the seismic wave coda from earthquakes as possible for further improved resolution at upper crustal depths and in a time-lapse manner, for future dyke imaging. Here, we demonstrate how advanced machine learning approaches to seismic imaging can play a critical role in utilising entire seismograms in the imaging process.

Numerous seismic datasets recorded in active and heterogeneous volcanic environments yield seismic data with highly complex codas due to multi-scale, high strength seismic scattering. For

example, Maass et al. (2025) performed passive reflection seismic imaging using nodal array data in the vicinity of the IDDP1 borehole near the Krafla caldera, northeast Iceland. They observed that seismograms recorded by neighbouring seismometers (30m spacing) exhibited marked incoherency and that the seismograms still contained significant complexity after applying resonance removal filters for near-surface site effects, indicating that they contain important information about the upper crust. It is important to stress that while seismic codas recorded in volcanic settings can exhibit complexity across a broad spectral range (Obermann et al. 2013), they still contain important information about the upper crust, which, in principle, has the potential to be leveraged by seismic imaging techniques. This presents an opportunity to investigate whether machine learning approaches can deliver seismic velocity models to match (and in future, extend beyond) state of the art seismic tomography resolution on volcanoes. Fourier Neural Operator-based (FNO) model architectures (Li et al. 2020) are an ideal learning tool because of their ability to learn multidimensional mappings between function spaces, making them conducive to solving partial differential equation-based problems, including the seismic wave equation. FNOs are also capable of predicting the inverse Green's functions needed for velocity inversion (O'Brien et. al. 2023): this forms the main focus of this work.

Previous studies have demonstrated that Fourier neural operator-based machine learning models can be applied to seismology. Yang et al. (2021) used an FNO to forward model acoustic wavefields across a wide range of 2D velocity models and source locations. Yang et al. (2023) demonstrated 2D full-waveform inversion by automatic differentiation using a U-shaped neural operator. Other studies have demonstrated the feasibility of using FNOs for 3D waveform forward modelling (Lehmann et al. 2023; Lehmann et al. 2024; Kong & Rodgers 2023; Kong et al. 2025). Factorised Fourier Neural Operators (FFNO), which replace the full complex weight matrix of the standard FNO with a low-rank tensor, have also been explored for efficient 3D waveform forward modelling and uncertainties (Lehmann et al. 2023; Lehmann et al. 2024). Multiple Input Fourier Neural Operators (MIFNO) have been used to encode model-to-waveform learning and source information in separate learning branches (Lehmann et al. 2025; Perrone et al. 2025). Other studies have shown that neural operators can be successfully applied to seismic datasets collected from the

physical world (Zou et al. 2025). Few studies have investigated using FNOs to perform the task of seismic inversion from seismic data directly. O'Brien et al. (2023) inverted synthetically-generated seismic waveforms for 2D velocity models of compressional velocity, shear velocity and density, using an FNO trained on a 20,000-sized dataset of 2D velocity models and seismograms.

This study builds from O'Brien et al. (2023) with the central objective to train a Fourier Neural Operator to predict a subsurface velocity model from seismic waveforms that have been highly scattered on multiple spatial scales, commensurate with that expected for volcanic environments. In this paper we use simulation-to-simulation (sim-to-sim) machine learning to determine the ability of FNOs to predict geologically feasible models from highly scattered seismic data. Here we attempt to show that adopting a 'sim-to-sim' paradigm is a starting point for future 'sim-to-real' applications for three practical reasons. Firstly, working with synthetically generated velocity models and their corresponding earthquake gathers enables us to qualitatively and quantitatively assess FNO prediction performance because the ground truth velocity model is a known entity. Secondly, adopting a sim-to-sim framework positions us to constrain the recovery resolution limits that we can reasonably expect to achieve from seismic data scattered at wavelengths and strengths encountered in the physical-world. Finally (and perhaps most importantly) it is possible to establish a direct statistical link between sim-to-sim machine learning model datasets and the properties of geological environments in the physical world.

Borehole geophysical measurements enable us to understand multi-scale fluctuations in seismic velocity and density, to build up a picture of the scattering physics expected for a given geological setting (Bean 1996; Holliger 1996). In section 2 we discuss the use of sonic log measurements from volcanic rocks in the physical world to constrain scattering strength and the scale lengths over which scattering takes place. We then discuss how we embed these scattering properties in a large databank of synthetic velocity models for FNO training, as well as the full waveform modelling required to generate a corresponding earthquake gather for each velocity model. Section 3 provides a summary of FNO model architecture and the performance results for the core training and validation datasets. In section 4 we show results based on the stress testing of the FNO model object trained for direct inversion for subsurface velocity structure from a single earthquake

gather. This includes appraising the FNO's ability to recover geological features of interest such as dykes. We show how fine-tuning the core training set with a small number of additional velocity models, embedded with geological features of interest, improves generalised feature recovery. In section 5 we discuss future potential directions for FNO-based velocity model inference based on the results shown here, including 'sim-to-real' velocity model recovery for initial subsurface diagnostics, rapid time-lapse seismic imaging in dynamic volcanic environments and related geo-risk mitigation.

2 DATASET AND METHODS

2.1 Fourier Neural Operator Overview

In this study we trained a Fourier Neural Operator (FNO) to perform seismic inversion i.e. to predict 2D crustal-scale velocity models from earthquake gather array data. The seismic inverse problem can be represented as

$$\mathbf{G}^{-1}\mathbf{d} = \mathbf{m} \tag{1}$$

where G represents the (elastodynamic) Green's function for $\mathbf{m}$, a given model of the subsurface. $\mathbf{d}$ represents the seismic wavefield which has propagated through model $\mathbf{m}$, recorded at discrete locations. In practice the inverse Green's function $\mathbf{G}^{-1}$ maps earthquake gather $\mathbf{d}$ back to its velocity model of origin, $\mathbf{m}$ (O'Brien et. al. 2023). In order to train an FNO to reliably predict a velocity model from a given time-distance earthquake gather, we created a large synthetic databank of seismic velocity models with statistical properties representative of highly scattering volcanic crust. From here we forward modelled the full wavefield through each unique velocity model in the databank to create a corresponding earthquake gather, with the earthquake source in the model varying in space.

In the following section we discuss the preparation of the large databank used to train the FNO to perform seismic inversion, including:

(i) the creation of 2D crustal-scale velocity models using Von Kármán series to represent the statistical properties of seismic velocity distribution in volcanic environments

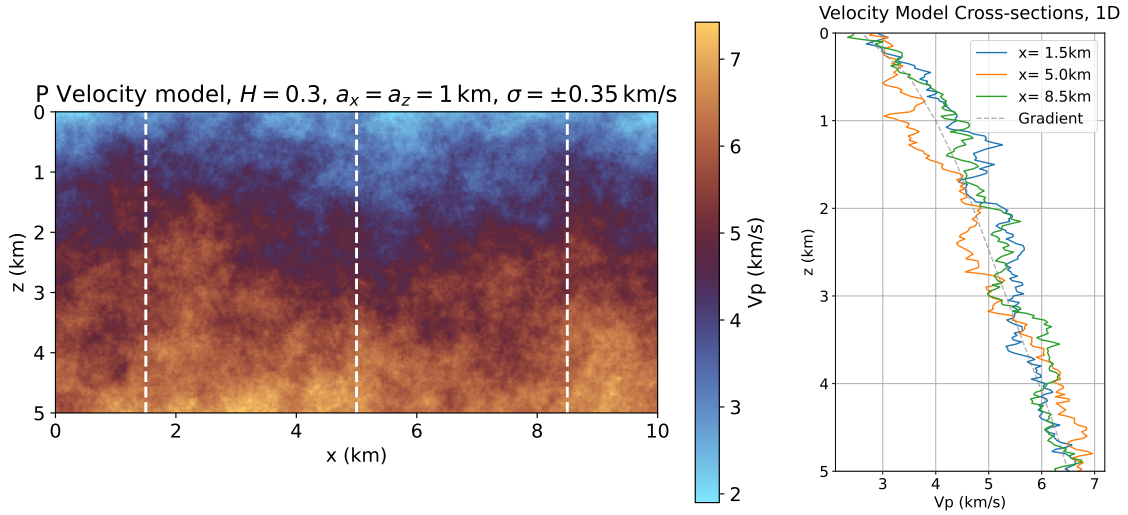

(a) (b)

Figure 1. (a) A typical example of a 2D P-wave velocity model from a larger 30,000-sized databank used in this study. Each velocity model is the sum of two parts; a 1D velocity gradient considered to be generally representative of Icelandic volcanic crust, and fractional perturbation up to 25% by Von Kármán series. The latter has a Hurst exponent of 0.3, isotropic correlation lengths in both dimensions of 1000m and 1 standard deviation $\pm 0.35 km/s$. (b) 1D cross sections through the 2D velocity model shown in (a) at x=1.5km, 5km and 8.5km respectively (dashed white lines). The reference velocity gradient common to all velocity model instances is the grey dashed line.

(ii) steps taken to base velocity model statistics on what is repeatedly observed for volcanic rocks in the physical world

(iii) full numerical forward modelling of the seismic wavefield through each velocity model in the databank, accounting for wave scattering and viscoelastic attenuation.

2.2 Velocity Model Databank

We created 30,000 unique seismic velocity models with 10km horizontal and 5km vertical extent. Each velocity model is the sum of two parts:

(i) a 1D velocity gradient in depth, broadly representative of upper crustal structure in Iceland, extrapolated across the lateral extent of the model. The starting velocity at the surface and base of each model was 2.5km/s and 6.5km/s respectively before applying von Kármán perturbations.

(ii) up to 25% perturbations to the background model using 2D von Kármán series.

An expanded explanation of our velocity model creation workflow and justification for our model parameter choices can be found in appendix **A**.

Figure 1 demonstrates one example model and 1D depthwise cross-sections at 1.5km, 5km and 8.5km along profile, to highlight the scale and magnitude of velocity heterogeneity.

The Hurst exponent H controls the degree of persistence in the velocity medium's structure and can be interpreted as a measure of medium roughness. Smaller values of H correspond to rougher, more rapidly fluctuating media and therefore stronger seismic scattering. We use H=0.3 to generate anti-persistent velocity models, consistent with field measurements of the Hurst exponent in volcanic rocks from borehole studies (e.g. Martini et al. (2005); Bean & Martini (2010); Nelson et al. (2015)). The correlation lengths a_x and a_z define the characteristic horizontal and vertical scales of heterogeneity within the velocity medium. Seismic scattering is strongest when the seismic wavelength is comparable to, or smaller than, the correlation length, because waves interact efficiently with heterogeneities of similar or smaller scale. Conversely, when the seismic wavelength is much larger than the correlation length, scattering becomes weak as the wavefield effectively averages over the small-scale structure. In this study, guided by borehole sonic log information, we chose large correlation distances based on Von Kármán noise with $a_x = a_z = 1000m$. See figure A1 for the power spectrum corresponding to an endmember velocity model realisation. We used the Hurst exponent and correlation length of the Von Kármán series in combination with its standard deviation (i.e. its amplitude relative to the starting 1D background velocity model gradient) to represent the velocity model expression of volcanic rocks in the physical world. We chose to constrain the standard deviation for each Von Kármán series by consulting borehole studies through volcanic rocks (Naylor et al. 2003). Figure A2 provides an example of such data. We chose a one standard deviation of $\pm$350m/s for each Von Kármán series in the synthetic dataset. This resulted in minimum and maximum permitted element velocities of 1.9km/s and 7.65km/s across all models in the dataset (von Kármán series were clipped at $\pm$3 standard deviations and further limited to $\pm$2 standard deviations for the upper 500m of each model). We relate S-wave velocity in each 2D Von Kármán model to P-wave velocity using

$$V_s = \frac{Vp}{\sqrt{3}} \tag{2}$$

and set the density for each model using

$$\rho = 1700 + 0.2V_p \tag{3}$$

as used for previous studies of visco-elastic wavefield propagation in volcanic environments (O'Brien & Bean 2009).

2.3 Full Waveform Forward Modelling (SPECFEM2D)

2.3.1 Waveform simulation overview

After compiling 30,000 velocity model realisations as described in section 2.2, our next task was to forward model the seismic wavefield through each model. We used SPECFEM2D (Komatitsch & Tromp 2002; Tromp et. al. 2008) software for this purpose. Given that this is a full wavefield simulator, the near and intermediate field effects (Lokmer & Bean 2010) are automatically included. We stored earthquake gathers from surface-recorded line arrays to be used as input for FNO training. Figure 2 shows an example of an earthquake gather and its corresponding velocity model of origin. We simulated 6 seconds in total after the origin time for each earthquake gather. As highlighted in section 2.1, each velocity model was discretised on a 401 x 201 grid with uniform grid spacing of 25m along each dimension. Each receiver recorded the horizontal and vertical velocity field as a function of time. We opted to record the wavefield with a dense receiver spacing for full downstream workflow flexibility.

We set source locations by drawing x,z co-ordinate pairs from a uniform random distribution to ensure that the FNO was exposed to a broad variety of earthquake gather stencil shapes and slopes. We intend to introduce context-specific source distributions and varied free surface topography into our earthquake gathers in future work. The earthquake source time function used for each wavefield was a delta pulse filtered between 2-25Hz (created outside of SPECFEM) in order to mimic low magnitude local seismicity and to introduce approximately broadband spectral content into each earthquake gather (figure 3). Figure 3b shows the corresponding power spectral density of representative seismograms from all 30,000 earthquake gathers, each of which was generated

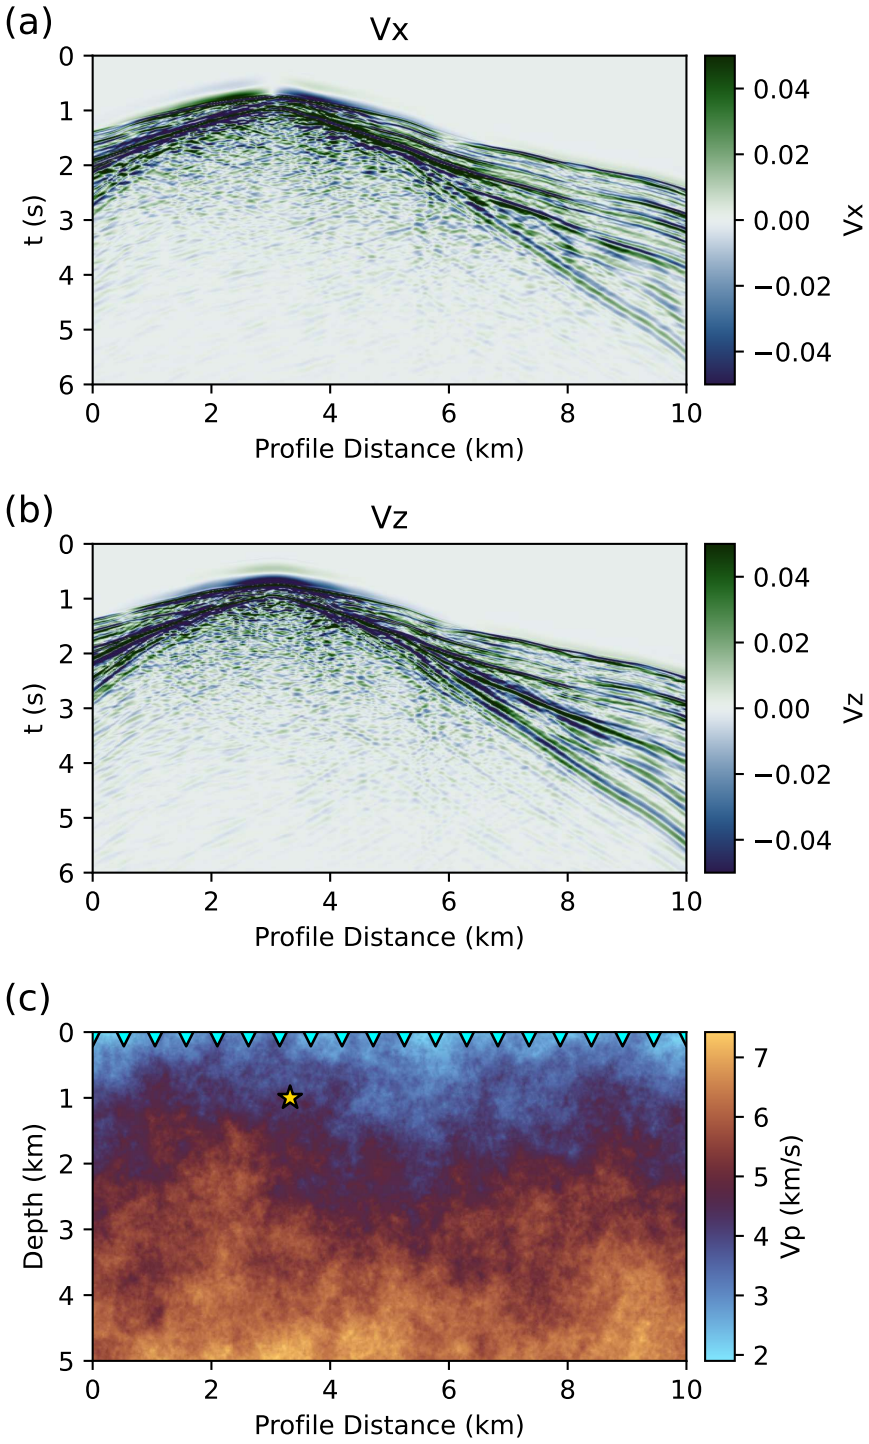

Figure 2. Example of a forward modelled earthquake gather and its velocity model of origin. In this study an FNO is trained to predict a velocity model when fed a single earthquake gather. It is trained using ≈ 30,000 unique earthquake gather and velocity model pairs, with a variable source location. **(a)** X-component velocity earthquake gather (analogous to radial component velocity seismograms in the physical world) **(b)** Z-component velocity earthquake gather **(c)** 2D velocity model of origin. The earthquake location is marked by a yellow star and receiver locations are marked by cyan inverted triangles. Only a subset of all receivers are shown for figure clarity. The actual receiver spacing is 25m, with 50m spacing used for FNO training- see text.

in a unique velocity model realisation. Here it can be seen that there is peak frequency content out to at least 20Hz. The source mechanism for each wavefield simulation used an isotropic moment tensor (M_{ii}=1).

2.3.2 SPECFEM simulation Parameterisation

For each earthquake gather generated in SPECFEM2D, we used a Newmark time stepping scheme on a Cartesian planar grid with equal vertical and horizontal grid spacings of 25m. Each simulation accounted for viscoelastic attenuation using the standard linear solid Q model with 3 standard linear solids. Here we set attenuation factors $Q_p = 60$ and $Q_s = 30$ (or $Q_\kappa \sim 100$ and $Q_\mu = 30$) uniformly throughout each velocity model domain. Our choice of Q_p and Q_s is in keeping with previous volcanic observations and simulations (Menke et. al. 1995; Tompkins & Christensen 2001; De Gori et. al. 2005; Del Pozzo et. al. 2006; O'Brien & Bean 2009). See appendix B for an expanded justification.

We used a time step of 0.285 milliseconds when forward modelling each wavefield, adhering to numerical stability criteria (see appendix B for more information).

3 FOURIER NEURAL OPERATOR MACHINE LEARNING TRAINING FOR SEISMIC INVERSION

3.1 Earthquake Gather Dataset Preparation

Having assimilated a 30,000-sized dataset of time-distance earthquake gathers and corresponding seismic velocity models of origin (see sections 2.2 and 2.3) our next task was to prepare a Fourier Neural Operator-based machine learning workflow for seismic inversion. Figure 2 illustrates the input-to-output data mapping actioned by the machine learning pipeline for our seismic imaging prediction task i.e. our primary aim was to predict a 2D P-wave velocity model from the complete seismic traces, including coda, associated with a single earthquake gather, using both the vertical and horizontal component channels of ground velocity. We divided our earthquake gathers and velocity models into a training dataset (the dataset used to tune the FNO) and a validation dataset. Here we used 90% of our total earthquake gather and velocity model databank as the training

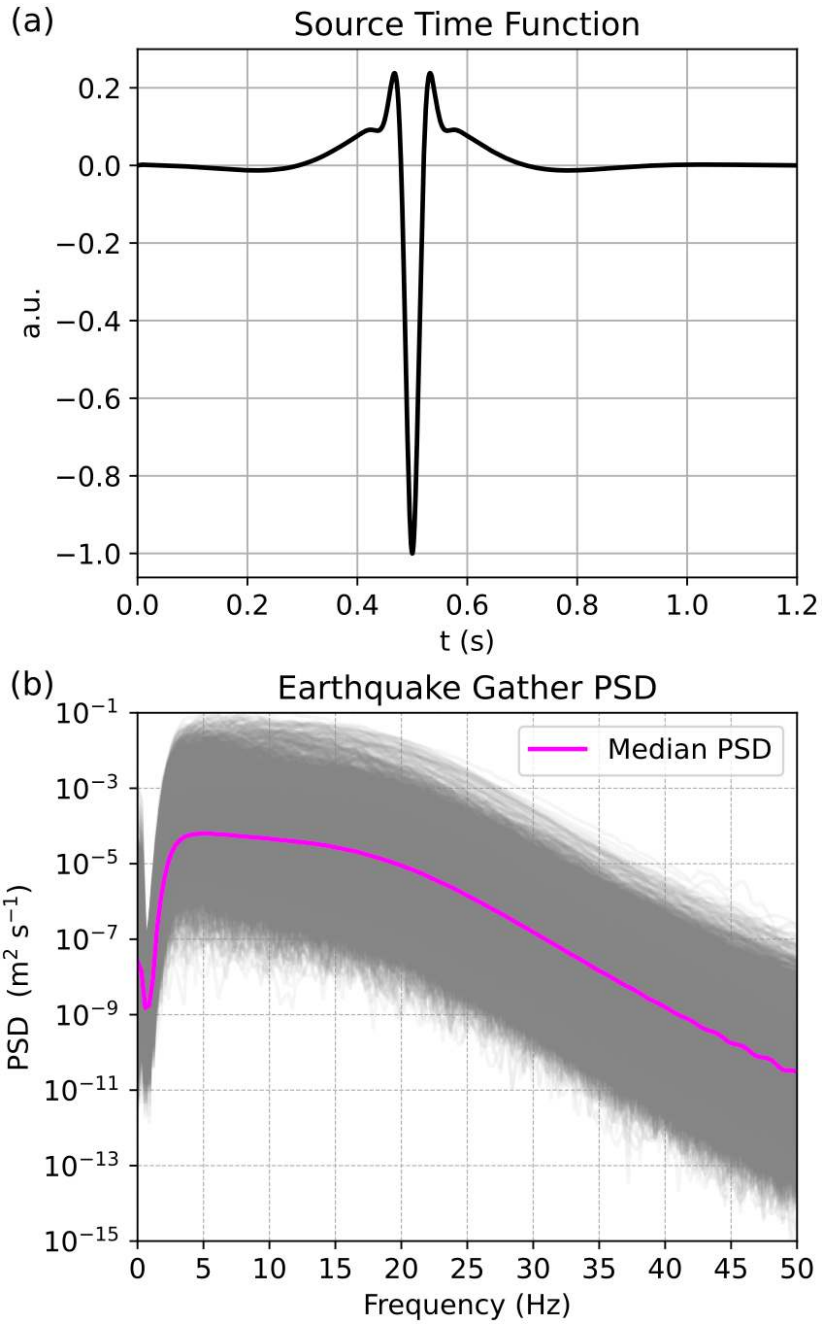

Figure 3. The broadband earthquake source time function used for each SPECFEM-generated wavefield created for machine learning training. **(a)** Delta pulse filtered between 2-25Hz **(b)** PSD plots of central seismograms from all 30,000 separate earthquake gathers modelled through each unique velocity model realisations. There are 401 seismograms per earthquake gather, with the central seismogram's PSD, recorded in each velocity model (x=5km) shown in grey. The median seismogram PSD is shown using the embossed magenta line.

set, reserving the remaining 10 % for the validation set (using a validation set size that comprises approximately 10-30% of the total databank has been used routinely in previous machine learning-based seismology other applications e.g. (Nguyen et al. 2021; Nooshiri et al. 2022; Yang et al. 2023).

Upon loading all 30,000 earthquake gathers into GPU memory, we further decimated each seismogram temporally to 75 Hz (or 450 time steps per seismogram). We decimated the earthquake gathers (the FNO model input data) spatially to 50m. We took this step because loading in tensors with larger dimensions increases machine learning model training time per epoch, as well as increasing GPU memory requirements. P-wave velocity models were not decimated and were discretised at 25m vertical and horizontal resolution.

3.2 FNO Model Workflow & Hyperparameters

Table A1 shows the main FNO model architecture used in this study and the tuning parameters for the trained FNO machine learning model are summarised in table A2. For an expanded outlay of the FNO architecture and justification for the final choice of model parameters, see appendix C.

3.3 FNO Training Results

Figure 4 shows the loss performance curve for both the training and validation datasets used to train the FNO. Here it can be seen that the training loss converged from an L2 loss per sample value of 0.00649 (after the first epoch) down to a final loss per sample of 0.00337. The validation loss per sample also converged stably from an L2 loss per sample of 0.00564 (after the first epoch) down to 0.00393. This represented reductions to 52.0% and 69.6% of initial loss per sample for the training and validation set respectively. Figure 5 shows a comparison between an endmember ground truth 2D velocity model from the validation set and its FNO-based prediction from its corresponding earthquake gather. While previous seismological applications of FNOs exhibited loss reductions of up to 70-90% (Yang et al. 2021; Yang et al. 2023; Zou et al. 2025), we note that the fine-scale velocity fluctuations created in the ground truth model by anti-persistent Von Kármán perturbation are not recovered in the FNO-based velocity model inference. Given that the FNO is

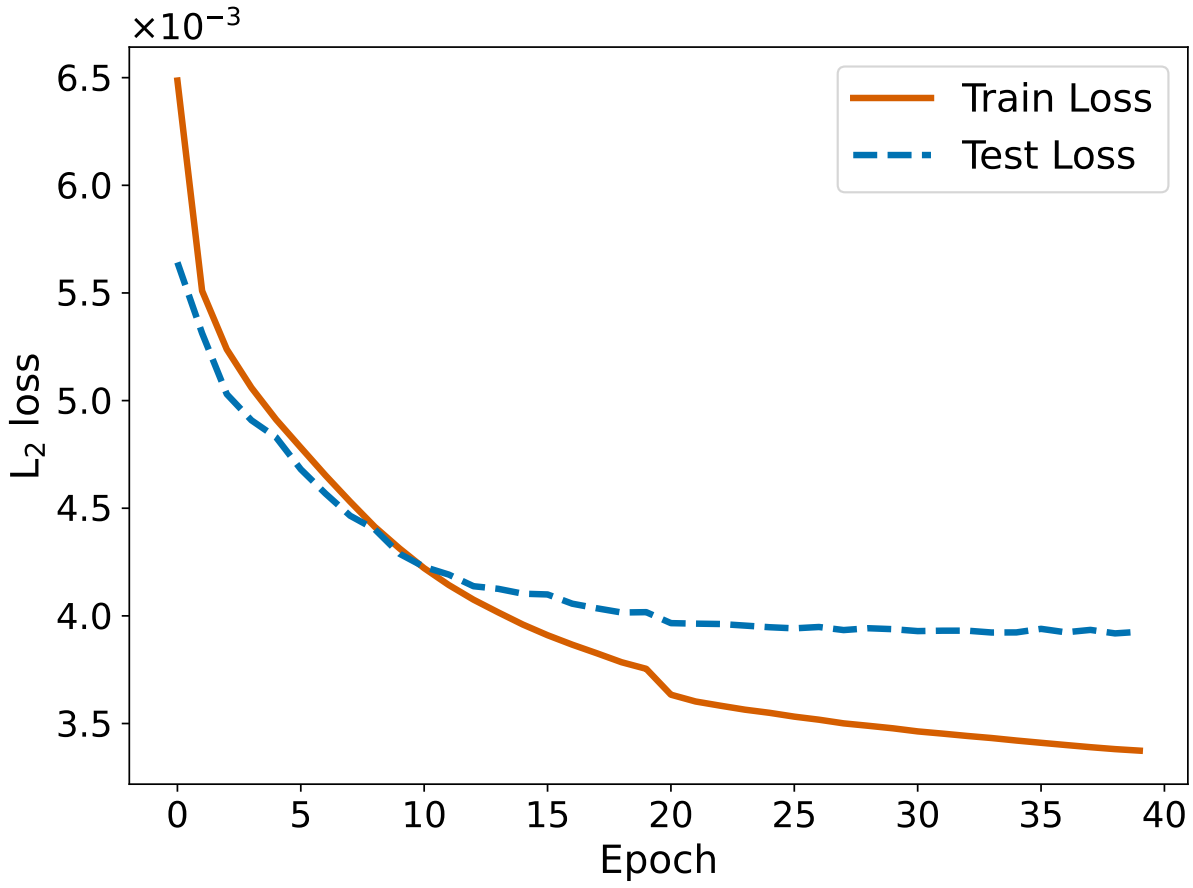

Figure 4. Loss per sample curve for the training set (orange line) and test set (dashed blue line) as a function of training epoch.

truncated at high frequencies, the roughness of the ground truth velocity model is on length scales that we do not expect to recover. This may in turn impose a cap on the fractional loss reduction seen during stable FNO training. We also note that similar fractional loss reductions have been encountered in other seismologically-based FNO studies for a forward modelling paradigm (e.g. Lehmann et al. (2023)).

In terms of the fidelity of the FNO-predicted velocity model to the ground truth velocity, it is possible to distinguish features present in both models down to scale lengths at and below the 1km correlation length. A 1D cross-sectional comparison (figure 5c) confirms that the broad-scale velocity structure has been recovered by the FNO down to as low as 250m. For example, velocity

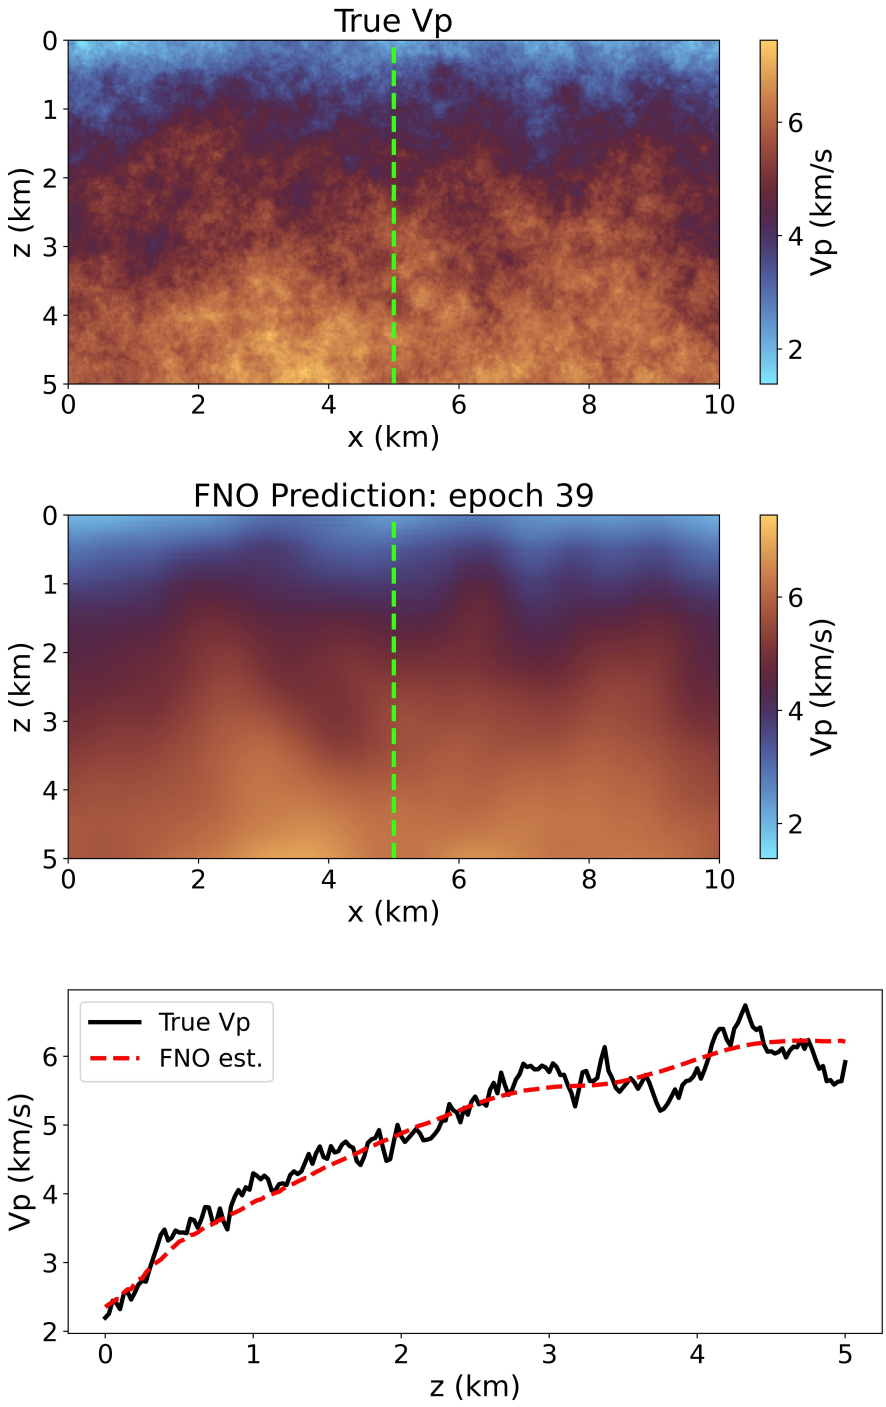

Figure 5. Velocity model prediction performance of the trained FNO model on earthquake gather data included as part of the larger validation. (a) The ground truth P-wave velocity model for the validation set. (b) The FNO predicted P-wave velocity model. (c) Cross-sections through each velocity model at x=5km, shown as a dashed black line in (a) and (b). Here the black solid line represents the 1D velocity profile through the ground truth velocity model. The red dashed line is the 1D velocity profile through the FNO-predicted velocity model.

structure highlighted in the vicinity of 1km depth in the ground truth model (figure 5a, blue-to-brown transition) is mostly recovered in the FNO prediction (figure 5b), down to ≈250m length scales. This is broadly in keeping with the length scales expected to be resolved by the seismic wavefield as a function of depth i.e. for higher expected wavefield frequency content of 25Hz and a mean velocity of 3km/s, at 1km depth in the model, one expects a seismic wavelength, λ, around ≈120m. At 4km depth, for a mean velocity of 6km/s, the minimum approximate wavelength rises to approximately 250m. It is possible that the mean free path approaches the seismic wavelength at higher frequencies, which makes traditional seismic imaging methods very difficult to perform. The mean free path defines the scale length where wave propagation transitions from quasi-direct transmission to complex scattering. For example, Calvet and Margerin (2013) predict sub-kilometre scale mean free paths for in strongly scattering media, between 3-6Hz, for both high and low Hurst exponents of 0.5 and 3 respectively. Furthermore Obermann et. al. (2014) determined that a scattering mean free path on the order of 100m is likely for Rayleigh waves in volcanic media. Indeed, Wegler & Lühr estimate mean free paths as low as 100m for S-waves from an active seismic experiment on Merapi volcano (Indonesia). While the mean free path may be contributing to the smoother velocity model predictions of the FNO, we note that the cap in the number of Fourier modes used and the FNO architecture (i.e. the number of Fourier operator layers and the dimensionality of channel uplift), imposed here for reasons of computational limits, likely plays a significant role. It is expected that finer model resolution could be achieved using a larger FNO network and higher Fourier mode truncation (the universal approximation theorem of neural operators states that an FNO can approximate any well-defined map given sufficient network complexity and an adequate training dataset (Kovachki et al. 2021)).

It is important to stress here that the trained FNO model object is able to adequately recover velocity structure from a single earthquake gather when using a line array spacing of 50m (i.e. 201 receivers over the whole 10km lateral domain extent). Through source-receiver reciprocity we would expect to achieve the same performance as shown in this study for fewer receivers, provided that there is sufficient seismicity. This latter scenario is desirable for future 'simulation-to-real' FNO applications where a large number of receivers is not necessary, provided that high

rates of distributed seismicity are present. This schema is plausible for terrestrial volcanic environments and is highly likely for oceanic volcanic environments, where one would expect to deploy relatively fewer seismic instruments.

4 FNO SEISMIC INVERSION APPLICATIONS

Having shown that an FNO can be trained to invert for 2D velocity models from earthquake gathers recorded in strongly scattering synthetic media, in this section we investigate the ability of the trained FNO model to recover geological features of interest. Here we focused on sub-vertically oriented columns, which were intended to emulate the velocity signature of magmatic intrusions such as dykes. Firstly, we diversify the FNO training dataset by incorporating a small number of additional velocity models embedded with specific geological features, along with their corresponding earthquake gathers. These account for approximately 5-15% of the total dataset size and can vastly improve feature recovery for unseen seismic data. Secondly, we stress test the resolution recovery limits of the FNO applied to earthquake gathers generated in strongly scattering velocity models embedded with dyke-like volcanic features. Thirdly, we investigate the ability of a trained FNO model to recover those same features when the earthquake gathers are embedded with synthetically-generated environmental noise. As a completeness exercise, we close by testing the FNO's ability to recover migrating, propagating dyke features.

4.1 Fine-tuning the FNO for enhanced recovery of magmatic intrusions

In this section we show how fine-tuning the core velocity model and earthquake gather training datasets can be used to improve geological feature recovery when the majority of velocity models used during FNO training are dominated by marked velocity gradients. Previous FNO-based seismology studies have shown that geological features of interest not included in the training set can be recovered when applying an FNO in inference mode. This applies to both the forward problem (e.g. Yang et al. (2021) successfully estimated waveform codas from velocity models not used in training) and the inverse problem using reverse-mode automatic differentiation. However, many of these studies adopted a more generalised approach to the population of velocity models

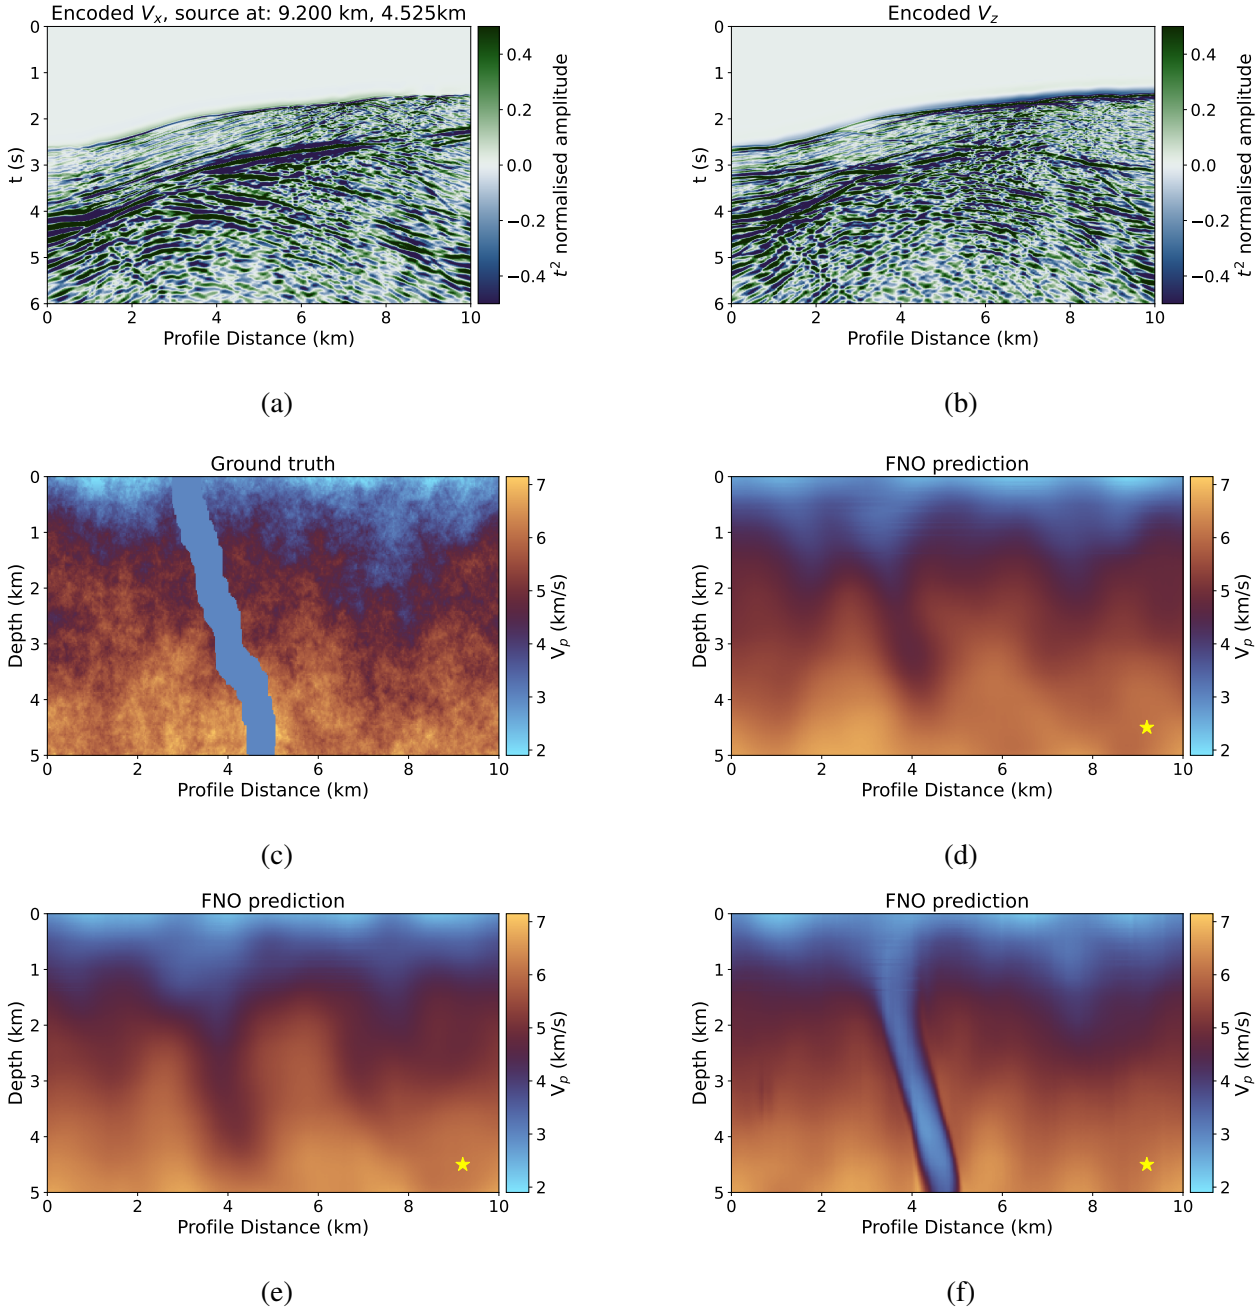

Figure 6. Appraisal of trained FNO velocity model prediction performance from an unseen earthquake gather using both core and fine-tuned datasets. Shown are results for a single earthquake gather located at x=9.2km, z = 4.53km. **(a)** X component of earthquake gather recorded by receivers on the top of the model domain, which has been encoded using Z-score normalisation **(b)** Z component for the earthquake gather with same normalisation **(c)** the ground truth velocity model comprised of; a von Kármán medium with depthwise positive gradient, and a 3km/s, 500m wide low velocity zone **(d)** predicted velocity model using an FNO trained on 30,000 core von Kármán models only **(e)** predicted velocity model using an FNO trained on the same 30,000 core von Kármán models plus an additional 1,000 fine-tuned velocity models with inverted velocity graident **(f)** predicted velocity model using an FNO trained on the same 30,000 core velocity models plus an additional 2,000 von Kármán-based models embedded with sub-vertical low velocity zones (3km/s) and approximate columnar width of 500m.

and waveforms. This included the use of Von Kármán series with a mean and standard deviation that covered the full range of scattering body velocities anticipated to be encountered at the time of FNO application. Other studies used a large family of depthwise velocity gradients with additive noise. Here we have highly tailored the velocity model population to average properties in Icelandic volcanic upper crust, both in terms of velocity gradient and the statistical scattering properties expected for a volcanic environment. One paradox that arises from this approach is that the magmatic intrusions that we seek to image can deviate markedly from these average volcanic velocity model properties. By extension, it is unlikely that scattering phases in earthquake gather codas, related to reflection and transmission of the wavefield through such intrusive features, have become a learned feature of the FNO. In light of this, we made the decision to introduce additional diversity in velocity models and earthquake gathers to the core FNO training set.

4.1.1 Updated FNO training with inverted velocity gradients

The first new population of velocity models generated to supplement the existing core population was embedded with the same Von Kármán perturbations (H=0.3, a_x=a_z=1000m, σ=$\pm$350m/s) but with an inverted 1D velocity gradient. The objective was to further generalise the training in an effort to account velocity inversion for slow velocities, deep in the model. We created 1,000 instances of such models and modelled the SPECFEM wavefield as per previous waveform modelling steps, recording the earthquake gather along a surface-based array. These 1,000 additional velocity models and gathers were then added to the existing 27,000 sized training dataset before retraining the FNO from scratch, resulting in 28,000 training earthquake gather and velocity model pairs. We repeated FNO training using this larger, more diverse population and then re-applied the updated FNO model to unseen (test) earthquake data, generated in a model with the same statistics as the training set, but that included an embedded 3km/s dyke-like feature. Figure 6e shows an updated prediction of the velocity model for a given shot location using the updated FNO model. When compared with the original prediction, it can be seen that adding 1,000 additional inverted velocity models to the training set had a minimal effect on the dyke feature recovery, as well as minimal change in the predicted background model. From this we conclude that diversifying the

training set using velocity models with inverted gradients is insufficient to improve FNO-based recovery of 'small-scale', specific geological features.

4.1.2 Updated FNO training with volcanic dyke intrusions

The next test undertaken was to directly include embedded dyke features in velocity models and to incorporate these along with their earthquake gathers into further FNO training updates. This time we created 1,000 such models with one new dyke embedded in each model realisation, each of which was a 500m wide, sub-vertical 3km/s feature, extending from the base to the top of each model. The path of each column through the model was determined by a Brownian walk with imposed lateral translation limits for each step upwards towards the surface. From here we modelled an earthquake gather corresponding to each velocity model realisation, each of which was assigned a source location drawn from a uniform random distribution across the entire 10km x 5km model domain. We note that the majority of seismicity in volcanic settings such as Iceland is consigned to depths below approximately 2km. Here we decided to leverage information from all possible source depths to maximise FNO sim-to-sim generalisation. After updating the FNO training set with these models and earthquake gathers, we observed that adding just 1,000 such diversified samples to a now 32,000 sized dataset vastly improved dyke feature recovery from using just a single earthquake gather. This can be seen in figure 6f where the left-dipping dyke can be clearly distinguished throughout the entire model. In other words, we have shown that recovery of a velocity model from a single earthquake gather is possible when an FNO is trained with the right population. Further to this, Figure A3 shows how randomising the earthquake gather data tensor corresponding to figure 6f results in a similarly scrambled velocity model output, showing that FNO velocity model recovery is not accidental.

4.1.3 Stress testing FNO inversion recovery resolution

While the demonstrated enhanced recovery of low velocity intrusions is encouraging for future field-based FNO applications, we acknowledge that individual dykes in the physical world can range from the order of 100m down to 1m in cross-sectional width. In what follows we stress test

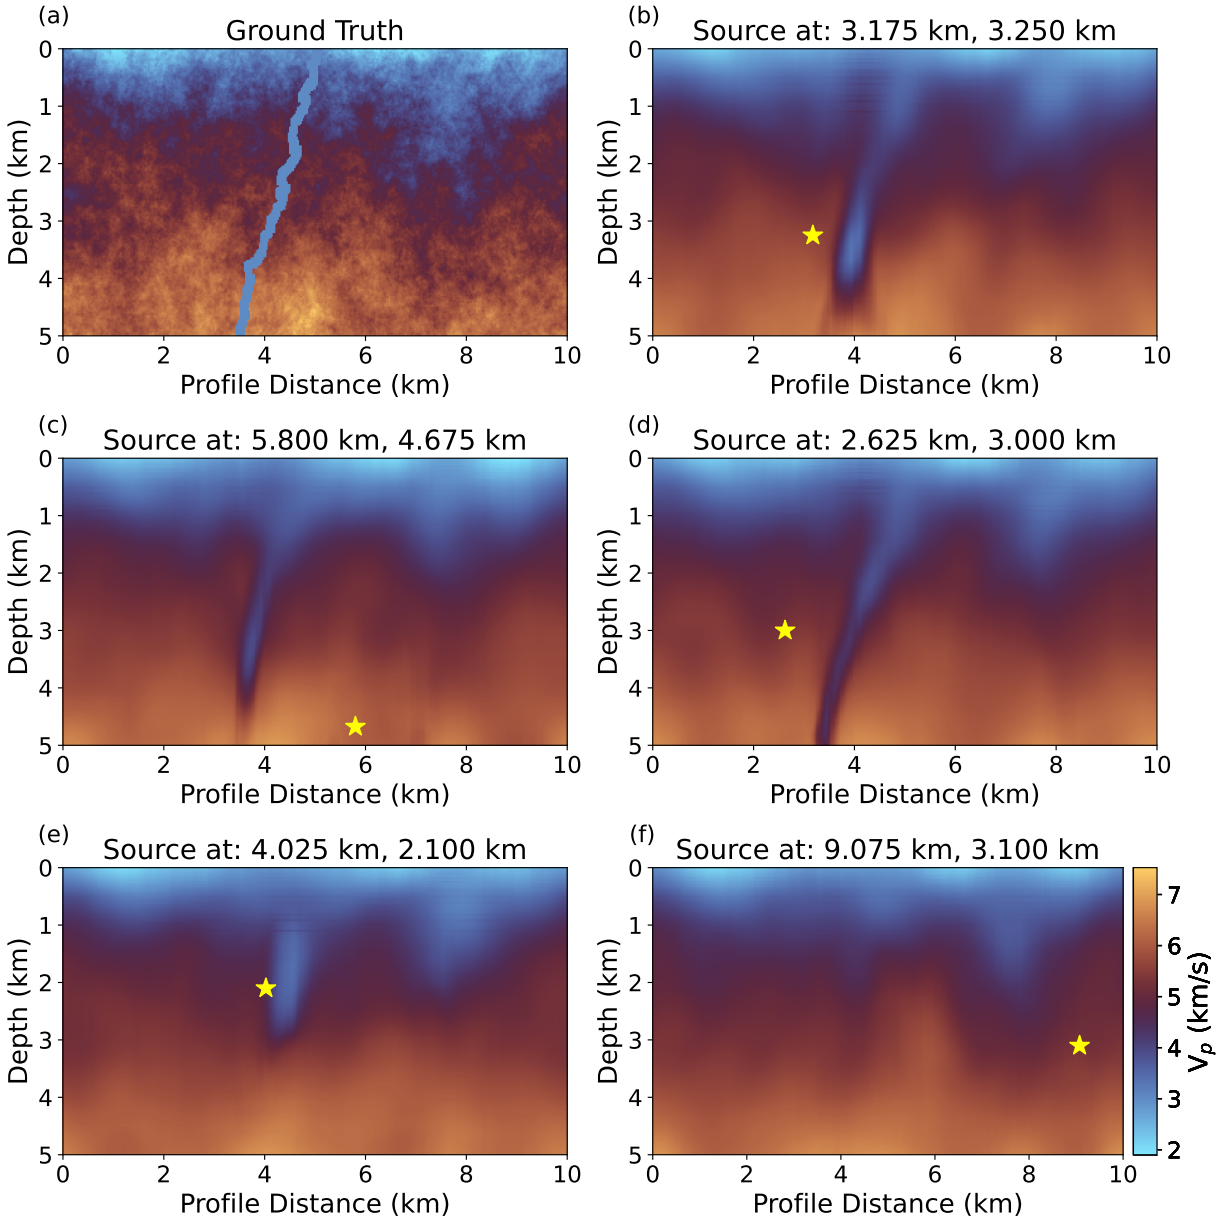

Figure 7. Stress testing the resolution limit of the trained Fourier Neural Operators shown in this study. **(a)** Ground truth velocity model with an embedded ≈ 3km/s low velocity feature with minimum cross-sectional width ≈200m **(b)** FNO prediction from a single earthquake gather recorded for a source located at x=3.175km, z=3.25km **(c)** FNO prediction from a single earthquake gather but with source located at x=5.8 km, z=4.675km **(d)** x=2.625 km, z=3km **(e)** x=4.025km, z=2.1km **(f)** x=9.075km, z=3.1km.

the resolution limits of recovery of our trained FNO model and introduce additional complexity into the velocity anomalies intended to represent magmatic intrusions. To this end we created 2,000 further volcanically representative velocity models embedded with dykes. This time, we allowed

the P-wave velocity inside the dyke to be assigned a value from a truncated normal distribution ranging from 2.5km/s to 3.5km/s. We also freed up the Vp/Vs ratio. Again we assigned a Vp/Vs ratio for each dyke in the updated training from a normal distribution with mean of 1.73 ($\sqrt{3}$) and one standard deviation of $\pm$0.05. The distribution was clipped at 1.73$\pm$0.13, resulting in a permitted Vp/Vs ratio range of 1.6 to 1.86. We chose this range based on what has been observed in the field during Icelandic dyke swarms. For example, Masihi and Fischer (2025) performed spatial analysis of the Vp/Vs ratio along the 2021 Fagradsfjall dyke swarm where they measured Vp/Vs clustering in the range of 1.6-1.76 for both pre and post-eruptive seismicity. Here we also included higher Vp/Vs ratios for maximum FNO learning generalisation. The final free parameter for each dyke was the permitted thickness. Here we used a left-skewed distribution. While we permitted dyke thicknesses as large as 400m, the modal dyke thickness in the distribution was 175m. The majority of the thickness distribution was concentrated at thicknesses less than this value, down to a minimum of 75-100m.

Having assigned a thickness, Vp and Vp/Vs ratio to each new dyke in the updated training set, each of which was embedded in a unique Von Kármán background velocity model, our next step was to include them in updated cycles of FNO training. To do this, an earthquake gather was generated for each new velocity model using SPECFEM, with source locations drawn from a uniform random distribution. The revised combined training and test set contained 34,000 earthquake gathers (inputs) and 34,000 corresponding velocity models (outputs) with the following statistical properties. The first 30,000 earthquake gathers were generated in velocity models with depthwise increasing velocity gradients perturbed by Von Kármán series; the next 1,000 earthquake gathers were generated in models with inverted (i.e. depthwise decreasing) velocity gradients; 1,000 further earthquake gathers corresponded to Von Kármán-perturbed models with depthwise positive velocity gradients and embedded dyke-like features with $\approx$500m width, and the remaining 2,000 gathers were generated in models embedded with dykes, drawn from thickness and Vp/Vs distributions, as described above. We then tested the updated FNO model on unseen earthquake gather data generated in velocity models embedded with individual dykes of thickness on the 100m-200m scale.

4.1.3.1 Recovery of Background Velocity Models Figure 7 evidences the robust ability of the FNO to recover the background velocity model, despite the addition of a dyke with different statistical properties, for different source locations. As a case in point, it can be seen that the topography along the approximate 4.5-5km/s velocity isocontour is spatially consistent across multiple independent, single earthquake gather FNO inferences. This suggests that the FNO is able to discriminate between scattered arrivals, related to the dyke and background model respectively, contained in a given earthquake gather.

4.1.3.2 Recovery of Dykes Figure 7 further highlights the recovery of a ≈200m thick dyke embedded in an ambient volcanic background model, for a variety of source locations, using the FNO inference from individual earthquake gathers. Here it can be seen that the single gather recovery is good for sources positioned within a lateral distance of a few kilometres relative to the dyke, across a range of source depths (figure 7 b-d). Figure 7e shows recovery from a source immediately adjacent to the dyke at 2.1km depth, where it can be seen that the low velocity anomalies have not been recovered in the bottom 2km of the model. There are several ways to explain this recovery, all of which are related to seismic wavefield sensitivity, or source-receiver kernel coverage (i.e. similar to the rationale shown in figure 12). Firstly, the strong impedance contrast between the ambient velocity model and the dyke results in high amplitude reflected waves dominating the earthquake gather, for receivers to the left of the low velocity anomaly, originating in the vicinity of the source location. Secondly, it is possible that for wavefields related to relatively more distant sources, the dyke has been sampled along a greater proportion of its length, resulting in a greater variety of reflected and transmitted arrivals in the earthquake gather, which are learnable features by the FNO. In figure 7f we observe that the dyke is poorly recovered for an earthquake gather corresponding to a distant source location. One possible explanation for this result is that high frequency wavefield attenuation for a very distant source renders it insensitive to the relatively thin low velocity anomaly. Alternatively, a relatively small travel time anomaly accumulated by the attenuated wavefield from the highly localised dyke may be distributed over the entire velocity model during FNO inversion. Nonetheless, figures 7b)-d) show that it is possible to recover dyke features along almost their entire length, along with their dip and approximate thickness. For fu-

ture simulation-to-real applications, one could envisage seeking consistent feature recovery from a population of models, derived from many earthquake source locations.

4.2 FNO Environmental Noise Resilience

Having shown that i) an FNO can be trained in a sim-to-sim paradigm for the task of seismic inversion for subsurface velocity models, when shown a single earthquake gather that contains signal from multiscale scattering ii) the FNO can recover geological features of interest (e.g. dykes) embedded in highly scattering velocity models designed to emulate volcanic environments, we now address a potential primary limiting factor for the trained FNO's sim-to-real performance in future work: environmental noise. In this study we trained the FNO to predict velocity models on noise-free, synthetic earthquake gathers. However, it is widely known that environmental noise is ubiquitous across seismic data recorded in volcanic environments. In light of this, seismic machine learning applications in volcanic environments need to demonstrate sufficient noise-related performance resilience. For example, Nooshiri et al. (2022) produced a noise model to match observational noise in the Hengill geothermal field, Iceland, when developing a neural network for point-source inversion from noise-augmented seismic data. In this section we stress test the noise-based resilience of our FNO, trained to invert for seismic velocity from noise-free data. In order to do this, we first choose a noise-free earthquake gather which resulted in an acceptable velocity model performance by the FNO relative to the ground truth velocity model, as shown in figure 8. The next step is to repeat the FNO predictions using the same earthquake gather, contaminated by simulated environmental noise with progressively lower SNR. Here we defined signal-noise ratio (SNR) as

$$SNR = \frac{\sigma_{x(t)}}{\sigma_n} \tag{4}$$

where $\sigma_{x(t)}$ is the standard deviation of the earthquake gather signal for a 1s time window around the first arrival and σ_n is the one standard deviation value of the environmental noise. Normally, one has acquired seismic waveforms from the physical world and quantifies the SNR based on a relation similar to the above. Here we must reverse engineer environmental noise into the SPECFEM forward-modelled earthquake gather of interest. In order to generate a noise field to this end we

use 2D Gaussian noise, setting its standard deviation relative to the standard deviation in an approximately 1s time window around the first arrival in the signal, to control the SNR. We pick the baseline value for standard deviation of the signal based on values taken across the length of our earthquake gather of choice.

Table 1 shows the different SNR values for the x and z component of the earthquake gather corresponding to figure 8 with a brief description of FNO velocity model prediction quality. For an SNR of 100 (figure 8b) we did not observe a significant depreciation in quality of either the background velocity model or the dyke, relative to the baseline recovery from the noise-free synthetics. For an approximate SNR of 50 (figure 8c) there appears to be no depreciation in dyke recovery at all, however, we observe that the background model recovery in the rightmost 3km of the model (7-10km) is reduced. This makes intuitive sense given that the known source location for this earthquake is located at a lateral position of x=1.35km i.e. in the leftmost portion of the model. It is therefore expected that signal information about the velocity model at greater distances from the source is contained in later time steps in the earthquake gather, where one would expect the noise field to pollute the coda signal to a greater extent. For an SNR of 33 (figure 8d) this rightmost portion of the background model further reduces in quality again and the entire background model looks to be affected by the noise at this point. In contrast, the dyke feature is still present. For an SNR of 25 (figure 8e) or lower the background velocity model is entirely degraded and the velocity values inside the dyke begin to tend towards the background model average. For this particular earthquake gather, we conclude that the FNO produces inferences of acceptable quality down to an SNR of 50. While this value seems to be relatively high, earthquake catalogue analyses from volcanic environments show that there are often a significant number of local earthquakes with SNR above this threshold. For example, the Krafla earthquake catalogue collated by Maass et al. (2025) yielded multiple events with SNR in the region of 50 and approximately 7 events with SNR greater than 100 to 1000, based on the definition given in 4. Furthermore, these events exhibited local magnitudes in the range of -0.5 to +1.5.

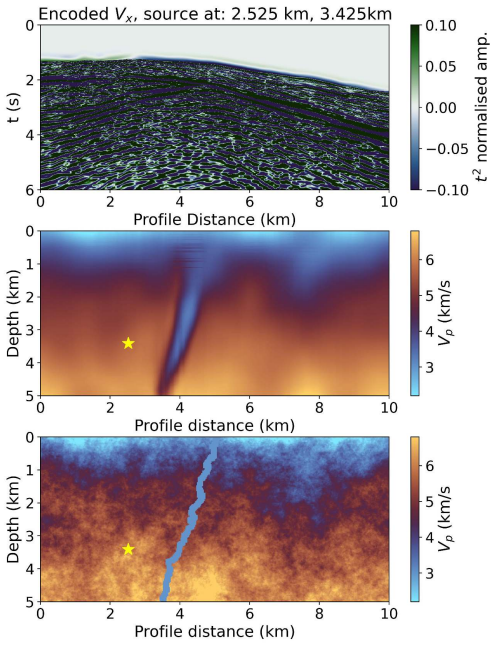

(a)

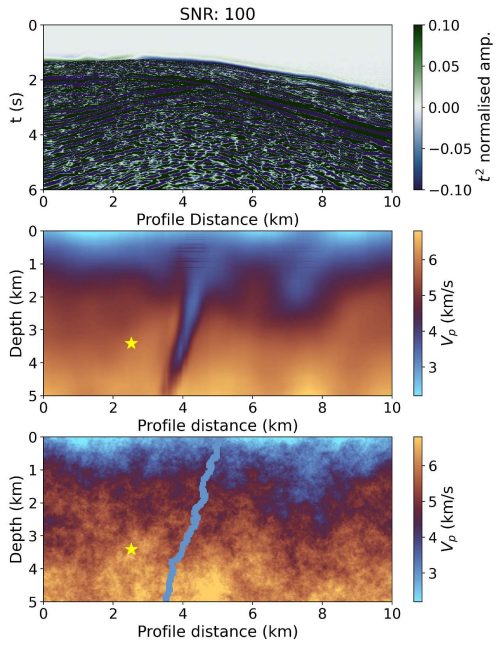

(b)

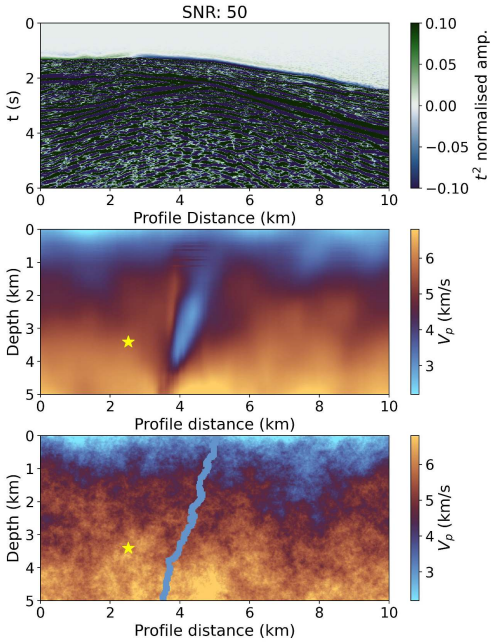

(c)

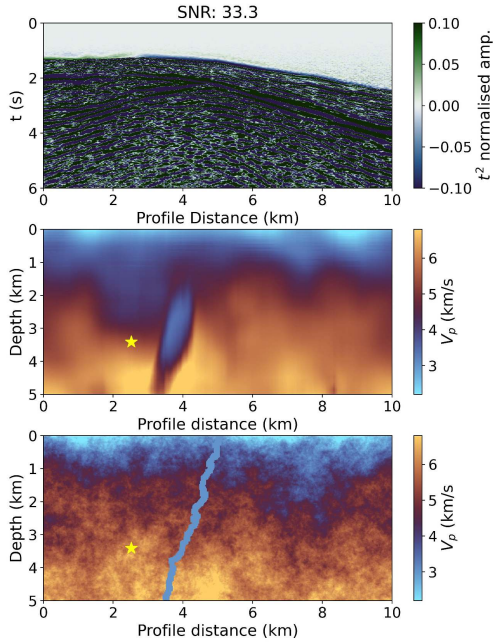

(d)

Figure 8. Demonstration of FNO resilience to synthetic environmental noise. **(a)** (top) X-component velocity earthquake gather for a noise free event at x=2.525km, z=3.425km (middle) the FNO velocity model prediction pertaining to the same earthquake gather, with the source location marked as a yellow star (bottom) the ground truth velocity model with ≈200m thick 3km/s dyke **(b)** FNO performance for the same earthquake gather, augmented with 2D Gaussian noise with approximate signal-to-noise ratio (SNR) of 100 **(c)** same as b but with SNR = 50 **(d)** SNR = 33.

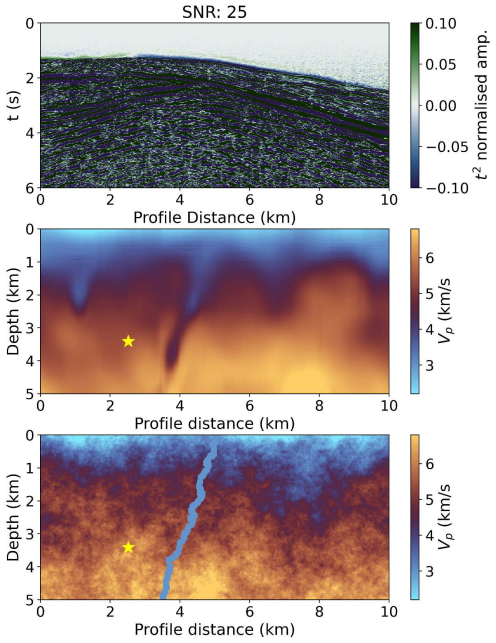

(e)

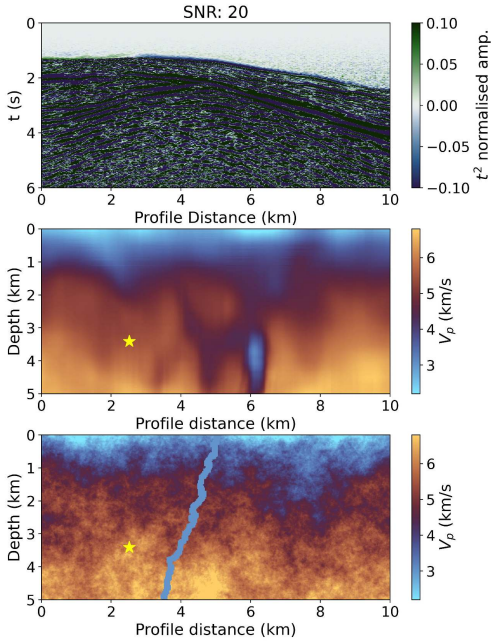

(f)

Figure 8. contd. **(e)** SNR = 25 **(f)** SNR = 20.

4.3 FNO Performance using frequency-ablated earthquake gathers

The previous section demonstrated the tolerances of the trained FNO when it was fed earthquake gathers with progressively lower signal-to-noise ratios. Here we perform frequency ablation tests on the input seismic data to analyse the contribution of different frequency components contained in the earthquake gather to the overall FNO prediction. The test is also intended to show that the FNO is not excessively reliant on low frequency earthquake gather content relative to higher frequency scattering information. Figure 9 shows the FNO performance for a single earthquake gather shown previously in figure 8a based on different low pass filters of the data input to the FNO. Each subfigure shows the velocity model prediction given the frequency-ablated earthquake gather and the difference between this prediction and the prediction based on the unfiltered earthquake gather. For low pass cutoff frequencies of 2Hz and 3Hz (figures 9a and b) the background velocity model remains more or less intact perceptually (albeit with non-insignificant amplitude differences) whereas the dyke is poorly recovered. When the gather is filtered using a 5Hz (figure

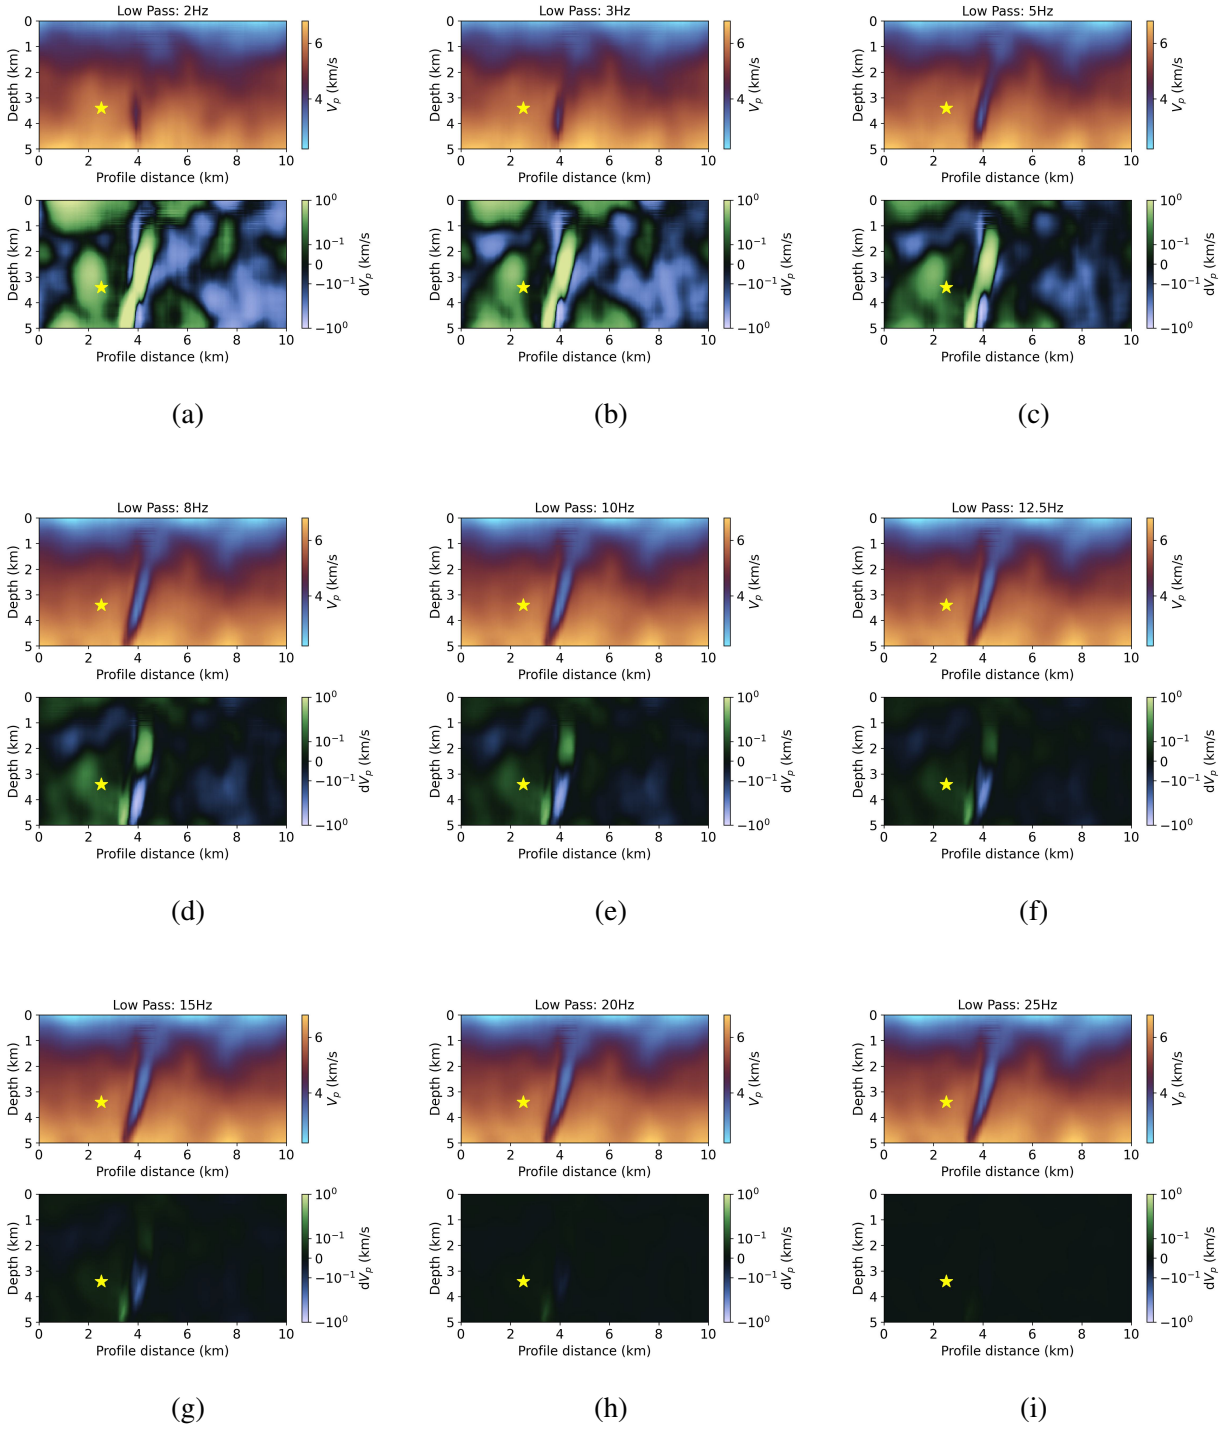

Figure 9. Frequency ablation test using different low pass filter frequencies on the earthquake gather shown in figure 8a before feeding the filtered gather to the FNO. The top panel in each figure shows the FNO-predicted velocity model when taking the low-passed earthquake filter as input. The bottom panel show the difference between this prediction and the original FNO prediction based on the unfiltered earthquake gather. **(a)** Low-pass filter frequency of 2Hz (b) 3Hz (c) 5Hz (d) 8Hz (e) 10Hz (f) 12.5Hz (g) 15Hz (h) 20Hz (i) 25Hz.

9c) cutoff frequency the dyke recovery is vastly improved. For cutoff frequencies of 8Hz, 10Hz and 12.5Hz respectively (figures 9d-f) the structural recovery of background model and dyke looks perceptually similar but amplitude residuals remain relative to the prediction based on the unfiltered earthquake gather. Figures 9g-h show that the inclusion of frequencies between 15-25Hz are in fact required to achieve like-for-like FNO velocity model predictions relative to the velocity model based on the unfiltered input. One interpretation of these frequency-ablation results is that the limiting factor for the FNO velocity model recovery resolutions shown in this paper are likely due to the capped modes and width of the FNO. Re-training such an FNO with a higher number of FNO Fourier modes is therefore a good opportunity for future studies. We also stress that it is important to avoid definitively attributing results from this frequency-ablation test solely to geophysical interpretation. It is more likely that the closer a given frequency-ablated earthquake gather is to the distribution of earthquake gathers used to train the FNO (i.e. quasi-broadband frequency content), the better the velocity model prediction.

4.4 Time-lapse imaging with FNO Engine

Thus far we have shown the ability of an FNO to perform the seismic inverse problem using earthquake gathers with high frequency codas, with a focus on minimum feature resolution recovery, as well as the FNO's prediction resilience when working with noise-embedded seismograms. In this section we attempt to showcase the potential of future FNO applications in volcanic settings for real-time imaging and how they can complement conventional seismic imaging methods, as well as emerging methodologies such as fibre-optic geodesy (Li et al. 2025). Figure 7 showed that our workflow is capable of recovering sub-vertical, low-velocity features that are embedded in heterogeneous, highly scattering background models, which extend from the base to the top of the model. To demonstrate time lapse FNO capabilities, we created a further 1,000 von Kármán perturbed 2D velocity models with the same statistics and depthwise background velocity gradient as shown in all previous sections of this paper. From here we embedded low velocity anomalies into each new velocity model realisation, with new characteristics. This time, each unique low velocity zone instance extended from the base of the 2D model domain at a subvertical angle be-

fore terminating at an intermediate depth in the model. These features were intended to resemble dyking events or other similar magmatic intrusions into host volcanic country rock, at progressing stages of upward migration towards the surface. For each unique low velocity zone occurrence, we constrained the termination depth by drawing a value from a uniform distribution that ranged from 4km depth to 500m depth, in 500m decrements. For simplicity we fixed the lateral width of each unique low velocity zone to 500m. We then modelled all corresponding earthquake gathers for each of the 1,000 new velocity model realisations, with a source location drawn from a uniform random distribution. Again, FNO retraining was implemented on the updated earthquake gather and velocity model dataset. We then tested the updated FNO model on 100 unseen earthquake gathers, which were modelled in von Kármán based velocity models with embedded low velocity anomalies terminating at different depths beneath the surface of each model's domain. Figure 10 shows the feature performance recovery achieved for a subsample of dyke-like features at different intrusion depths, for single earthquake gather-based recovery. Here it can be seen that the recovery of each unique feature is more than acceptable for sources located within $\approx$1km distance of the dyke. This is conducive to physical world applications, given that seismicity locations often cluster in the vicinity of fissures and fractures, during dyking events.

4.5 Demonstrating wavefield coda contributions to the overall FNO velocity model inference

As a final completeness exercise, we present FNO performance results for a single earthquake gather inference, muted for different time windows. This test was designed to analyse the ability of the FNO to reconstruct a velocity model image using only parts of the wavefield coda. This was carried out by zeroing out different time windows of a given earthquake gather before feeding it to the FNO to assess the change in velocity model prediction performance. For this test we used the ground truth velocity model and its corresponding earthquake gather shown in figure 10b, with the following main observations:

- 4-5s (figure 11c) - a significant part of the coda is muted. Although present, the dyke is mislocated. The background velocity model is of borderline quality with significant artefacts.

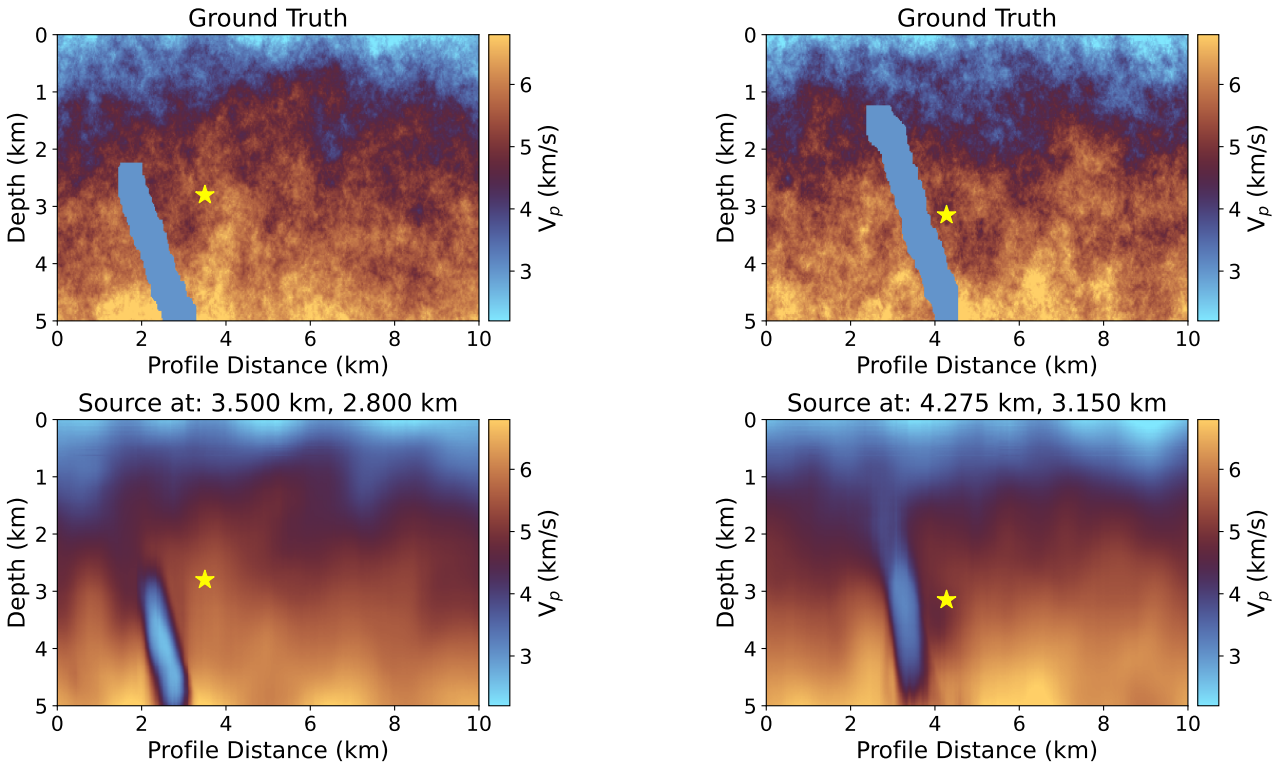

Figure 10. Motivating the application of the FNO in a time-lapse context. Shown are single earthquake gather inferences for four unique, training set-independent velocity model realisations. Each velocity model contains a unique dyke at different stages of intrusion into the background model. **a)** (top) ground truth P-wave velocity model from a gather for an earthquake source located at x=3.5km, z=2.8km (bottom) the corresponding FNO-predicted velocity model **b)** Same as a) but for another unique velocity model and dyke, with earthquake source at x=4.275km, z=3.15km.

- 5-6s (figure 11d) - here the lattermost coda is missing. The FNO inference is comparable to the full 0-6s inference shown in figure 10c.

From the above observations, we deduce that the FNO model in its current trained state performs best when shown both the ballistic arrivals and coda. The full coda plays an integral role in constructing the most high fidelity velocity model inference possible for the FNO in its current trained state. For the velocity model domain sizes presented here, using earthquake gathers between 0-5s is sufficient for recovery. In summary, feeding earthquake gather data to the FNO, which contains ballistics and only part of the coda, results in blurry images akin to classical tomography, whereas feeding a gather containing ballistics and a complete coda extends the model

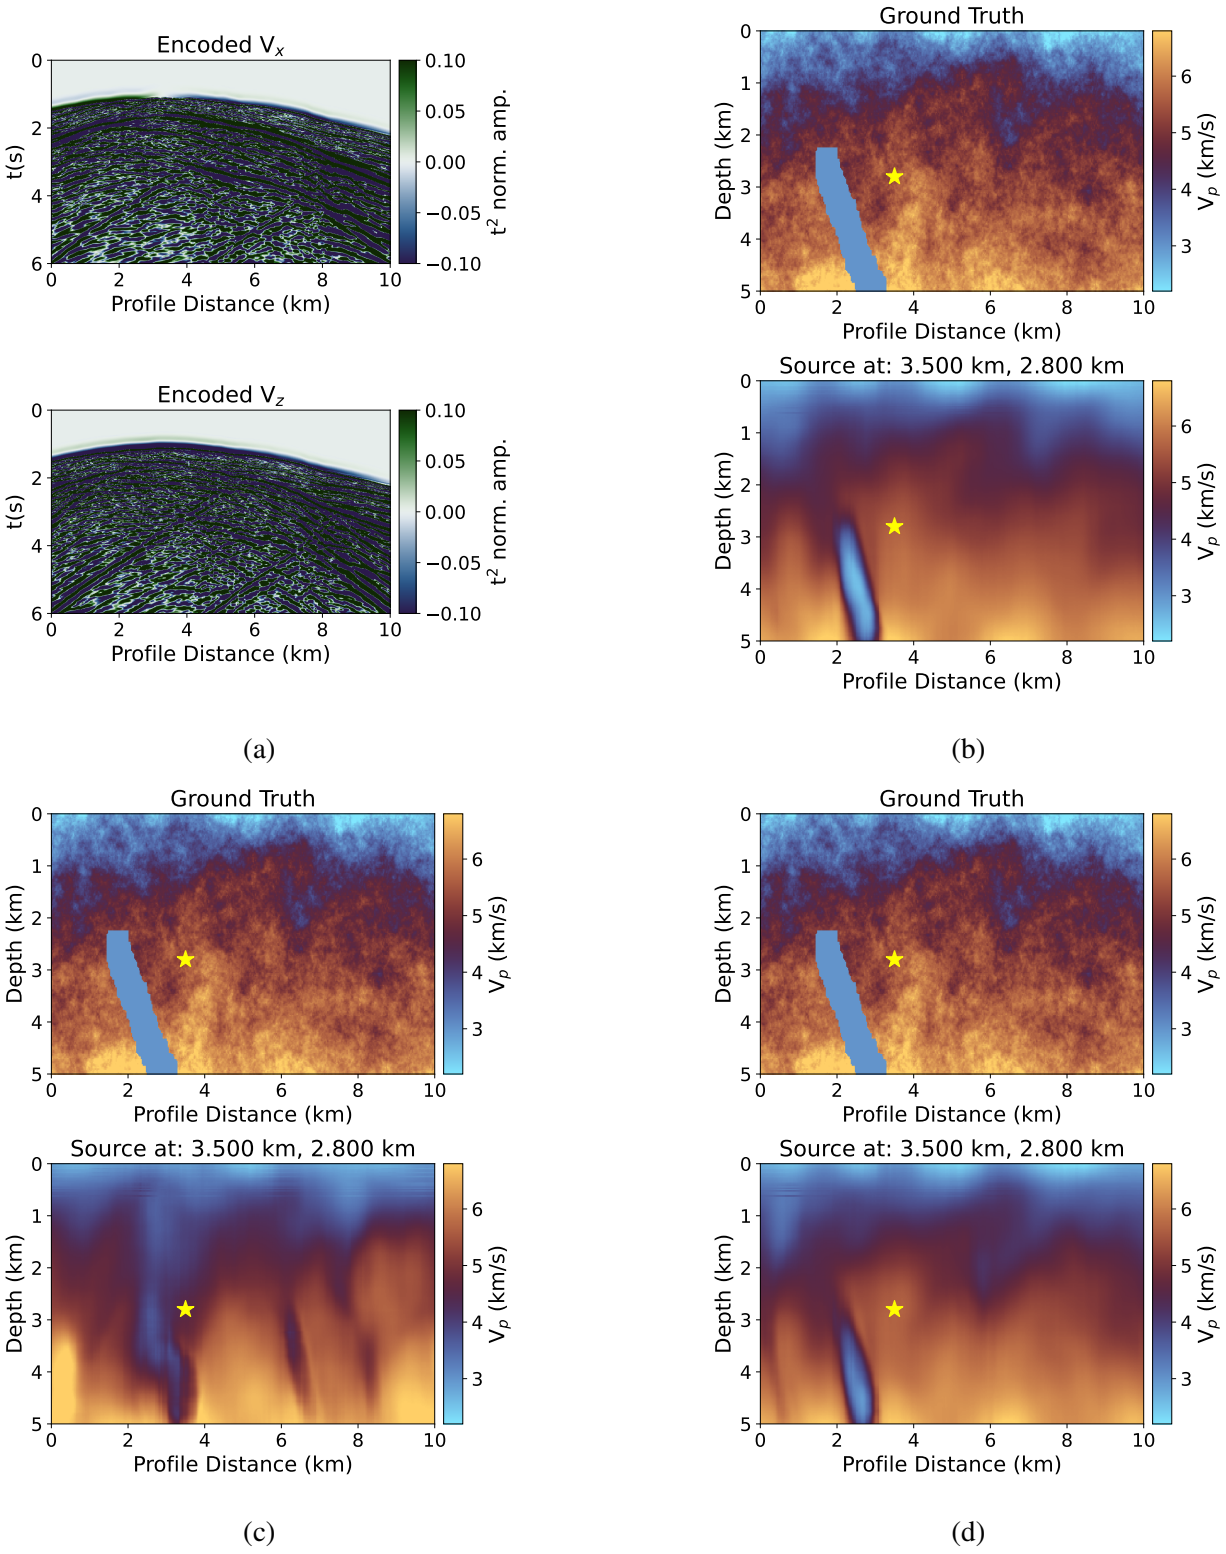

Figure 11. FNO performance with muted earthquake gather coda for different time windows. Shown here is the same single earthquake gather inference as shown in figure 10a, repeated for different gather muting time windows. **a)** The original unmuted and encoded earthquake gather with both the x and z velocity components shown for reference **b)** FNO inference when the coda has not been muted at all **c)** muted between 4-5s and **d)** 5-6s.

Test	SNR_x	SNR_z	Model Qual.	Feat. Qual.
a	noise free	-	-	-
b	100	100	good	good
c	50	60	medium	good
d	33	33	medium	good
e	25	25	low	medium
f	20	20	low	low

Table 1. Qualitative FNO prediction quality as a function of decreasing signal-to-noise ratio. Test indexes a-f correspond to labels in figure 8. Model Qual. refers to the quality of the recovery background velocity mode and feat. qual. refers to recovery quality of the embedded geological feature (dyke).

recovery to higher resolution. Given that 5 seconds of data fully captures a velocity model of this scale, it is possible that the lattermost 1 second of seismic coda is much lower in amplitude due to wavefield attenuation. This can be examined using sensitivity kernels. In the following section (5.1), we use sensitivity tests to show that coda waves are required to fingerprint the entire velocity model domain. See appendix E for an expanded presentation of how truncating different earthquake gather time windows impacts FNO velocity model prediction results, as shown in figure A4.

5 DISCUSSION

5.1 Quantifying Whole Model Wavefield Sensitivity

In order for an FNO machine learning model to converge towards a successful training state and to generalise to unseen data, it is important that each unique earthquake gather contains a comprehensive set of features that can be mapped to its corresponding velocity model of origin. Furthermore, there must be sufficient learnable feature diversity across the combined training and validation input datasets. Here we have sought to invert surface-recorded earthquake gather data for seismic velocity, using both the phase and amplitudes of both ballistic arrivals and coda waves. The inclusion of coda waves in the inversion procedure is necessary for better velocity model images. It is therefore important to demonstrate that our SPECFEM-modelled wavefields contain sensitivity across the entire domain, corresponding to their velocity model of origin. Specifically, we seek

to quantify the spatially-dependent velocity model sensitivity contained in different parts of the ballistic and scattered wavefield. To do this, we computed travel-time sensitivity kernels based on the method of Tromp et. al. (2005) and implemented in SPECFEM. Martínez Montesinos et. al. (2021) previously used sensitivity kernels to explore the relationship between different seismogram packets and structural models of Mt. Etna. Sensitivity kernels express the Fréchet derivative of seismic travel-times relative to changes in seismic velocity model parameters. They are calculated by forward-modelling the source-to-receiver full wavefield with its adjoint (time-reversed) receiver-to-source wavefield, integrated over the time window of interest. We direct the inclined reader to the adjoint methodology used in Tromp et. al. (2005) for the derivation behind the kernels used here.

Figure 12 shows spatially-stacked P-wave travel-time kernels based on different earthquake gather time windows, for a given velocity model realisation. We retained P-wave, S-wave and density travel-time kernels (S-wave and density kernels are not shown here for brevity). The kernel stack is the spatially-dependent sum of individual source-receiver kernels for like-for-like, origin time relative temporal windows, computed for every 25^{th} seismogram in the earthquake gather (400m spacing). Given that we are principally interested in the relative information corresponding to constituent parts of the velocity model, as contributed by different time packets in the earthquake gather, we do not show the true polarity of the kernel sums, instead focusing on the absolute kernel stack values. It can be seen in figure 12d that the ballistic arrivals, which dominate the first 2 seconds of the earthquake gather, provide strong sensitivity across a large proportion of the model domain. However, the kernel stack has an approximate 'banana-doughnut' shape often seen for direct arrival-based kernels, with poor seismic illumination at the base of the model, particularly in the leftmost and rightmost bottom corners. Between 2-4 seconds, which contains secondary arrivals and a large portion of multi-frequency scattered seismic coda, it can be seen that the kernel stack has more uniform sensitivity across the vast majority of the velocity model domain. By 4-6 seconds, the magnitude of sensitivity has achieved a more uniform equilibrium across the model, with a slight (although not unexpected) sensitivity bias at depths above the source and below the receivers. From this figure we deduce that the scattered wavefield codas

contained in our broadband, SPECFEM-generated earthquake gathers, make an active contribution to completing the velocity model sensitivity. In addition to completing sensitivity across the entire model domain, the coda shows particular sensitivity to the dyke structure (figure 12). We further deduce that earthquake gather ballistics do not provide sufficient model sensitivity on their own. In light of this, we conclude that, by including the coda, we have amassed an earthquake gather databank that is sufficiently comprehensive and diverse for FNO training for the seismic inversion task, that includes finer scale structures.

5.2 Summary of Key Findings

We have shown that it is possible to train an FNO to directly invert for a velocity model from a surface-recorded earthquake array gather containing a multifrequency, high strength scattered seismic coda, excited by a broadband seismic source, in highly heterogeneous, volcanically-representative media. Using travel-time sensitivity kernels, we verified that the ballistic arrivals in each earthquake gather are insufficient to complete whole-model sensitivity, especially for fine scale structures. Furthermore, we showed that the seismic coda is required to illuminate the entire model (using 5-6 seconds of data after the earthquake origin time). In addition to achieving generalised recovery of 2D velocity models of P-wave velocity down to the order of $\approx$200m from a single earthquake gather, we were able to recover low velocity zones with first order morphology approximating what is expected for dyke propagation, down to similar length scales. With the addition of noise to an input earthquake gather, the trained FNO returned acceptable velocity models for signal-to-noise (SNR) ratios $\geq$50. This value range aligns with local volcanic seismic datasets recorded near the Krafla volcano, Iceland, where dozens of earthquakes exceeded the threshold SNR for our FNO over a two month instrument deployment (Maass et al. 2025). We also found that our FNO object was able to recover dyke-like velocity features at different stages of their intrusion into country rock (i.e. the background model).

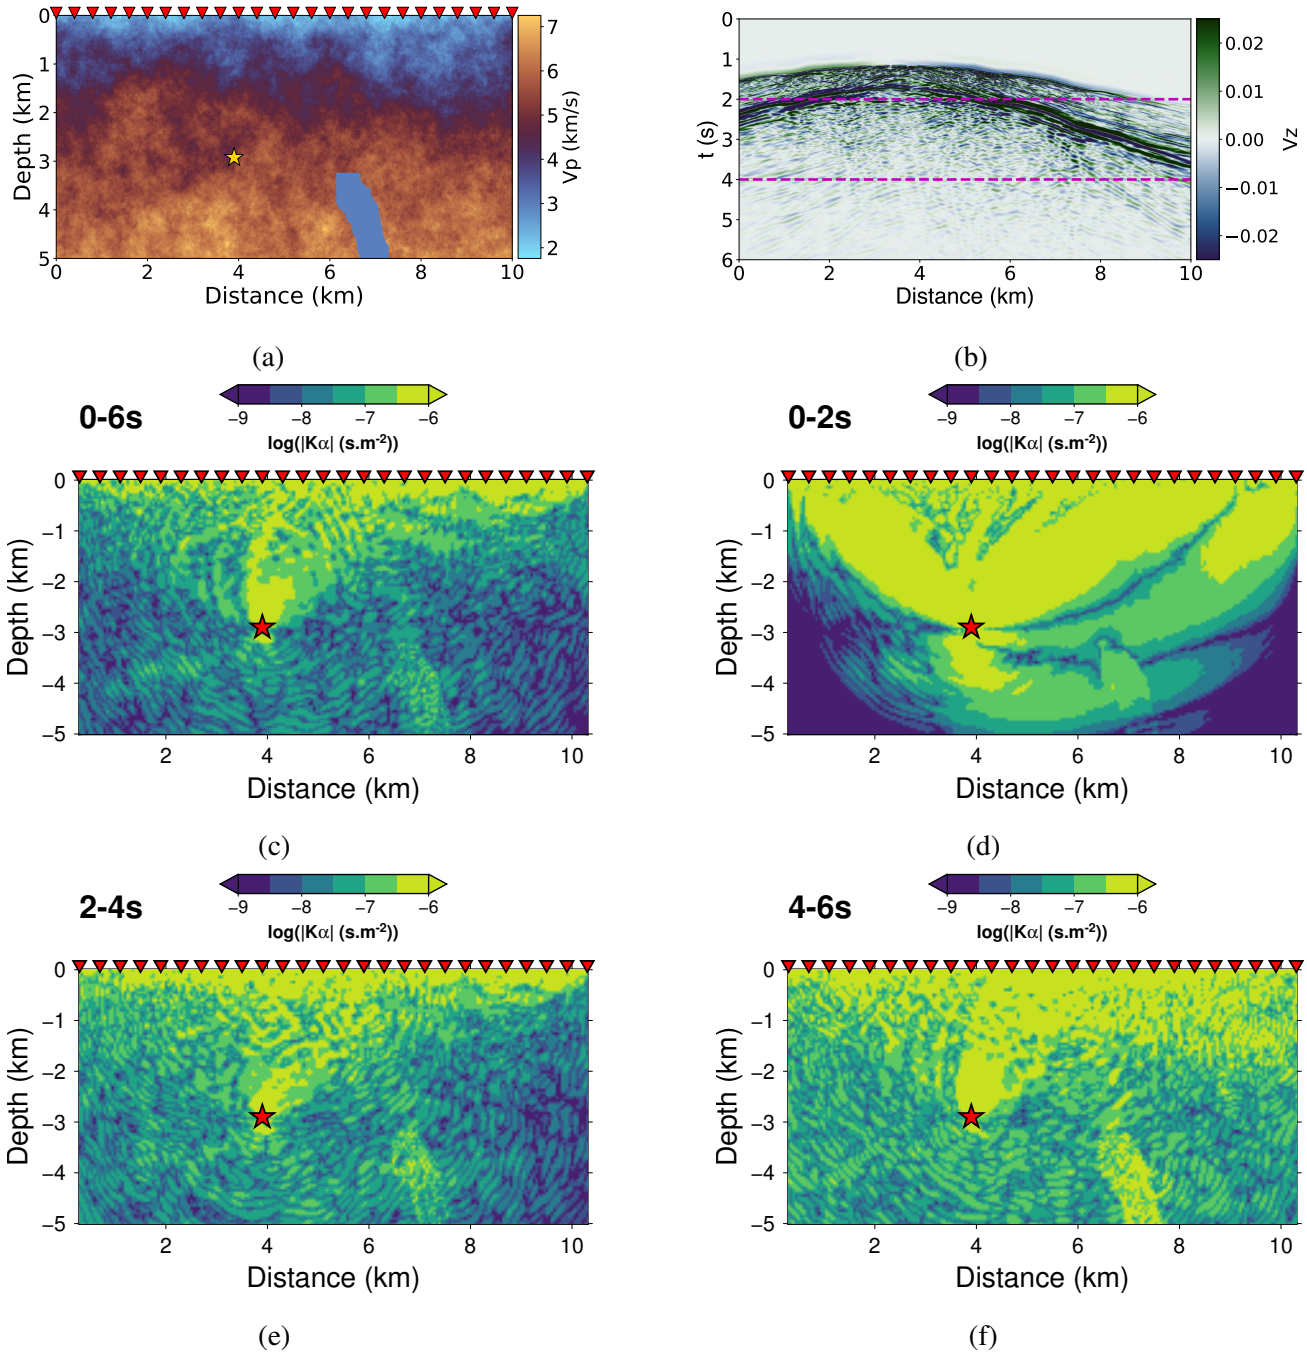

Figure 12. Demonstration of seismic sensitivity across the whole velocity model domain using P-wave travel-time sensitivity kernels. **a)** An example of a strongly scattering velocity model with a statistically distinct, embedded low velocity anomaly, of 0.5km width and 2km height. The source location is marked with a yellow star at x=3.9km, z=2.925km. The subset of 26 receivers, whose seismograms are used to compute the sensitivity kernel stacks shown in c)-f) are marked by red triangles. **b)** The corresponding earthquake gather (z component of velocity) for the model shown in a). The dashed magenta lines divide the receiver gather into 3 time windows (2 seconds duration), corresponding to the sensitivity kernel time window integrations shown in d)-f) **c)** The travel-time sensitivity kernel stack over the entire duration of the earthquake gather using individual kernels corresponding to the receivers shown in a). Shown is the logarithm of the absolute value of the kernel sum. **d)** Same as c) but for a sensitivity kernel integration time window between 0-2s **e)** 2-4s **f)** 4-6s.

5.3 Implication of Findings for Seismic Imaging

Our results show how machine learning can be used to support and extend state-of-the-art approaches to seismic imaging. Sensitivity analysis of the earthquake gathers in our training set demonstrates that the scattered part of the wavefield plays a significant role in 'illuminating' the entire velocity model domain, as well as catching sensitivity to low velocity anomalies or regions with high local impedance contrasts. In other words, the non-ballistic part of the earthquake gather is an essential requirement for velocity model image reconstruction. Given that the successful recovery tests for the background velocity model and embedded low velocity features are from just a single earthquake gather, the Fourier Neural Operator appears to be utilising the ballistics and seismic coda efficiently to invert for P-wave seismic velocity, with encouraging fidelity to the ground truth. This strengthens the rationale for future research into seismic imaging using the FNO as an 'inference engine'. Machine learning-driven inversion has the potential to extend the decades-long progression towards entire wavefield imaging, including coda waves directly in the imaging process, whose foundations are rooted in phase-picked travel-time tomography and full-waveform inversion. In addition to harnessing the unique fingerprint of a given earthquake gather for subsurface insight, FNO-based inversion has the potential for rapid model recovery (in $<$ 1sec), once the FNO is trained. We note that the caveat of this is the significant task of assembling a training and validation set population with acceptable generalisation potential to the target application. This will be a dominant challenge for any future simulated training-to-physical world imaging paradigm: that is, building training models that capture the broad characteristics of the physical world. We also acknowledge that in real volcanic environments, individual earthquake gather quality is typically affected by a combination of path and site effects, environmental noise and individual instrument responses. The current conclusions shown in this paper apply primarily to synthetic data. Future real-world applications will likely require FNO inversion using multiple seismic events or stacking strategies using multiple earthquakes. It is of note that, as is expected to be encountered in physical world volcanic environments, the inverse problem that we have sought to solve is ill-posed. Given that we have not included the source location as an input channel during FNO training, one potential limiting case for FNO performance is the existence of two velocity

models and unique sources which result in almost identical wavefields. This could result in an inability of the FNO to distinguish between two models. However, given that we are dealing with multiple scattering and multi-pathing of the wavefield, it is unlikely that two different models with different sources will produce the same results.

5.4 Next Steps and Future Research

We determined that the training dataset needed a combination of broad generic models and a population with specific features of interest to achieve the generalisation and stable learning of the FNO shown here, when trained using 16 Fourier modes (M=16), as well as a data channel uplift dimension of 32 (P=32). The final overall master dataset of 35,000 earthquake gathers generated in SPECFEM (30,000 generic models and 5,000 additional models with specific features) required approximately 4TiB of storage before decimation. Furthermore, each earthquake gather required between 2.5-3.0 CPUh compute time (i.e. one simulation took between 2-5 minutes to run on the computational resources available for this study). The Fourier transformed earthquake gathers input to the FNO were also truncated at the threshold number of Fourier modes of the trained model (M=16). We therefore expect that this modal truncation results in some under-utilisation of each earthquake gather's unique coda wave fingerprint during inversion. We argue that the results shown here are encouraging for future studies that can facilitate the inclusion of more Fourier modes during FNO training, with possible opportunities for higher fidelity velocity model recovery.

The FNO workflow shown here inverts for a 2D velocity model. Wavefield scattering in the physical world inevitably occurs in 3D and the coda wave amplitudes are also determined by 3D geometric spreading of the wavefront. However, we contend that there are a number of physical world problems where in-plane scattering serves as a sufficient approximation, such as linear nodal arrays that are perpendicular to geological strike. We note that alternative machine learning architectures such as multiple input FNOs, which have been shown to deliver computational speedups for the seismic forward problem (Perrone et al. 2025), could enable access to 3D inversion solutions, in future work.

Future research will build on the results presented in this study in an effort to leverage sim-to-sim training towards successful sim-to-real applications. This would require further extension of the training dataset of velocity models and earthquake gathers shown here by increasing the diversity of dyke features. Features to incorporate include a wider range of dip angle, dyke sinuosity and the inclusion of multiple dykes in a given velocity model realisation, using detailed geological field maps as a starting basis. Adapting the FNO architecture to invert for multiple channels including V_s could also help to capture the properties of fine scale features in the physical world. Future applications could invert for both P-wave and S-wave velocity, facilitating estimates of Poisson's ratio (i.e. through $\frac{V_p}{V_s}$) across the entire model. This could include diversifying dyke and other magmatic features in the velocity model set, with Poisson's ratio ranges constrained by laboratory studies of representative magma types. Most significantly, this would enable the characterisation of magma properties along a given dyke.

FNO-driven seismic inversion has the potential to harness the highest frequencies available in the seismic coda for improved subsurface insight, moving towards entire wavefield imaging. Given that dykes and other physical world features in volcanic settings can be as thin as 2m in cross-sectional width, possible lines of inquiry include training on high pass filtered earthquake gathers or adopting a multi-stage training approach (e.g. Kong et. al. (2025)). The latter approach entails training an FNO to invert for seismic velocity to a baseline resolution, before using the residual misfit between ground truth and prediction to improve model recovery at higher frequencies. Small aperture, quasi-continuous seismic data sampling, facilitated by recent advances in distributed acoustic sensing (DAS), has been used for seismic monitoring on volcanoes (Jousset et. al. 2022). Targeting FNO training for DAS dataset applications could offer further opportunities for improved access to the unique fingerprint of earthquake gathers. We note, however, that DAS records provide a uni-axial strain measurement that changes with the direction of the cable. This introduces further complications if one were to amass a training dataset to this end and may limit training to specific cable geometries. Perhaps most importantly of all is the opportunity to use FNO based inversion to perform seismic imaging in remote locations on earth where there is

limited station coverage. This approach could also be applied to planetary seismology where there is very limited station coverage (e.g. Stähler et. al. (2021)).

To close, the need for more inference-based approaches to seismic imaging has been rightly identified by Tsai et. al. (2023; 2023). They argue that Earth imaging has reached a state of maturity whereby sufficiently many plausible structures have been imaged across a range of geophysical scales, such that the field would benefit from inference-oriented frameworks such as geological feature inversion. While the imaging workflow shown here (e.g. dataset diversification using a range of dyke stencils and shapes) is in keeping with this philosophy, we believe that FNO-based imaging also retains the original (perhaps essential) spirit of early discovery-oriented waveform and tomography studies e.g Grand and Helmberger (1984). Provided that improved utilisation of unique coda wave fingerprints using machine learning approaches progresses sufficiently, the geophysics community is in a strong position to survey the subsurface on multiple spatial scales in an entirely new way.

6 CONCLUSION

In this paper we investigated the ability of a Fourier Neural Operator (FNO) machine learning network to invert for 2D seismic velocity models from strongly scattered earthquake gathers, designed to emulate multi-scale scattering expected in volcanic environments. We created 30,000 training velocity models with volcanic setting-representative 1D gradients, perturbed by anti-persistent von Kármán 2D series. We chose the correlation lengths of the perturbation fields and their standard deviation based on sonic log observations from the physical world. For each velocity model, we used SPECFEM2D to forward model a corresponding earthquake gather with up to ≈25Hz coda wave spectral content, induced by multifrequency scattering in each model and a broadband earthquake source time function. Having trained an FNO for the velocity inversion task, our main findings are summarised below:

(i) The FNO trained on the core dataset (size 30,000) recovered the broadscale 2D velocity structure for the test set and also when applied to unseen earthquake gather data, from single earthquake gathers, resolving structure down to as low as 250m.

(ii) By embellishing the core training set with ≈ 5,000 additional velocity models with embedded geological features resembling intrusive dykes in country rock, FNO re-training from scratch resulted in sharper dyke recovery. We stress tested the recovery resolution down to thicknesses on the order of 200m. Intrusive features included in the re-training dataset exhibited wide Vp and Vp/Vs ranges informed by seismicity and tomography observations from Iceland.

(iii) We tested the resilience of our fine-tuned FNO model to environmental noise. This entailed adding 2D Gaussian noise into earthquake gathers during the FNO extraction, transformation and loading stage (ETL). Acceptable dyke feature recovery persisted down to SNRs as low as 25 for individual earthquake gathers, with acceptable background model recovery down to SNRs ≈50.

Future work will focus on improving FNO resolution recovery, investigating the feasibility of 3D FNO inversion and exploring the possible application of the workflow shown here to real seismic data.

7 DATA AVAILABILITY STATEMENT

The velocity models and earthquake gathers used for this research are available on request. Please contact EJT on etotten@cp.dias.ie to arrange this. SPECFEM2D software (Tromp et. al. 2008) was used to generate earthquake gathers using waveform modelling. The codes for velocity model data generation and launching multiple SPECFEM2D simulations, along with machine learning models related to this research, are available at: `https://git.dias.ie/eoghan_totten/dias_fno`

Borehole data is provided by GSNI under the Crown Copyright and Database Right 2024 – SpatialNI - A service provided by Ordnance Survey of Northern Ireland. More information is available at `https://mapapps2.bgs.ac.uk/gsni_geoindex/home.html`

ACKNOWLEDGMENTS

This publication has emanated from research conducted with the financial support of Taighde Éireann – Research Ireland under Grant number GOIPD/2023/1241. We wish to express sincere thanks to Dr. Shane Maloney from the Astrophysics section of the Dublin Institute for Advanced

Studies for permitting us to perform many machine learning training runs on an NVIDIA RTX A6000 GPU. Sincere thanks to Dr. Regina Maass for helping EJT access Krafla earthquake catalogue data, as well as computing earthquake signal-to-noise ratios for comparisons with the results shown here. Thanks are also due to Klaus Holliger for thoughtful email discussions on using sonic borehole log data to inform the creation of synthetic velocity models in 2D. EJT and CJB thank Philippe Grange and Jean-Francois Bucas of DIAS IT for their excellent IT support. EJT expresses thanks to the DIAS Geophysics research group for their useful feedback at regular group meetings.

Fourier Neural Operator training was computed mainly on an NVIDIA RTX A6000 GPU (48GiB). Trained machine learning models were applied on an NVIDIA Tesla P100-PCIE-16GB GPU. We mainly used perceptually uniform colour maps (Crameri et. al. 2020) to create figures shown in this paper.

The authors declare that there are no competing interests.

REFERENCES

Bean, C.J., 1996. On the cause of 1/f-power spectral scaling in borehole sonic logs. *Geophysical Research Letters*, 23(22), pp.3119-3122.

Bean, C.J., Marsan, D. and Martini, F., 1999. Statistical measures of crustal heterogeneity from reflection seismic data: the role of seismic bandwidth. *Geophysical Research Letters*, **26(21)**, pp.3241-3244.

Bean, C.J. and Martini, F., 2010. Sub-basalt seismic imaging using optical-to-acoustic model building and wave equation datuming processing. *Marine and Petroleum Geology*, **27(2)**, pp.555-562.

Bracale, M., Campillo, M., Shapiro, N.M., Brossier, R. and Melnik, O., 2025. Multiple scattering of seismic waves in a heterogeneous magmatic system and spectral characteristic of long period volcanic earthquakes. *EarthArXiv*, https://eartharxiv.org/repository/view/8889/. Preprint.

Brougois, A., Bourget, M., Lailly, P., Poulet, M., Ricarte, P. and Versteeg, R., 1990, May. Marmousi, model and data. In EAEG workshop-practical aspects of seismic data inversion (pp. cp-108). European Association of Geoscientists & Engineers.

Calvet, M. and Margerin, L., 2013. Lapse-time dependence of coda Q: Anisotropic multiple-scattering models and application to the Pyrenees. *Bulletin of the Seismological Society of America*, **103(3)**, pp.1993-2010.

Crameri, F., Shephard, G.E. and Heron, P.J., 2020. The misuse of colour in science communication. *Nature communications*, **11(1)**, p.5444.

De Gori, P., Chiarabba, C. and Patanè, D., 2005. Qp structure of Mount Etna: Constraints for the physics of the plumbing system. *Journal of Geophysical Research: Solid Earth*, **110(B5)**.

Del Pezzo, E., Bianco, F. and Zaccarelli, L., 2006. Separation of Qi and Qs from passive data at Mt. Vesuvius: A reappraisal of the seismic attenuation estimates. *Physics of the Earth and Planetary Interiors*, **159(3-4), pp.202-212**.

De Siena, L., Calvet, M., Watson, K.J., Jonkers, A.R.T. and Thomas, C., 2016. Seismic scattering and absorption mapping of debris flows, feeding paths, and tectonic units at Mount St. Helens volcano. *Earth and Planetary Science Letters*, **442**, pp.21-31.

Dolan, S.S., Bean, C.J. and Riollet, B., 1998. The broad-band fractal nature of heterogeneity in the upper crust from petrophysical logs. *Geophysical Journal International*, **132(3)**, pp.489-507.

Eibl, E.P., Bean, C.J., Vogfjörd, K.S., Ying, Y., Lokmer, I., Möllhoff, M., O'Brien, G.S. and Pálsson, F., 2017. Tremor-rich shallow dyke formation followed by silent magma flow at Bárdarbunga in Iceland. *Nature Geoscience*, **10(4)**, pp.299-304.

Flood, R.D., Piper, D.J.W., Klaus, A., et al. [1995] Proceedings of the Ocean Drilling Program, Initial Reports. College Station, TX, Ocean Drilling Programme.

Glastonbury-Southern, E., Winder, T., Rawlinson, N., White, R.S., Greenfield, T., Bacon, C.A., Ágústsdóttir, T., Brandsdóttir, B., Gudnason, E.Á., Hersir, G.P. and Fischer, T.J., 2025. Pre-existing structures control the orientation of strike-slip faulting during the 2021 dike intrusion at Fagradalsfjall, Iceland. *Journal of Geophysical Research: Solid Earth*, **130(6)**, p.e2024JB030162.

Glück, E., Garambois, S., Vandemeulebrouck, J., Muzellec, T., Virieux, J., Mortensen, A.K., Gudnason, E.Á. and Ágústsdóttir, T., 2025. Seismicity patterns and multi-scale imaging of Krafla, N–E Iceland, with local earthquake tomography. *Journal of Geophysical Research: Solid Earth*, 130(5), p.e2024JB030425.

Grand, S.P. and Helmberger, D.V., 1984. Upper mantle shear structure of North America. *Geophysical Journal International*, **76(2)**, pp.399-438.

Gudmundsson, A., 2025. Statistical Physics of Fissure Swarms and Dike Swarms. *Geosciences*, **15(8)**, p.301.

Holliger, K., 1996. Upper-crustal seismic velocity heterogeneity as derived from a variety of P-wave sonic logs. *Geophysical Journal International*, **125(3)**, pp.813-829.

Jousset, P., Currenti, G., Schwarz, B., Chalari, A., Tilmann, F., Reinsch, T., Zuccarello, L., Privitera, E. and Krawczyk, C.M., 2022. Fibre optic distributed acoustic sensing of volcanic events. *Nature communications*, **13(1)**, p.1753.

Komatitsch, D. and Tromp, J., 2002. Spectral-element simulations of global seismic wave propagation—I.

Validation. *Geophysical Journal International*, **149(2)**, pp.390-412.

Kong, Q. and Rodgers, A., 2023. Feasibility of using Fourier neural operators for 3D elastic seismic simulations (No. LLNL-TR-854521). Lawrence Livermore National Laboratory (LLNL), Livermore, CA (United States).

Kong, Q., Zou, C., Choi, Y., Matzel, E.M., Azizzadenesheli, K., Ross, Z.E., Rodgers, A.J. and Clayton, R.W., 2025. Reducing Frequency Bias of Fourier Neural Operators in 3D Seismic Wavefield Simulations Through Multi-Stage Training. arXiv preprint arXiv:2503.02023.

Kovachki, N., Lanthaler, S., Mishra, S., 2021. On Universal Approximation and Error Bounds for Fourier Neural Operators. *Journal of Machine Learning Research*, **22**, pp.1-76.

Lehmann, F., Gatti, F., Bertin, M. and Clouteau, D., 2023. Fourier neural operator surrogate model to predict 3D seismic waves propagation. *arXiv preprint arXiv:2304.10242*.

Lehmann, F., Gatti, F., Bertin, M. and Clouteau, D., 2024, April. Quantifying uncertainties in seismic waves propagation with a Fourier Neural Operator surrogate model. In *MASCOT-NUM 2024*.

Lehmann, F., Gatti, F. and Clouteau, D., 2025. Multiple-input Fourier Neural Operator (MIFNO) for source-dependent 3D elastodynamics. *Journal of Computational Physics*, **527, p.113813**.

Li, Z., Kovachki, N., Azizzadenesheli, K., Liu, B., Bhattacharya, K., Stuart, A. and Anandkumar, A., 2020. Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*.

Li, J., Biondi, E., Heimisson, E.R., Puel, S., Zhai, Q., Zhang, S., Hjörleifsdóttir, V., Wei, X., Bird, E., Klesh, A. and Kamalov, V., 2025. Minute-scale dynamics of recurrent dike intrusions in Iceland with fiber-optic geodesy. *Science*, **388(6752)**, pp.1189-1193.

Lokmer, I. and Bean, C.J., 2010. Properties of the near-field term and its effect on polarisation analysis and source locations of long-period (LP) and very-long-period (VLP) seismic events at volcanoes. *Journal of Volcanology and Geothermal Research*, **192(1-2)**, pp.35-47.

Maass, R., Li, K.L. and Bean, C.J., 2025. Improving passive reflection seismic imaging in complex geological settings through site effect reduction: application to Krafla volcano, Iceland. *Geophysical Journal International*, **241(2)**, pp.756-769.

Martin, G.S., Wiley, R. and Marfurt, K.J., 2006. Marmousi2: An elastic upgrade for Marmousi. *The leading edge*, **25(2)**, pp.156-166.

Martínez Montesinos, B., Bean, C.J. & Lokmer, I., 2021. Quantifying strong seismic propagation effects in the upper volcanic edifice using sensitivity kernels, *Earth Planet. Sci. Lett.*, **554, 116683**.

Martini, F., Hobbs, R.W., Bean, C.J. and Single, R., 2005. A complex 3D volume for sub-basalt imaging. *First Break*, **23(7)**.

Masihi, A. and Fischer, T., 2025. In-Situ Vp/Vs ratio variations in seismic swarms as indicator of magmatic processes: Fagradalsfjall volcanic activity, SW Iceland. *Geophysical Journal International*, **241(3)**, pp.1725-1739.

Menke, W., Levin, V. and Sethi, R., 1995. Seismic attenuation in the crust at the mid-Atlantic plate boundary in south-west Iceland. *Geophysical Journal International*, **122(1)**, pp.175-182.

Naylor, D., Philcox, M.E. and Clayton, G., 2003. Annaghmore-1 and Ballynamullan-1 wells, Larne-Lough Neagh basin, Northern Ireland. *Irish Journal of Earth Sciences*, **21(1)**s, pp.47-69.

Nelson, C.E., Hobbs, R.W. and Rusch, R., 2015. On the use of fractal surfaces to understand seismic wave propagation in layered basalt sequences. *Pure and Applied Geophysics*, **172(7)**, pp.1879-1892.

Nguyen, Q.H., Ly, H.B., Ho, L.S., Al-Ansari, N., Le, H.V., Tran, V.Q., Prakash, I. and Pham, B.T., 2021. Influence of data splitting on performance of machine learning models in prediction of shear strength of soil. *Mathematical Problems in Engineering*, **2021(1)**, p.4832864.

Nooshiri, N., Bean, C.J., Dahm, T., Grigoli, F., Kristjánsdóttir, S., Obermann, A. and Wiemer, S., 2022. A multibranch, multitarget neural network for rapid point-source inversion in a microseismic environment: examples from the Hengill Geothermal Field, Iceland. *Geophysical Journal International*, **229(2)**, pp.999-1016.

Obermann, A., Planès, T., Larose, E. and Campillo, M., 2013. Imaging preeruptive and coeruptive structural and mechanical changes of a volcano with ambient seismic noise. *Journal of Geophysical Research: Solid Earth*, **118(12)**, pp.6285-6294.

Obermann, A., Larose, E., Margerin, L. and Rossetto, V., 2014. Measuring the scattering mean free path of Rayleigh waves on a volcano from spatial phase decoherence. *Geophysical Journal International*, **197(1)**, pp.435-442.

O'Brien, G. S. & Bean, C. J., 2009, Volcano topography, structure and intrinsic attenuation: Their relative influences on a simulated 3D visco-elastic wavefield, *Journal of Volcanology and Geothermal Research*, **183**, 122–136.

O'Brien, G.S., Bean, C.J., Meiland, H. and Witte, P., 2023. Imaging and seismic modelling inside volcanoes using machine learning. *Scientific Reports*, **13(1), p.630**.

Paulatto, M., Hooft, E.E., Chrapkiewicz, K., Heath, B., Toomey, D.R. and Morgan, J.V., 2022. Advances in seismic imaging of magma and crystal mush. *Frontiers in Earth Science*, **10**, p.970131.

Perrone, N., Lehmann, F., Gabrielidis, H., Fresca, S. and Gatti, F., 2025. Integrating fourier neural operators with diffusion models to improve spectral representation of synthetic earthquake ground motion response. *arXiv preprint arXiv:2504.00757*.

PyTorch, 2025. torch.nn.Dropout2d — PyTorch 2.8 documentation. Available at: *https://docs.pytorch.org/docs/2.8/generated/torch.nn.Dropout2d.html* [Accessed 1 Jun. 2025].

Rahimi Dalkhani, A., Ágústsdóttir, T., Gudnason, E.A., Hersir, G.P., Zhang, X. and Weemstra, C., 2024. Transdimensional ambient-noise surface wave tomography of the Reykjanes Peninsula, SW Iceland. *Geophysical Journal International*, **236(1), pp.621-643**.

Santurkar, S., Tsipras, D., Ilyas, A. and Madry, A., 2018. How does batch normalization help optimization?

Advances in neural information processing systems, **31**.

Sato, H., 2019. Power spectra of random heterogeneities in the solid earth. *Solid Earth*, **10(1)**, pp.275-292.

Stähler, S.C., Khan, A., Banerdt, W.B., Lognonné, P., Giardini, D., Ceylan, S., Drilleau, M., Duran, A.C., Garcia, R.F., Huang, Q. and Kim, D., 2021. Seismic detection of the martian core. *Science*, **373(6553)**, pp.443-448.

Tompkins, M.J. and Christensen, N.I., 2001. Ultrasonic P-and S-wave attenuation in oceanic basalt. *Geophysical Journal International*, **145(1)**, pp.172-186.

Tromp, J., Tape, C. and Liu, Q., 2005. Seismic tomography, adjoint methods, time reversal and banana-doughnut kernels. *Geophysical Journal International*, **160(1)**, pp.195-216.

Tromp, J., Komattisch, D., Liu, Q., 2008. Spectral-element and adjoint methods in seismology. *Commun. Comput. Phys*. **3**, 1–32.

Tsai, V.C., 2023. The future of Earth imaging. *Seismological Research Letters*, **94(5)**, pp.2119-2128.

Tsai, V.C., Huber, C. and Dalton, C.A., 2023. Towards the geological parametrization of seismic tomography. *Geophysical Journal International*, **234(2)**, pp.1447-1462.

Turcotte, D.L., 1997. Fractals and chaos in geology and geophysics. *Cambridge university press*.

Wegler, U. and Lühr, B.G., 2001. Scattering behaviour at Merapi volcano (Java) revealed from an active seismic experiment. *Geophysical Journal International*, **145(3)**, pp.579-592.

White, J.C., 2009. Development and application of the phase-screen seismic modelling code (Doctoral dissertation, Durham University).

Yang, Y., Gao, A.F., Castellanos, J.C., Ross, Z.E., Azizzadenesheli, K. and Clayton, R.W., 2021. Seismic wave propagation and inversion with neural operators. *The Seismic Record*, **1(3)**, pp.126-134.

Yang, Y., Gao, A.F., Azizzadenesheli, K., Clayton, R.W. and Ross, Z.E., 2023. Rapid seismic waveform modeling and inversion with neural operators. IEEE Transactions on Geoscience and Remote Sensing, 61, pp.1-12.

Zou, C., Ross, Z.E., Clayton, R.W., Lin, F.C. and Azizzadenesheli, K., 2025. Ambient Noise Full Waveform Inversion with Neural Operators. *arXiv preprint arXiv:2503.15013*.

APPENDIX A: VELOCITY MODEL CREATION

A1 Velocity Model Population: 1D Deterministic Structure

For the 1D deterministic velocity structure, we chose a modified gradient from the starting model of Eibl et al. (2017) related to the volcanic zone around the Vatnajökull glacier, southeast Iceland (figure 1). We removed the near-surface velocity signature of the cryosphere layer as this study

focuses on seismic imaging in the upper crust only. We also increased the starting surface P-wave velocity to 2.6km/s and lowered the base velocity at 5km depth to 6.5km/s.

A2 Velocity Model Population: 2D Statistical Perturbations

A2.1 Von Kármán Field Power Spectra

In addition to the definitions laid out in Dolan et. al. (1998) we considered implementations by White (2009) and Sato (2019) when creating the von Kármán models used in this study. The power spectrum of a 2D Von Kármán series used to perturb each velocity model can be expressed as (Dolan et. al. 1998):

$$P(k) = 2\sqrt{\pi a}\frac{\Gamma(H+0.5)}{(1+k^2a^2)^{H+0.5}\Gamma(H)} \tag{A.1}$$

where H is the Hurst exponent and a_x and a_z are the correlation lengths in lateral (x) and vertical (z) directions respectively.

For previous studies in volcanic settings, Martini et al. (2005) used a Hurst exponent of 0.3 for waveform modelling studies in intrusive basalt layers as constrained by sonic borehole log measurements in basalts as part of the Ocean Drilling Project (Flood et al. 1995). Furthermore, Bean & Martini (2010) imaged a Tertiary basalt outcrop in County Antrim, Northern Ireland, in order to perform optical to acoustic waveform modelling for sub-basalt imaging. They deduced a Hurst exponent of 0.2 for this outcrop, although Nelson et al. (2015) note that the outcrop in question is at the base of a basalt succession overlying limestone. This may have further altered the medium roughness due to limestone erosion and subsequent lava flow infill.

The slope of the power spectrum of a Von Kármán series falls off as

$$-2H - E \tag{A.2}$$

at wavelengths below the correlation length, where E is the Euclidean dimension i.e E=2 in 2D (Turcotte 1997). For wavelengths greater than the correlation length, the power spectrum of a Von Kármán series approximately scales as white noise (Bean et. al. 1999). Figure A1 shows example

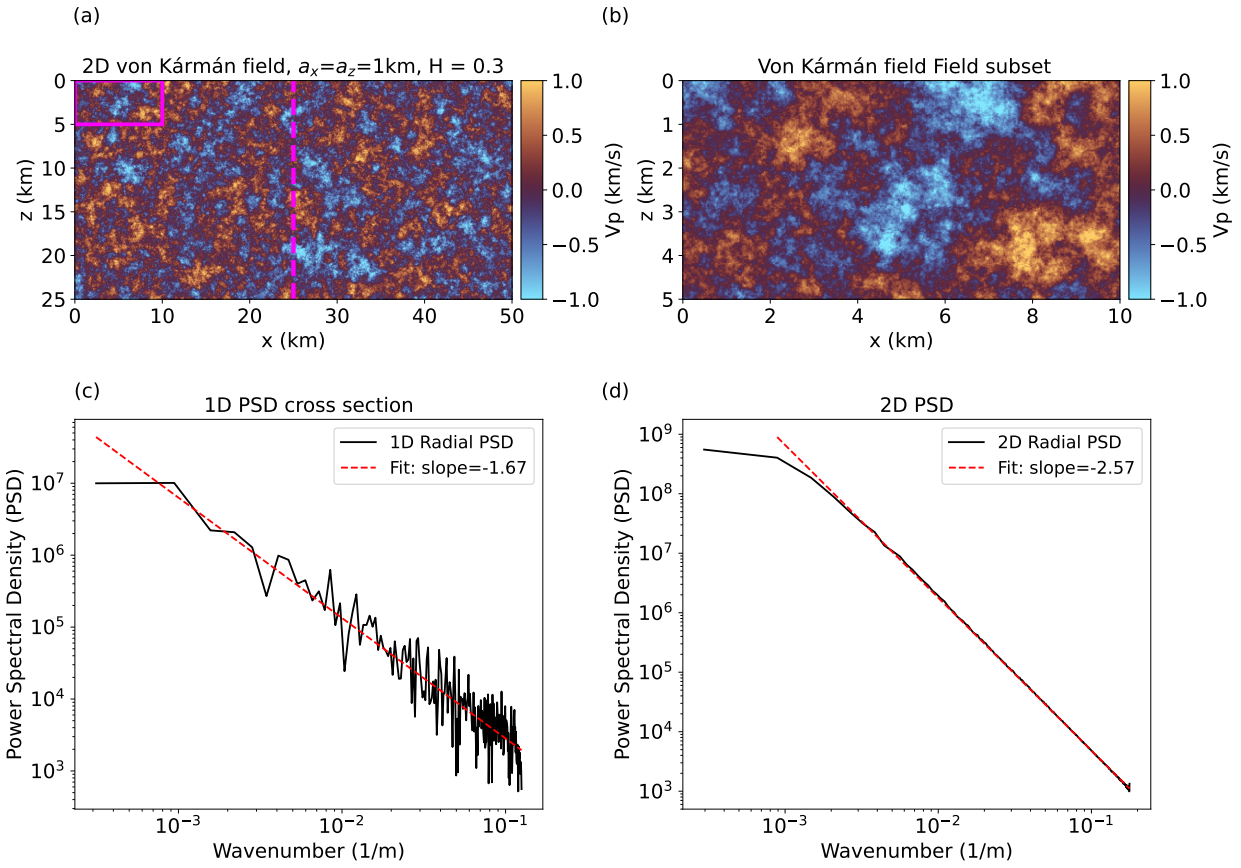

Figure A1. Illustration of a von Kármán medium and the behaviour of its power spectral density. **(a)** 25 km deep × 50 km wide master P-wave velocity anomalies described by a von Kármán medium, with a 1km isotropic correlation length and scattering persistence described by a Hurst exponent of 0.3. The magenta rectangle is the 5 km deep × 10 km wide model subset shown in (b). The dashed magenta line shows the cross-section used to calculate its 1D power spectrum shown in (c). **(b)** 5 km × 10 km wide velocity model subset of the master model shown in (a). **(c)** 1D power spectrum as a function of wavenumber for the cross-section through the master model shown in (a). The red dashed line is the line of best fit through the power spectrum at wavelengths less than the correlation length (i.e., wavenumbers greater than 10^{-3} m^{-1}), with slope ≈ -2H-E (E=1) **(d)** 2D power spectrum of the master model shown in (a) with slope ≈ -2H-E (E=2).

power spectra in 1D and 2D for the Von Kármán series used in this study.

A2.2 Field-derived Variances for Von Kármán Fields

We analysed data from three sonic borehole logs hereafter referred to as the Annaghmore, Ballynamullan and Ballymackilroy boreholes. The Annaghmore and Ballynamullan borehole locations

are situated within one kilometre of each other, with the Ballymackilroy borehole situated approximately 15 kilometres to the northeast of Annaghmore and Ballynamullan. Figure A2 demonstrates how we extracted a standard deviation value from each of the three borehole logs using Ballymackilroy as an example. The raw data for each borehole was sampled in approximately 15cm increments. We detrended each borehole log before applying a running mean of 24.9m. This step was carried out to align field-derived measurements of seismic velocity in volcanic rocks with the representation resolution of the numerical grids on which we project each velocity model for waveform forward modelling. At this stage of the workflow we used a grid with spacing of 25m for the lateral and vertical extent of velocity models. This 25m grid spacing was chosen to ensure numerical stability during full waveform modelling of the seismic wavefield through each velocity model (see section 2.3 for more information). After applying an approximately 25m running mean to each detrended borehole log, we calculated standard deviations of $\pm$240m/s, $\pm$530m/s and $\pm$375m/s for the Annaghmore, Ballynamullan and Ballymackilroy boreholes respectively.

In the interests of computational tractability when computing a wavefield through each velocity model, we clipped each Von Kármán instance at $\pm 3\sigma$. In the uppermost 1 kilometre of each model, we further constrained the distribution to $\pm 2\sigma$. This was done to strike a balance between velocity model grid spacing and the number of grid points available to represent the shortest S wavelengths, in low velocity zones at upper model depths, when modelling the seismic wavefield, ensuring numerical stability.

APPENDIX B: FULL WAVEFORM FORWARD MODELLING (SPECFEM2D)

B1 Justifying Velocity Model Attenuation Factors

O'Brien et. al. (2009) noted low frequency ($<$10Hz) broad spectral content for synthetic seismograms generated in strongly attenuating, fractal based velocity models ($Q_p = 40$ and $Q_s = 20$, H=0.3). For Iceland, Menke et. al. (1995) expect $Q_p \approx 60$ in the upper 4km of crust due to near-surface heterogeneity attributed to fissures, faulting and porosity changes, with lower values expected in volcanic regions of the island. Regarding laboratory measurements of attenuation from basalt drill cores, Tompkins & Christensen (2001) report a Q_p range of 14-167 for water-saturated

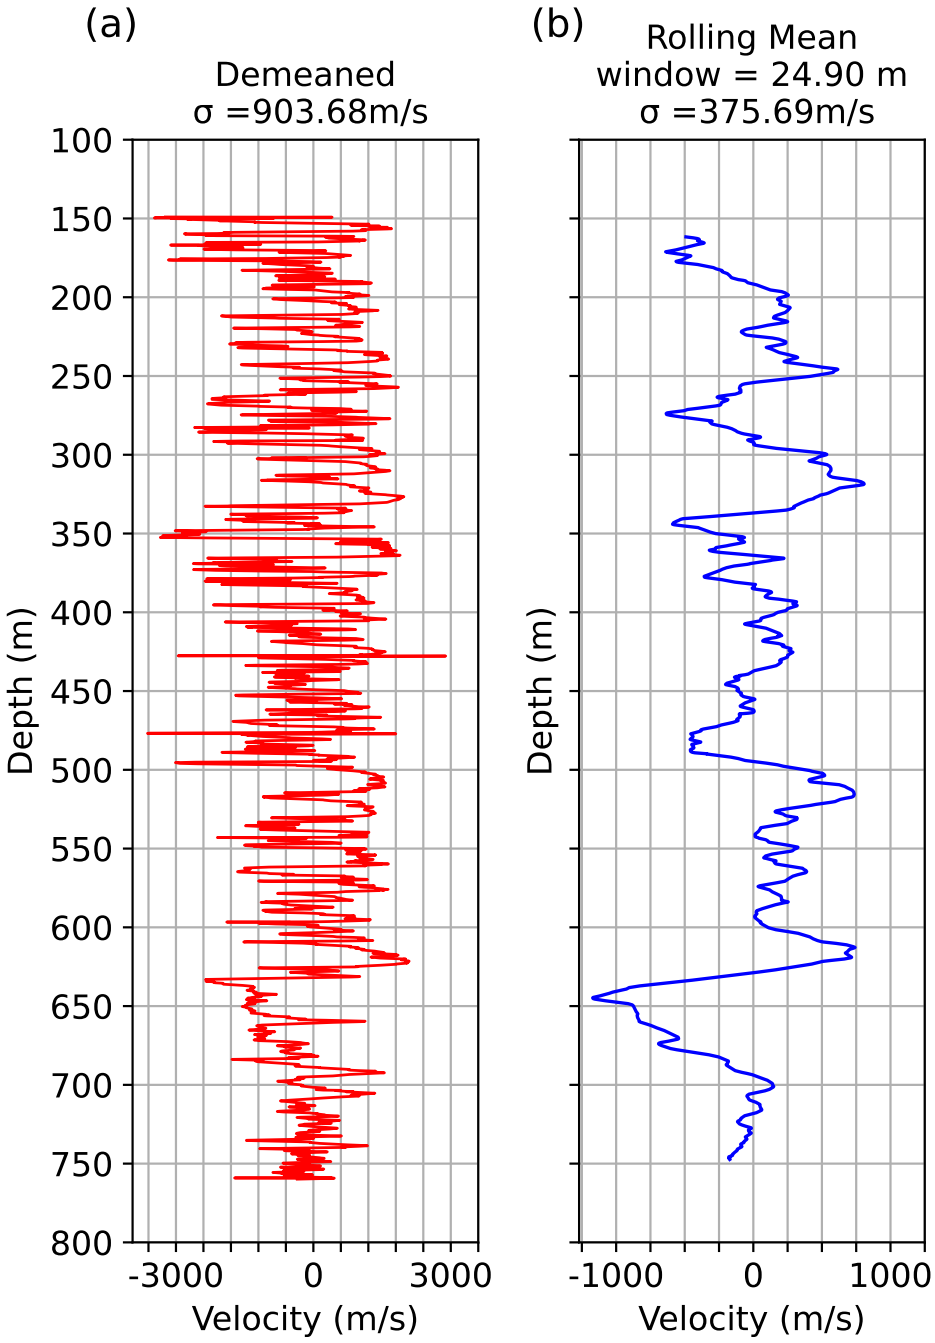

Figure A2. P-wave sonic borehole log through basaltic lavas for the Ballymackilroy borehole in county Antrim, Northern Ireland. The data is plotted down to 760m depth, where the basaltic lavas contact the Ulster white limestone formation. **(a)** Demeaned borehole log using the original sampling interval of 0.15m **(b)** Demeaned borehole log with a running mean applied to the data. The running mean window length is 24.9m, which broadly matches the 25m resolution velocity model grids upon which we forward model SPECFEM wavefields for FNO training. The demeaned log has a one standard deviation and range of approximately 375m/s and $\pm$1100m/s respectively.

oceanic basalts, and a corresponding Q_s range between 8-27. Attenuation factors have also been measured in other volcanic settings. Del Pozzo et al. (2006) observed that the total quality factor on Mount Vesuvius, Q_i, is highly frequency dependent, ranging from as low as 2 to as much as 1000 at higher frequencies. For Mt Etna, De Gori et al. (2005) observed a Q_p range between 30-80 down to 3km beneath the surface, with a wider range out to as high as 120 at 6km depth.

B2 Ensuring Numerical Stability

We needed to take numerical stability into account when forward modelling a SPECFEM wavefield for each velocity model realisation. The Courant-Friedrichs-Lewy (CFL) criterion is used to ensure partial differential equation convergence and can be defined as

$$CFL = \frac{v\Delta t}{\Delta x} \leq CFL_{max} \tag{B.1}$$

where v is is the seismic velocity at a given grid point, Δt is the simulation time step, Δx is the grid step length and CFL_{max} is the threshold CFL value. Staying within the CFL criterion mitigates for the introduction of numerical artefacts into the wavefield. In order to ensure that a CFL criterion value $\leq$0.5 was achieved for every simulation throughout the model domain, we used a time step of 0.285 milliseconds i.e. an approximate sampling frequency of 3.5kHz (SPECFEM2D automatically calculates the CFL criterion and suggests the corresponding time step based on the fastest voxels in the model). Following this protocol, we used a time step of 0.285 milliseconds. Once a given simulation is complete, we save earthquake gathers of size 21052 (time steps) x 401 (number of receiver objects), before further decimating temporally to approximately 150Hz sampling (i.e. each 6 second seismogram is represented by 900 time steps).

APPENDIX C: FNO MODEL WORKFLOW AND HYPERPARAMETERS

C1 FNO Workflow

Training and test datasets were fed in mini batches to the workflow during training and validation. Each 4D input tensor had dimensions $batch\ size \times 2N_{rec} \times N_t$, where $batch\ size$ denotes the chosen mini-batch size, N_{rec} the number of receivers per earthquake gather, N_t the number of time steps

per earthquake gather. Here we converged on an optimal mini-batch size of 20, used 201 receivers per earthquake gather (for each component of ground motion, V_x and V_z) and decimated to 450 time steps per earthquake gather. Instead of feeding in V_x and V_z as two separate input channels, we concatenated each V_x and V_z earthquake gather to form one image. Feeding the earthquake gather data in this configuration had little to no effect on the overall performance of the FNO during training.

Next, we padded each earthquake gather in the mini-batch from 402×450 to 512×512 to mitigate for any artefacts which may arise from Fourier transform operations. This aligns with previous studies where FNO input tensors did not have periodic boundaries (Tang et al., 2024). From here the earthquake gathers are uplifted to a higher dimension using a fully connected layer in PyTorch, which applies a linear transformation (P) to the earthquake gathers (u(x)) of the form

$$v_0(x) = P(u(x)) \tag{C.1}$$

where P applies the operation

$$v_0(x) = u(x)A^T + b \tag{C.2}$$

with A being the weight matrix and b the bias term. Here we uplifted to a width of 32. Increasing the width beyond this value caused the FNO to overfit the training set.

Once the earthquake gathers have been uplifted as described above, they are passed through several Fourier neural operator layers. Each of these can be expressed as

$$v_{l+1}(x) = \sigma((Kv_l)(x) + Wv_l(x)) \tag{C.3}$$

where W is a local linear transform on the data, actioned as a pointwise convolution by PyTorch's conv1d layer, K is a Fourier space-based operator and σ is the non-linear GeLU activation function applied after batch normalisation of $(Kv_l)(x) + Wv_l(x)$. It is a common practice to use batch normalisation before activation layers in neural networks to promote more stable and faster training by smoothing the training optimisation landscape (Santurkar et al. 2018). The Fourier space-based operator, K, can be thought of as a spectral convolution operation and can be expressed as (Li et al. 2020)

$$(Kv_l)(x) = \mathcal{F}^{-1}(R.T_M(\mathcal{F}v_l))(x) \tag{C.4}$$

Layer	**Operation**	**Tensor Shape**
-	Input earthquake gathers	20, 402, 450, 1
Padding	Pad earthquake gathers	20, 512, 512, 1
Channel Uplift (P)	Cat, Linear Operator	20, 512, 512, 32
-	Permute	20, 32, 512, 512
FOL1	SpectralConv2D, Linear, Add, GELU	20, 32, 512, 512
FOL2	SpectralConv2D, Linear, Add, GELU	20, 32, 512, 512
FOL3	SpectralConv2D, Linear, Add, GELU	20, 32, 512, 512
FOL4	SpectralConv2D, Linear, Add	20, 32, 512, 512
Project to target (Q)	Linear operator	20, 401, 201,1

Table A1. FNO model architecture used in this study. SpectralConv2D represents one pass through a Fourier Operator Layer (FOL). GELU is the Gaussian Error Linear activation unit.

where $\mathcal{F}$ and $\mathcal{F}^{-1}$ denote Fourier transform and inverse Fourier transform operations respectively and T_M is a filter which retains the first M modes from the input. R is a learnable operation applied in Fourier space. R multiplies the input tensors with an input-to-output, spatially-varying weight matrix, for each feature at every spatial grid-point and is actioned using PyTorch's einsum operation. We chose M=16 for this implementation.

The input data tensors are passed through L many layers (in this study, L=4). Each Fourier operator layer (FOL) consists of a 2D spectral convolution and a local linear transform (pointwise convolution), followed by batch normalisation and non-linear activation. Following this, the outputs are transformed to the target output dimensions (in our case, batches of 2D models of P-wave velocity of lateral size 401 (x) $\times$ vertical size 201 (z)) by a series of successive linear operations, again using PyTorch conv1d.

Full-scale FNO training on the 30,000 sized earthquake gather and velocity model dataset was computed on an NVIDIA RTX A6000 GPU with 48GiB memory. The application of the trained FNO machine learning model objects to unseen data, as shown in what follows in this paper, were mainly carried out using a smaller NVIDIA Tesla P100-PCIE-16GB GPU.

Parameter	**Value/Choice**
Batch Size	20
Optimiser	AdamW + lookahead
Weight decay	0.0001
Initial learning rate	0.001
Learning rate decay factor	2
Learning rate update frequency	20 epochs
Activation Function	GeLu
2D FOL dropout	p=0.1
Loss function	L_2 (PyTorch's MSELoss)
No. Fourier Modes	16
FNO Width	32

Table A2. FNO model hyperparameters used in this study. Abbreviation 'FOL' refers to Fourier Operator Layers. 2D FOL dropout is applied after layers 1,2 and 3 only. GELU is the Gaussian Error Linear activation unit.

C2 FNO Hyperparameters

The tuning parameters for the trained FNO machine learning model are summarised in table A2. We chose a batch size of 20. In practice this means that loss is computed between the ground truth 2D P-wave velocity model and the FNO-predicted velocity model for all 20 predicted models in a given batch. The derivative of the average of all losses in the batch with respect to the FNO model's weights is then computed, followed by updating the FNO model weights using backpropagation. We used Pytorch optimiser's AdamW with lookahead as our optimiser of choice. We choose a weight decay of $1e^{-4}$ and an initial learning rate of $1e^{-3}$, reducing the learning rate by a factor of 2 every 20 epochs. Pytorch's MSELoss was used to compute the loss between ground truth and FNO-predicted velocity models. At the end of the first, second and third Fourier Operator Layer we used 2D dropout to omit some of the 32-channel uplifted channels (probability $p = 0.1$) for additional regularisation of the FNO model. This means that on every forward pass through the FNO model during training, channels are zeroed out independently with 10% probability, drawn from a Bernoulli distribution (PyTorch 2025). This mitigated for slight overfitting of the test set.

APPENDIX D: FNO INVERSION PERFORMANCE ON 2D-RANDOMISED EARTHQUAKE GATHER INPUTS

As a further completeness exercise, we investigated how scrambling the earthquake gather FNO inputs affects the predicted velocity model. Specifically, we wanted to ensure that the FNO did not return a velocity model with increasing values as a function of depth, for input data tensors without any physical meaning. Figure A3 shows the same earthquake gather used in figure 6 which has been randomised in space and time for each component of ground velocity, before being t^2 normalised and encoded for FNO input. The resulting FNO velocity model prediction (figure A3c) contains alternating regions of high and low velocity that do not correlate with depth in the model. This result confirms that the trained FNO model objects shown here are not over-reliant on depthwise increasing velocity gradients when inverting for velocity structure from arbitrary data. It further implies that the FNO is sensitive toe particular combinations of ballistic arrivals and coda waves in a given earthquake gather.

APPENDIX E: DEMONSTRATING WAVEFIELD CODA CONTRIBUTIONS TO THE OVERALL FNO VELOCITY MODEL INFERENCE (CONTD.)

This paper has been produced using the Blackwell Scientific Publications GJI LaTeX2e class file.

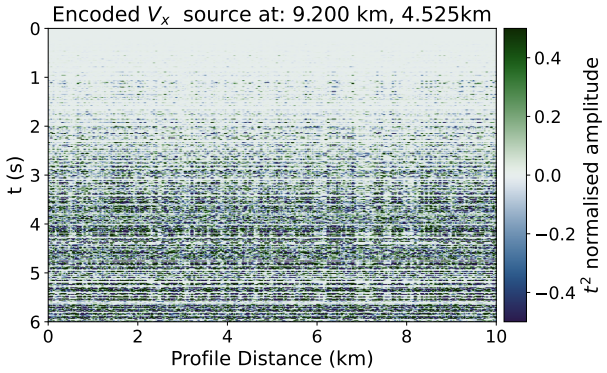

(a)

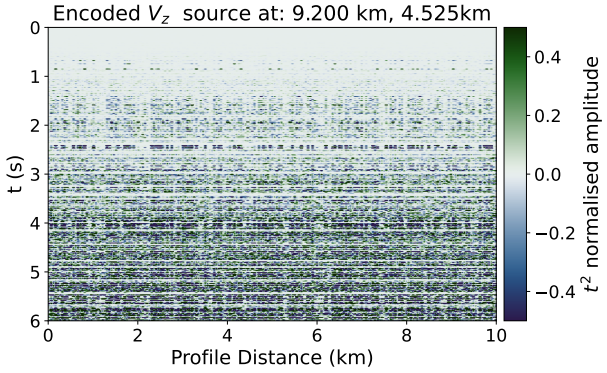

(b)

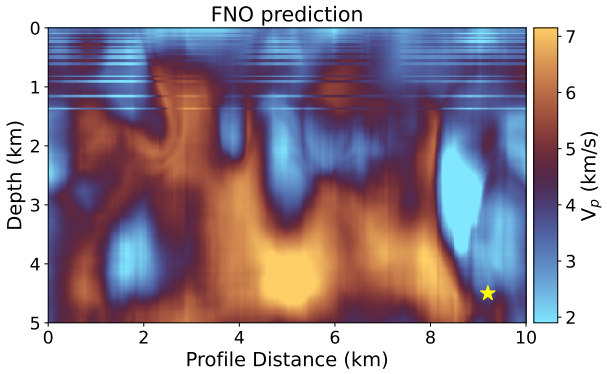

(c)

Figure A3. Spatio-temporal randomisation of the earthquake gathers shown in figure 6 and the subsequent FNO velocity model prediction. Shown are prediction results for an FNO trained on 30,000 core velocity models plus 2,000 additional von Kármán-based models embedded with sub-vertical low velocity zones (3km/s) and approximate columnar width of 500m. **(a)** Randomised, t^2 normalised and encoded V_x earthquake gather **(b)** Spatio-temporally randomised, t^2 normalised and encoded V_z earthquake gather **(c)** FNO prediction. The source location is marked by the yellow star.

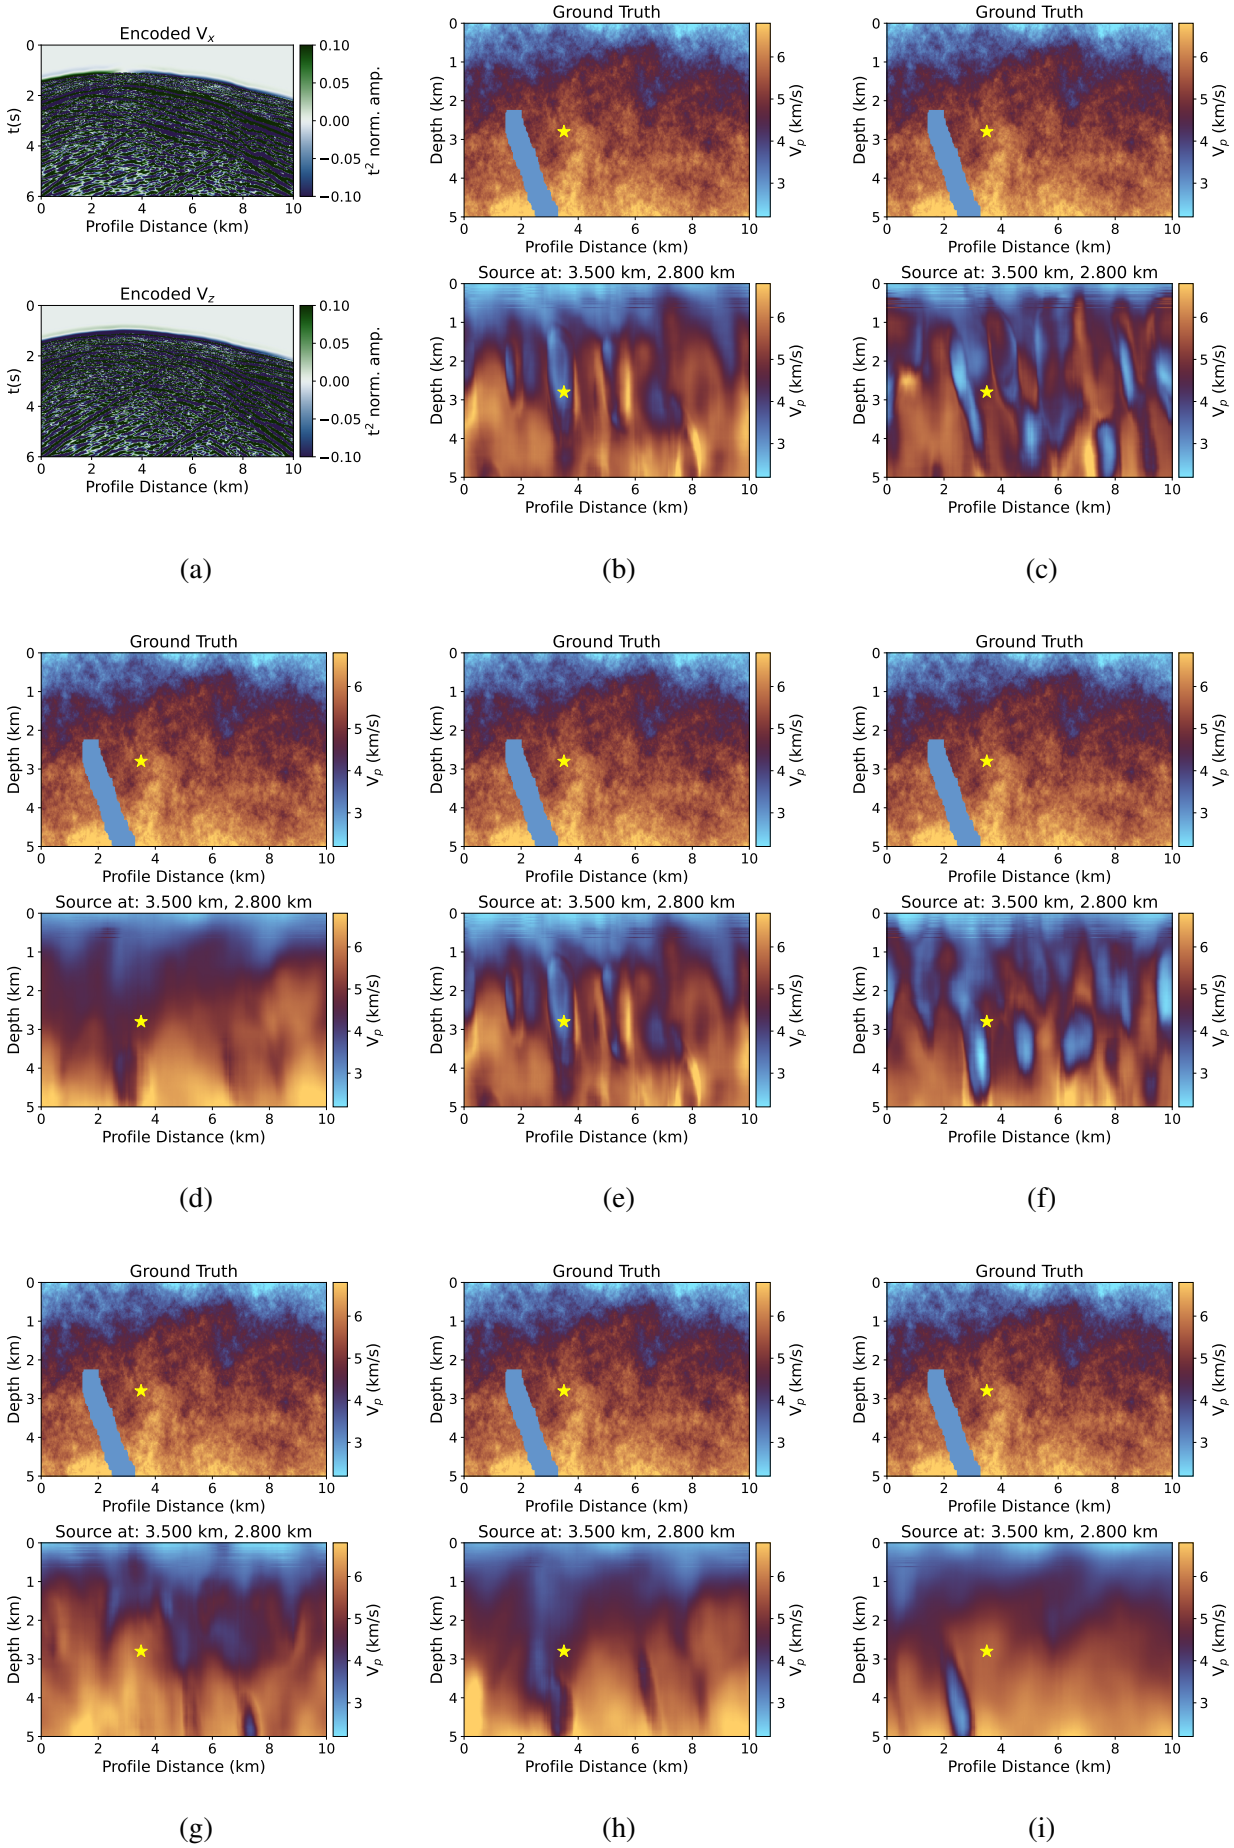

Figure A4. FNO performance with muted earthquake gather coda for different time windows. **(a)** Original unmuted gather **(b)** 0-2s mute **(c)** 2-4s **(d)** 4-6s **(e)** 1-2s **(f)** 2-3s **(g)** 3-4s **(h)** 4-5s **(i)** 5-6s.