
Velocity continuation with Fourier neural operators for accelerated uncertainty quantification

Ali Siahkoobi, Mathias Louboutin, and Felix J. Herrmann
Georgia Institute of Technology
{alisk, mlouboutin3, felix.herrmann}@gatech.edu

Abstract

Seismic imaging is an ill-posed inverse problem that is challenged by noisy data and modeling inaccuracies—due to errors in the background squared-slowness model. Uncertainty quantification is essential for determining how variability in the background models affects seismic imaging. Due to the costs associated with the forward Born modeling operator as well as the high dimensionality of seismic images, quantification of uncertainty is computationally expensive. As such, the main contribution of this work is a survey-specific Fourier neural operator surrogate to velocity continuation that maps seismic images associated with one background model to another virtually for free. While being trained with only 200 background and seismic image pairs, this surrogate is able to accurately predict seismic images associated with new background models, thus accelerating seismic imaging uncertainty quantification. We support our method with a realistic data example in which we quantify seismic imaging uncertainties using a Fourier neural operator surrogate, illustrating how variations in background models affect the position of reflectors in a seismic image.

1 Introduction

Seismic imaging involves estimating the short-wavelength component of the Earth’s subsurface squared-slowness model—known as the seismic image—given shot records and an estimation of the smooth background squared-slowness model. This linearized imaging problem is challenged by the computationally expensive forward operator as well as presence of measurements noise, linearization errors, modeling errors, and the nontrivial nullspace of the linearized forward Born modeling operator [1–3]. These challenges highlight the importance of uncertainty quantification (UQ) in seismic imaging, where instead of finding one seismic image estimate, a distribution of seismic images is obtained that explains the observed data [4], consequently reducing the risk of data overfit and enabling UQ [5–15].

The seismic imaging uncertainty can be attributed to two main sources [16–18]: (1) errors in the data, which include measurement and linearization errors; and (2) modeling errors, which include errors in the estimation of the background squared-slowness model. In this paper, we focus on uncertainties with respect to the background model as it has the main contributing factor to imaging uncertainty due to its effect on reflector positioning [18–20]. To this end, we assume there are no measurement or linearization errors and consider the map from shot records to the seismic image to be represented by the deterministic reverse-time migration (RTM) algorithm. In this setup, quantifying the uncertainty in seismic imaging—due to errors in the background model—involves computing numerous RTMs with all the background model posterior samples [8, 9]. Due to the high-dimensionality of seismic images, the number of seismic imaging posterior samples required to obtain accurate estimations of the imaging posterior moments is large, making UQ with respect to the background model computationally expensive for 2D and unfeasible for large 3D problems. These costs can be alleviated via velocity continuation methods [21–25] that map seismic images associated with one background model to another, without directly solving the imaging problem. Velocity continuation methods are

designed to be frugal compared to computing a new RTM for each new background model, which makes them suitable for large-scale seismic imaging uncertainty quantification [19, 20].

While several velocity continuation methods have been proposed [21–25], they typically involve solving partial differential equations (PDEs), which can be still costly in the context of seismic imaging UQ. To address this challenge, we propose a neural network surrogate for velocity continuation that is capable of mapping seismic images associated with one background model to another with negligible computational cost. Motivated by the success of Fourier neural operators [FNOs, 26] in approximating the solution operator of PDEs [27–29], we chose them as the architecture for our neural network surrogate. Due to our main interest in accelerating velocity continuation in the context of UQ, we train a survey-specific FNO that acts as a surrogate for velocity continuation for the specific survey at hand. This choice, while not offering generalizations across different survey areas in the Earth, can speed up seismic imaging UQ for the survey at hand, which involves computing seismic images associated with many background model posterior samples. We show that the FNO can be trained using 200 pairs of background and seismic image pairs, making the survey-specific training procedure computationally viable. To scale this method to industry size problems, transfer learning [30] can further reduce the upfront costs of training the FNO. After training, the FNO can be used to obtain samples from the imaging posterior almost free of cost.

In the next section, we describe seismic imaging by introducing the forward Born modeling operator, through which the seismic image relates to the background model. Next, we define velocity continuation, followed by describing our proposed FNO-based approach for accelerating seismic imaging UQ. Finally, we evaluate the performance of the trained neural operator on a realistic dataset, and we demonstrate how the imaging uncertainty affects the positioning of the reflectors in the seismic image.

2 Theory

We introduce a deep-network surrogate for velocity continuation in order to enable faster quantification of uncertainty—due to errors in the background velocity model—in seismic imaging. We begin with an introduction to seismic imaging and the linearized forward model associated with it.

2.1 Seismic imaging

The inverse problem that we tackle involves the process of estimating the short-wavelength component of the Earth’s unknown subsurface squared-slowness model given measurements recorded at the surface. This problem, also known as seismic imaging, can be formulated as a linear inverse problem by linearizing the nonlinear relationship between shot records and the squared-slowness model, governed by the wave-equation. In its simplest acoustic form, the linearization with respect to the slowness model—around a background squared slowness model \mathbf{m}_0 —leads to a linear inverse problem for estimating the ground truth seismic image $\delta\mathbf{m}^*$ with the following forward model,

$$\delta\mathbf{d}_i = \mathbf{J}(\mathbf{m}_0, \mathbf{q}_i)\delta\mathbf{m}^*. \quad (1)$$

In the above expression, $\delta\mathbf{d} = \{\delta\mathbf{d}_i\}_{i=1}^{n_s}$ are n_s linearized shot records and $\mathbf{J}(\mathbf{m}_0, \mathbf{q}_i)$ represents the linearized Born scattering operator. This operator is parameterized by the source signature \mathbf{q}_i and the background squared-slowness model \mathbf{m}_0 , which is typically estimated in the previous inversion steps [8, 9]. In this work, we focus on potential inaccuracies in the background model, the main source of variability and uncertainty in seismic imaging [18–20]. Therefore, we assume there are not measurement and linearization errors, which leads to a deterministic mapping from $\{\delta\mathbf{d}_i\}_{i=1}^{n_s}$ to $\delta\mathbf{m}$ for a given background model. We use RTM for this deterministic map, which involves applying the adjoint linearized Born scattering operator to the linearized shot records for all source experiments,

$$\delta\mathbf{m}_{\text{RTM}} = \sum_{i=1}^{n_s} \mathbf{J}(\mathbf{m}_0, \mathbf{q}_i)^\top \delta\mathbf{d}_i. \quad (2)$$

Since this map is deterministic, the uncertainty in the background model can be translated into uncertainty in seismic imaging by evaluating Equation 2 for the background models sampled from the posterior distribution $p(\mathbf{m}_0 \mid \mathbf{d})$, where \mathbf{d} represents the shot records with no linearization (the field data) [8, 9]. Bayesian inference involving high-dimensional posterior distributions, for example seismic imaging and full-waveform inversion, requires many samples from the posterior distribution for accurate Monte Carlo approximation of high-dimensional integrals [31]. As a result, the mapping in Equation 2 must be evaluated over numerous background model samples from $p(\mathbf{m}_0 \mid \mathbf{d})$, which

is computationally expensive due to costs of applying $\mathbf{J}(\mathbf{m}_0, \mathbf{q}_i)^\top$, $i = 1, \dots, n_s$. To address this computational challenge, our proposed method involves a FNO-based velocity continuation approach that is capable of mapping seismic images associated with one background model to another, effectively replacing a costly demigration-migration. We describe this approach in the next section.

2.2 Velocity continuation

At its core, velocity continuation is a process which alters a seismic image based on changes in the background model [21, 22, 24, 25]. Using this technique, an initial seismic image associated with an initial background model \mathbf{m}_{init} is altered to approximate the seismic image associated with a target background model $\mathbf{m}_{\text{target}}$ without computing Equation 2. This process can be interpreted as a map [23]:

$$\mathcal{T}_{(\mathbf{m}_{\text{init}}, \mathbf{m}_{\text{target}})} : \delta\mathcal{M} \rightarrow \delta\mathcal{M}, \quad (3)$$

where $\delta\mathcal{M}$ denotes the space of seismic images, and the velocity continuation map is parameterized by \mathbf{m}_{init} and $\mathbf{m}_{\text{target}}$. We take advantage of recent advances in deep learning to train a surrogate model for velocity continuation that approximates $\mathcal{T}_{(\mathbf{m}_{\text{init}}, \mathbf{m}_{\text{target}})} = \mathbf{J}(\mathbf{m}_{\text{target}}, \mathbf{q}_i)^\top (\mathbf{J}(\mathbf{m}_{\text{init}}, \mathbf{q}_i)^\top)^\dagger$, which can be evaluated virtually for free instead of solving four PDEs. Through this approach, we are able to significantly accelerate velocity continuation, trading wave-equation solves for a simple neural network inference. This is of considerable importance in the context of seismic imaging UQ. Before demonstrating the benefits of our approach, we introduce FNOs in the context of velocity continuation.

2.3 Fourier neural operators for velocity continuation

In light of their success in learning mesh-free solution operators to PDEs [27–29], we choose FNOs [26] as a surrogate model for velocity continuation—a process that can be interpreted as a double linearized PDE solve. The main components of FNOs are the Fourier layers, which involve a Fourier transform over the spatial dimensions of their input, followed by a learned pointwise multiplication and an inverse Fourier transform. These layers act as long-kernel convolutional layer, akin to pseudo-spectral methods, which explains the representation power of FNOs [29]. To train a FNO as a surrogate model for velocity continuation, we define it as a map

$$\mathcal{G}_{\mathbf{w}} : \mathcal{M} \times \delta\mathcal{M} \rightarrow \delta\mathcal{M}, \quad (4)$$

where $\mathcal{G}_{\mathbf{w}}$ denotes the FNO with weights \mathbf{w} , and \mathcal{M} is the space of background models. In words, we design $\mathcal{G}_{\mathbf{w}}$ as a neural network that takes as input the target background model and the initial seismic image, and outputs the target seismic image. This choice is analogous to the structure of the velocity continuation map $\mathcal{T}_{(\mathbf{m}_{\text{init}}, \mathbf{m}_{\text{target}})}$ (cf. Equation 3), except that we fix the initial background model \mathbf{m}_{init} to an arbitrary background model posterior sample, hence making the dependence of $\mathcal{G}_{\mathbf{w}}$ to \mathbf{m}_{init} implicit. With this choice of inputs and output for $\mathcal{G}_{\mathbf{w}}$, the FNO’s task is to perturb the input initial seismic image according to the provided target background model in order to predict the target seismic image.

Due to our interest in accelerating velocity continuation in the context of UQ, we train a survey-specific FNO that acts as a surrogate for velocity continuation for the specific survey at hand. This choice, while not offering generalizations across different survey areas in the Earth, can speed up seismic imaging UQ for the survey at hand, which involves computing seismic images associated with many posterior background model samples $\mathbf{m}_0 \sim p(\mathbf{m}_0 \mid \mathbf{d})$. By training the FNO on a small number of these background model and seismic image pairs, i.e., approximately 200, we can accelerate the velocity continuation process for the rest of the background model posterior samples while limiting the risk of introducing generalization errors due to the strong heterogeneity of Earth. To achieve this, we construct a set of N training input-output pairs in the form of

$$\left\{ \left((\mathbf{m}_0^{(i)}, \delta\mathbf{m}_{\text{init}}), \delta\mathbf{m}_{\text{RTM}}^{(i)} \right) \mid i = 1, \dots, N \right\}, \quad (5)$$

where $(\mathbf{m}_0^{(i)}, \delta\mathbf{m}_{\text{init}})$ is the input target background and initial seismic image training pair, and $\delta\mathbf{m}_{\text{RTM}}^{(i)}$ is the associated target seismic image. Training involves minimizing the squared ℓ_2 -norm of the difference between the FNO output and the target seismic image with respect to FNO weights,

$$\mathbf{w}^* = \arg \min_{\mathbf{w}} \frac{1}{N} \sum_{i=1}^N \|\mathcal{G}_{\mathbf{w}}(\mathbf{m}_0^{(i)}, \delta\mathbf{m}_{\text{init}}) - \delta\mathbf{m}_{\text{RTM}}^{(i)}\|_2^2. \quad (6)$$

We solve optimization problem 6 with the Adam stochastic optimization algorithm [32]. After training, the trained FNO approximates the velocity continuation map for a fixed homogeneous initial background model, i.e.,

$$\mathcal{G}_{\mathbf{w}^*}(\mathbf{m}_{\text{target}}, \delta\mathbf{m}_{\text{init}}) \approx \mathcal{T}_{(\mathbf{m}_{\text{init}}, \mathbf{m}_{\text{target}})}(\delta\mathbf{m}_{\text{init}}) = \delta\mathbf{m}_{\text{target}}, \quad (7)$$

where \mathbf{m}_{init} is fixed to an arbitrary background model posterior sample and $\delta\mathbf{m}_{\text{target}}$ denotes the seismic image associated with $\delta\mathbf{m}_{\text{init}}$. Given background model samples $\mathbf{m}_0 \sim p(\mathbf{m}_0 | \mathbf{d})$ as input, the FNO outputs samples from the imaging posterior, $p(\mathbf{m}_0 | \delta\mathbf{d})$. This process accelerates seismic imaging uncertainty quantification as no further RTMs (Equation 2) need to be computed. Training the FNO with as little as 200 training pairs makes the survey-specific training procedure computationally viable. To further accelerate the process, training of the FNO can be started at the same time as the background model posterior sampling phase, using the already collected posterior samples as training data. In the next section, we show the results of approximating the velocity continuation map with a FNO via a quasi-real seismic experiment.

3 Numerical experiments

The purpose of the presented numerical experiments here is to demonstrate the ability to approximate the velocity continuation map (Equation 3) using a FNO. We further show how this trained FNO can be used for UQ by showing the effect of imaging uncertainty on the positioning of the reflectors. We begin by describing the training setup, including the seismic acquisition geometry.

3.1 Acquisition geometry and training configuration

Our examples involve imaging a 2D subset of the [Parihaka](#) [33, 34] prestack Kirchhoff migration field dataset. We use this 2D section to create linearized data according to the linear forward model in Equation 1, where we consider no measurement and linearization errors. The model is discretized on a 12.5 m vertical and 20 m horizontal grid, and the data is acquired with 102 equally spaced sources and 204 fixed receivers. We use a Ricker wavelet with a central frequency of 30 Hz as the source signature. For presentation purposes, we augment a 125 m water column on top of these models to limit the near source imaging artifacts. We create background models by assuming access to an oracle, which provides geologically consistent background models with the 2D section. Given 200 background models, we migrate the simulated seismic data via Equation 2 to obtain the corresponding seismic images. Figures 1a and 1b show the vertical profile of five randomly selected background and seismic images at 2.50 km, respectively. These images, displaying strong amplitude and phase differences, highlight the importance of quantifying the uncertainty in imaging when dealing with errors in the background model. To reduce the costs associated with UQ, we train the FNO surrogate using the 200 training pairs (cf. Equation 5) for 500 epochs with the Adam optimizer [32].

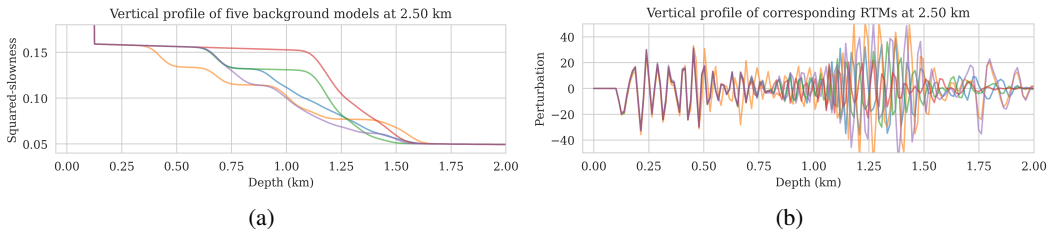


Figure 1: Variation among five (a) background; and (b) seismic images (RTMs), plotted as a vertical profile at 2.50 km.

3.2 Results

To evaluate the accuracy of the trained FNO in predicting seismic images, we compare its output to the target seismic image, i.e., the RTM image obtained via Equation 2 when using the target background model (Figure 2a) to parameterize the Born scattering forward operator. This is conducted over the testing dataset, which is derived using the same procedure as the training dataset using the background model creating oracle. Figures 2b and 2c show the target and predicted seismic images, respectively, and Figures 2d includes the difference between them. We observe that the network has accurately predicted the target image, where errors are mostly due to amplitude differences. The accuracy of

the prediction in terms of phase and amplitude can be further confirmed by focusing on two vertical profiles, at 2.50 km (Figure 2e) and 4.125 km (Figure 2f) horizontal locations, which include regions with fault and tortuous reflectors. We use the trained FNO for UQ by providing testing background models and predicting the associated seismic images. We visualize the obtained uncertainties by showing the variability in the location of reflectors, determined via an automatic horizon tracking software [35]. We pass the seismic images predicted by the FNO to the horizon tracker for 25 selected horizons. As a result, we obtain multiple instances of each horizons, from which we compute pointwise mean and standard deviations. Figure 3 indicates the result where the solid lines correspond to the mean among different instances of each horizons and shaded areas indicate the mean plus and minus the pointwise standard deviation. These shaded areas indicate uncertainties in the location of the reflectors, which are due to the variability in the background models. As expected, we find a general increase of uncertainty with depth. We also observe that the areas of high uncertainty are correlated with areas of poor illumination, faults and tortuous reflectors.

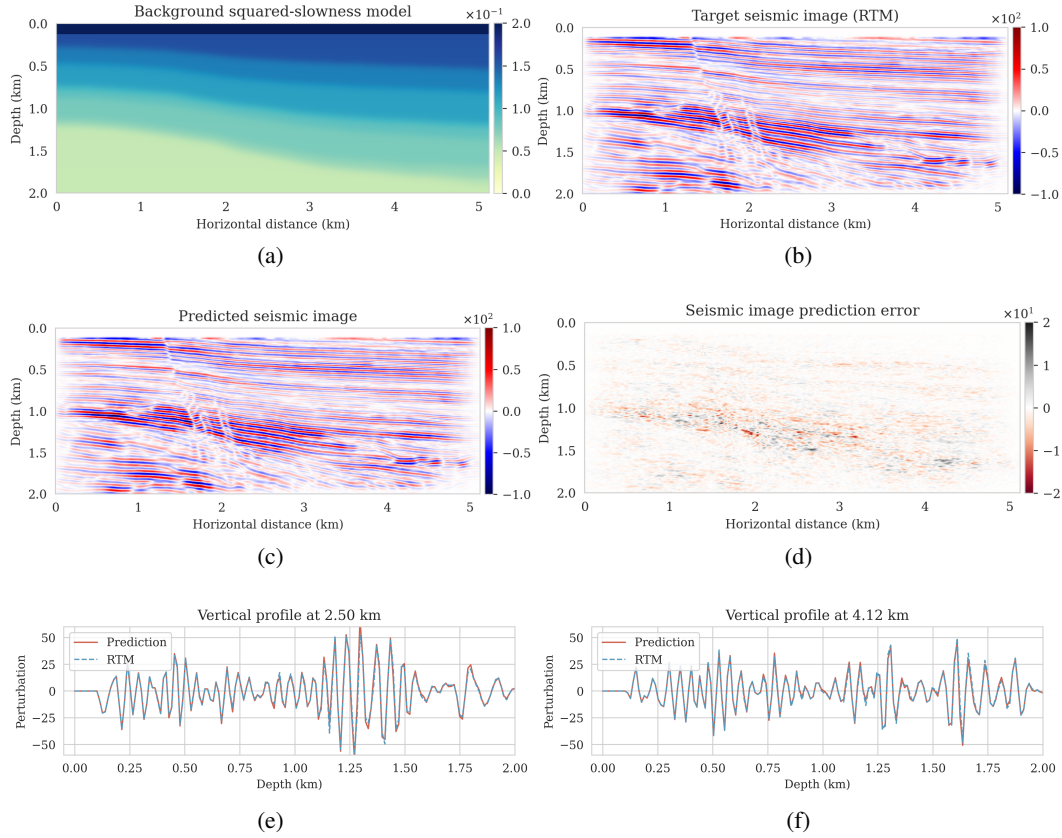


Figure 2: Velocity continuation with FNOs. Target (a) background; and (b) seismic images. (c) Predicted seismic image with the FNO. (d) Difference between target and prediction. (e) and (f) Vertical profile comparisons between target (dashed blue) and predicted (red) seismic images.

In this example, we use Devito [36, 37] for the wave-equation based simulations. We based our PyTorch FNO implementation on the [original implementation](#). The code to reproduce our results are made available on [GitHub](#).

4 Conclusions and discussion

Quantifying the uncertainty in seismic imaging due to errors in the background model involves solving many seismic imaging problems that vary in the parameterization of the background model of the forward operator. To reduce the computational cost of this process—mainly due to the computational costs of the forward operator—we proposed to train a survey-specific Fourier neural operator surrogate that mimics velocity continuation. This surrogate model maps seismic images associated with one background model to another virtually for free, which has the benefit of accelerating uncertainty

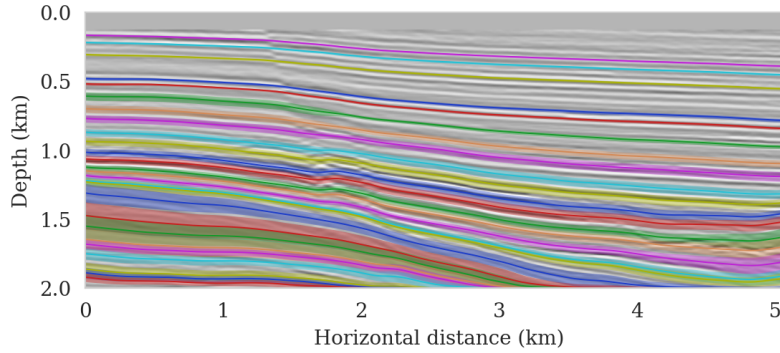


Figure 3: Uncertainty in the tracked horizons.

quantification. We showed that this surrogate model can be trained with as few as 200 training pairs while still providing a good seismic image prediction accuracy. Further research is required in training a reliable global surrogate, being able to generalize across other survey areas and more realistic physics.

5 Acknowledgment

This research was carried out with the support of Georgia Research Alliance and partners of the ML4Seismic Center.

References

- [1] Gilles Lambaré, Jean Virieux, Raul Madariaga, and Side Jin. Iterative asymptotic inversion in the acoustic approximation. *Geophysics*, 57(9):1138–1154, 1992.
- [2] Gerard T Schuster. Least-squares cross-well migration. In *SEG Technical Program Expanded Abstracts 1993*, pages 110–113. Society of Exploration Geophysicists, 1993.
- [3] Tamas Nemeth, Chengjun Wu, and Gerard T Schuster. Least-squares migration of incomplete reflection data. *GEOPHYSICS*, 64(1):208–221, 1999. doi: 10.1190/1.1444517.
- [4] Albert Tarantola. *Inverse problem theory and methods for model parameter estimation*. SIAM, 2005. ISBN 978-0-89871-572-9. doi: 10.1137/1.9780898717921.
- [5] Alberto Malinverno and Robert L Parker. Two ways to quantify uncertainty in geophysical inverse problems. *GEOPHYSICS*, 71(3):W15–W27, 2006. doi: 10.1190/1.2194516.
- [6] James Martin, Lucas C. Wilcox, Carsten Burstedde, and OMAR Ghattas. A Stochastic Newton MCMC Method for Large-scale Statistical Inverse Problems with Application to Seismic Inversion. *SIAM Journal on Scientific Computing*, 34(3):A1460–A1487, 2012. URL <http://epubs.siam.org/doi/abs/10.1137/110845598>.
- [7] Anandaroop Ray, Sam Kaplan, John Washbourne, and Uwe Albertin. Low frequency full waveform seismic inversion within a tree based Bayesian framework. *Geophysical Journal International*, 212(1):522–542, 10 2017. ISSN 0956-540X. doi: 10.1093/gji/ggx428. URL <https://doi.org/10.1093/gji/ggx428>.
- [8] Hejun Zhu, Siwei Li, Sergey Fomel, Georg Stadler, and Omar Ghattas. A bayesian approach to estimate uncertainty for full-waveform inversion using a priori information from depth migration. *Geophysics*, 81(5):R307–R323, 2016.
- [9] Zhilong Fang, Curt Da Silva, Rachel Kuske, and Felix J. Herrmann. Uncertainty quantification for inverse problems with weak partial-differential-equation constraints. *GEOPHYSICS*, 83(6):R629–R647, 2018. doi: 10.1190/geo2017-0824.1.
- [10] Georgia K Stuart, Susan E Minkoff, and Felipe Pereira. A two-stage Markov chain Monte Carlo method for seismic inversion and uncertainty quantification. *GEOPHYSICS*, 84(6):R1003–R1020, 11 2019. doi: 10.1190/geo2018-0893.1.

- [11] Felix J. Herrmann, Ali Siahkoohi, and Gabrio Rizzuti. Learned imaging with constraints and uncertainty quantification. In *Neural Information Processing Systems (NeurIPS) 2019 Deep Inverse Workshop*, 12 2019. URL <https://arxiv.org/pdf/1909.06473.pdf>.
- [12] Ali Siahkoohi, Gabrio Rizzuti, and Felix J. Herrmann. A deep-learning based bayesian approach to seismic imaging and uncertainty quantification. *82nd EAGE Conference and Exhibition 2020*, 6 2020. URL <https://slim.gatech.edu/Publications/Public/Submitted/2020/siahkoohi2020EAGEdlb/siahkoohi2020EAGEdlb.html>.
- [13] Ali Siahkoohi, Gabrio Rizzuti, and Felix J. Herrmann. Uncertainty quantification in imaging and automatic horizon tracking—a Bayesian deep-prior based approach. In *SEG Technical Program Expanded Abstracts 2020*, pages 1636–1640, 9 2020. doi: 10.1190/segam2020-3417560.1.
- [14] Ali Siahkoohi and Felix J Herrmann. Learning by example: fast reliability-aware seismic imaging with normalizing flows. In *First International Meeting for Applied Geoscience & Energy*, pages 1580–1585. Society of Exploration Geophysicists, 2021.
- [15] Ali Siahkoohi, Gabrio Rizzuti, and Felix J Herrmann. Deep bayesian inference for seismic imaging with tasks. *arXiv preprint arXiv:2110.04825*, 2021.
- [16] Pierre Thore, Arben Shtuka, Magali Lecour, Taoufik Ait-Ettajer, and Richard Cognot. Structural uncertainties: Determination, management, and applications . *Geophysics*, 67(3):840–852, 2002.
- [17] Konstantin Osypov, Yi Yang, Aimé Fournier, Natalia Ivanova, Ran Bachrach, Can Evren Yarman, Yu You, Dave Nichols, and Marta Woodward. Model-uncertainty quantification in seismic tomography: method and applications. *Geophysical Prospecting*, 61(6-Challenges of Seismic Imaging and Inversion Devoted to Goldin):1114–1134, 2013.
- [18] Gregory Ely, Alison Malcolm, and Oleg V. Poliannikov. Assessing uncertainties in velocity models and images with a fast nonlinear uncertainty quantification method. *GEOPHYSICS*, 83(2):R63–R75, 2018. doi: 10.1190/geo2017-0321.1. URL <https://doi.org/10.1190/geo2017-0321.1>.
- [19] Sergey Fomel and Evgeny Landa. Structural uncertainty of time-migrated seismic images. *Journal of Applied Geophysics*, 101:27–30, 2014.
- [20] Oleg V Poliannikov and Alison E Malcolm. The effect of velocity uncertainty on migrated reflectors: Improvements from relative-depth imaging. *Geophysics*, 81(1):S21–S29, 2016.
- [21] Sergey Fomel. Time-migration velocity analysis by velocity continuation. *Geophysics*, 68(5):1662–1672, 2003.
- [22] Sergey Fomel. Velocity continuation and the anatomy of residual prestack time migration. *Geophysics*, 68(5):1650–1661, 2003.
- [23] Anton A Duchkov and Maarten V De Hoop. Velocity continuation in the downward continuation approach to seismic imaging. *Geophysical Journal International*, 176(3):909–924, 2009.
- [24] Tristan van Leeuwen and Felix J Herrmann. Wave-equation Extended Images-Computation and Velocity Continuation. In *74th EAGE Conference and Exhibition incorporating EUROPEC 2012*, pages cp–293. European Association of Geoscientists & Engineers, 2012.
- [25] Mengmeng Yang, Marie Graff, Rajiv Kumar, and Felix J Herrmann. Low-rank representation of omnidirectional subsurface extended image volumes. *Geophysics*, 86(3):S165–S183, 2021.
- [26] Zongyi Li, Nikola Borislavov Kovachki, Kamyar Azizzadenesheli, Burigede liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier Neural Operator for Parametric Partial Differential Equations. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=c8P9NQVtmnO>.
- [27] J Quetzalcóatl Toledo-Marín, Geoffrey Fox, James P Sluka, and James A Glazier. Deep learning approaches to surrogates for solving the diffusion equation for mechanistic real-world simulations. *Frontiers in Physiology*, 12, 2021.

- [28] T Konuk and J Shragge. Physics-guided deep learning using fourier neural operators for solving the acoustic vti wave equation. In *82nd EAGE Annual Conference & Exhibition*, volume 2021, pages 1–5. European Association of Geoscientists & Engineers, 2021.
- [29] Nikola Kovachki, Samuel Lanthaler, and Siddhartha Mishra. On universal approximation and error bounds for fourier neural operators. *Journal of Machine Learning Research*, 22:Art–No, 2021.
- [30] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How Transferable Are Features in Deep Neural Networks? In *Proceedings of the 27th International Conference on Neural Information Processing Systems*, pages 3320–3328, 2014. URL <http://dl.acm.org/citation.cfm?id=2969033.2969197>.
- [31] Andrew Gelman and Donald B. Rubin. Inference from Iterative Simulation Using Multiple Sequences. *Statistical Science*, 7(4):457–472, 1992. ISSN 08834237. URL <http://www.jstor.org/stable/2246093>.
- [32] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. URL <https://arxiv.org/pdf/1412.6980.pdf>.
- [33] Veritas. Parihaka 3D Marine Seismic Survey - Acquisition and Processing Report. Technical Report New Zealand Petroleum Report 3460, New Zealand Petroleum & Minerals, Wellington, 2005.
- [34] WesternGeco. Parihaka 3D PSTM Final Processing Report. Technical Report New Zealand Petroleum Report 4582, New Zealand Petroleum & Minerals, Wellington, 2012.
- [35] Xinming Wu and Sergey Fomel. Least-squares horizons with local slopes and multi-grid correlations. *GEOPHYSICS*, 83:IM29–IM40, 2018. doi: 10.1190/geo2017-0830.1. URL <https://doi.org/10.1190/geo2017-0830.1>.
- [36] F. Luporini, M. Lange, M. Louboutin, N. Kukreja, J. Hückelheim, C. Yount, P. Witte, P. H. J. Kelly, F. J. Herrmann, and G. J. Gorman. Architecture and performance of devito, a system for automated stencil computation. *CoRR*, abs/1807.03032, jul 2018. URL <http://arxiv.org/abs/1807.03032>.
- [37] M. Louboutin, M. Lange, F. Luporini, N. Kukreja, P. A. Witte, F. J. Herrmann, P. Velesko, and G. J. Gorman. Devito (v3.1.0): an embedded domain-specific language for finite differences and geophysical exploration. *Geoscientific Model Development*, 12(3):1165–1187, 2019. doi: 10.5194/gmd-12-1165-2019. URL <https://www.geosci-model-dev.net/12/1165/2019/>.