PROCEEDINGS, INDONESIAN PETROLEUM ASSOCIATION
Fiftieth Annual Convention & Exhibition, May 2026

Application of Fourier Neural Operators (FNO) to Achieve Fast-cycle Update of Velocity Model

Jyoti Kumar[1]
Øystein Kormso[1]
Nizar Chemingui[1]

[1]TGS

ABSTRACT

This paper looks at how deep learning can improve the initial Full Waveform Inversion (FWI) velocity model and reduce the need for traditional tomography in velocity model building (VMB), which can be time-consuming. The main goal is to make the VMB process faster and more efficient, but deep learning also opens new possibilities — allowing better data handling, quicker convergence, and improved performance in areas where ray-based methods struggle, especially in complex, high-contrast environments.

In this paper, we present a deep neural network designed to function as a tomographic back-projection operator, effectively replacing conventional tomographic inversion engines. The network architecture is based on Fourier Neural Operators (FNOs), which enable the learning of long-range dependencies within the data. This allows the model to map velocity errors independently of where the effects of velocity errors are observed. Training is performed using migrated gathers without the need for residual moveout (RMO) picks, horizon interpretations, or salt masks, allowing the network to capture a wide spectrum of both smooth and sharp velocity features. The implementation operates iteratively in the image domain, requiring re-migration after each prediction to progressively refine the velocity model.

We demonstrate the capability of this approach to both accelerate and enhance the velocity model building (VMB) workflow using field data examples from the offshore Agung and Norwegian sea area. The results show that the FNO-based network can effectively condition the initial FWI model by mitigating cycle-skipping and generating a structurally consistent velocity model directly from a 1-D starting profile. The method exhibits substantially faster convergence compared to conventional techniques, attributable to the network's ability to reconstruct the true velocity model from diverse, imperfect initial conditions. Moreover, the approach demonstrates improved performance in high-contrast geological settings, where ray-based tomographic methods often fail to provide reliable solutions.

INTRODUCTION

Velocity Model Building (VMB) is a critical yet computationally and operationally intensive component of the seismic imaging workflow, requiring significant expertise in Full Waveform Inversion (FWI), depth imaging, geological interpretation, and tomography. Although FWI has proven highly effective for deriving accurate velocity models, its performance is strongly dependent on the quality of the initial velocity model. In frontier exploration settings, such models are often poorly constrained and may rely solely on legacy 2D seismic data. As a result, the VMB process remains heavily dependent on manual intervention, creating inefficiencies and bottlenecks in both time and cost.

In recent years, machine learning (ML) has become an increasingly important and integrated component of seismic data processing workflows, with most applications focused on preprocessing and interpretation tasks. In contrast, relatively few studies have applied ML techniques to VMB using field data. This limited adoption is primarily due to the strong variability of subsurface geology, the high dimensionality of seismic acquisition geometries, and the inherent difficulty of capturing long-range velocity dependencies across both image-domain and data-domain gathers.

In this study, we employ a deep neural network architecture combining Fourier Neural Operators (FNOs), convolutional neural networks (CNNs), and fully connected neural networks (FCNs) to estimate velocity errors in the image domain. The network is trained in synthetic datasets to perform tomographic-style updates, effectively replacing conventional reflection tomography engines. We validate the approach using multiple field data examples from different sedimentary basins worldwide, including Agung area located north of Bali. The application demonstrates that the proposed methodology can complement or enhance traditional tomography workflows and provide improved conditioning of the velocity model for FWI.

MOTIVATION

While this work is primarily motivated by accelerating and supporting the VMB sequence, deep learning also offers the potential to interpret and handle the data in fundamentally different ways, helping to overcome limitations inherent in conventional velocity model–building approaches such as reflection tomography and FWI. Conventional reflection tomography depends on residual moveout (RMO) picks and a priori information to guide velocity updates. However, flat gathers do not necessarily imply a correct velocity model. Indicators such as uncollapsed diffractions and event crossings in the common-angle or offset domain provide additional insights on velocity errors. Moreover, ray-based methods are known to perform poorly in high-contrast media, often requiring specialized treatment or interpretation to handle geobodies and other non-smooth features. While FWI is, in principle, capable of addressing these challenges, it remains computationally expensive, particularly when high-frequency elastic modeling is involved. Moreover, both tomography and FWI linearize a highly nonlinear inverse problem through incremental gradient updates, resulting in many iterations and high computational cost.

In this study, a deep neural network is used as a tomographic operator, effectively replacing the traditional tomographic inversion engine. The network is trained solely on migrated gathers, without the use of RMO picks, horizon constraints, or salt masks, and is designed to identify both smooth and sharp velocity variations. This strategy allows the network to recover the velocity model from a range of inaccurate initial models more efficiently and to apply larger velocity updates than those typically achievable with traditional reflection tomography. Furthermore, the proposed framework allows for training with diverse data types, including datasets containing multiples, thereby offering the potential to further accelerate the velocity model–building workflow.

METHODOLOGY

Our methodology adopts a hybrid network architecture that couples integral operator blocks (IOBs) with convolutional neural networks (CNNs). In contrast to standard CNNs, IOBs are explicitly designed to model long-range dependencies in the data – that is, relationships between seismic features whose informative context spans large spatial separations and/or wide time (depth) and offset ranges, beyond the limited receptive field of local convolutions. Seismic data exhibits long-range dependencies, as reflections from deeper events are often influenced by overburden complexity, so shallow structure can imprint on the appearance and kinematics of deeper reflections across the gather/image. The inclusion of Fourier Neural Operators (FNOs) enables non-local feature extraction, global representation learning, and mesh-independent behavior. Such neural operator formulations have shown considerable promise in solving complex and computationally demanding problems, including partial differential equations (Li *et al.*, 2020).

Integral Operator Blocks

Integral operator blocks are computational constructs that provide structured, often modular, representations of integral operators. Such operators map functions to functions through kernel-based integral formulations and are widely used in numerical analysis and computational physics. More recently, they have gained prominence in machine learning as a means of learning continuous operators directly from data. Integral operator blocks implement these mappings as modular units, with each block encapsulates a component of the underlying operator. This modular formulation facilitates numerical implementation, network design, and theoretical analysis, while allowing complex operators to be composed from simpler building blocks.

In seismic imaging, integral operators naturally arise from reformulations of partial differential equations (PDEs) through Green's function representations, which recast local differential equations into global integral equations. Such formulations are often advantageous in the presence of complex geometries or infinite domains. Within this framework, operator blocks enable the decomposition of complex integral expressions into tractable components, with individual blocks representing distinct physical processes such as wave propagation, scattering, or reflection. These blocks can be discretized using appropriate quadrature schemes or numerical integration techniques. In our usage, integral operator blocks are interpreted as network layers that perform inherently non-local operations. Unlike standard convolutional layers, which are local and translation-invariant, these blocks capture complex global structures, a capability that is essential for seismic velocity inversion, where spatially distant regions can exert strong mutual influence.

Fourier Neural Operators

Neural operators constitute a class of deep learning architecture designed to learn mappings between infinite-dimensional function spaces, rather than between finite-dimensional vectors. Unlike conventional neural networks that operate in Euclidean parameter spaces, neural operators accept functions as inputs and outputs, making them well suited for learning solution operators of partial differential equations (PDEs), such as those governing seismic wave propagation and tomography.

Fourier Neural Operators (FNOs) represent a particular kind of neural operator in which the integral kernel is parameterized directly in Fourier space by leveraging Fourier transforms. By performing convolution in the spectral domain, FNOs provide an efficient and expressive architecture capable of capturing complex, non-local interactions inherent in nonlinear dynamical systems. This formulation enables accurate and scalable modeling of challenging PDEs, including those associated with turbulent flows, while offering substantial computational advantages over traditional numerical solvers—often achieving speedups of several orders of magnitude (Li *et al.*, 2023).

FNOs have been successfully applied across a range of scientific and engineering domains, including computational fluid dynamics, computational mechanics, and geosciences. Notable applications include the prediction of multiphase flow in porous media and the simulation of carbon dioxide migration, both of which are central to subsurface and environmental studies. The core mechanism of FNOs involves transforming the input kernel into the Fourier domain, applying a learned spectral kernel, and then subsequently transforming the result back to the original space. Training is performed by optimizing the parameters of the Fourier kernel using gradient-based optimization methods to minimize the error between predicted and reference PDE solutions.

Seismic Reflection Tomography

In the proposed methodology, the machine learning components described above are integrated to efficiently address the central objective of seismic reflection tomography: updating the seismic velocity model such that migrated gathers—typically produced using Kirchhoff migration—are flattened in the image domain. Conventional reflection tomography workflows rely on the picking of dense, multi-offset moveout information, which is subsequently converted into velocity updates through the solution of linearized partial differential equations (Brittan and Yuan, 2005; Luo *et al*., 2014). During this process, subsurface velocities are updated by back-projecting travel-time residuals along ray paths, with the residuals themselves being extracted directly from the seismic data. The strategy used to generate these picks constitutes a critical component of any velocity model–building workflow.

In contrast, the approach presented here eliminates the need for explicit moveout picking, thereby removing one of the primary bottlenecks in tomographic turnaround—namely, manual or semi-manual human intervention. Although reflection tomography is a well-established and mature technology that is often stabilized through *a priori* constraints such as interpreted structure or well control (Luo *et al.*, 2017), such constraints have not yet been incorporated into the FNO-based framework described in this study. This choice does not preclude the future integration of prior information; rather, the incorporation of geological or well-based constraints remains an active area of ongoing development within the proposed workflow.

Training

In the initial phase of this work, the training strategy focused on learning a direct mapping between velocity errors inferred from the data and their manifestation in the image domain (Huang *et al.*, 2023). More recently, the approach has been refined in two keyways: (1) the network has been reformulated to predict velocity errors directly within the image domain, and (2) the training procedure has been expanded through the generation and incorporation of additional labeled datasets. These refinements have led to a framework in which the network takes as input depth-migrated gathers together with the corresponding velocity model and produces a predicted velocity update as output (Crawley *et al.*, 2024). The resulting micro- and macro-scale network architectures are illustrated in Figure 1.

Within each IOB, forward propagation begins by lifting the data, that increases the dimensionality of the input via a fully connected mapping, denoted $v(x)$. The lifted representation is then transformed into the Fourier domain using a forward Fourier Transform (FT), where a subset of the most energetic spectral modes is retained and subsequently mapped back to the spatial domain through an inverse transform. In parallel, the input data is also passed outside the FT route, and the two representations are combined prior to the application of an activation function, thereby completing the FNO layer. The resulting features are normalized and passed to a multilayer perceptron (MLP), another fully connected network. Dropout and skip connections are employed as regularization mechanisms, enabling the construction of deeper networks required to approximate complex nonlinear operators. Multiple IOBs are subsequently interleaved with CNN layers, and final velocity updates are generated through another MLP.

An image-domain formulation has been adopted in this work for several reasons. First, the migrated image resides in the same spatial domain as the velocity model, whereas shot gathers are defined in the data domain. Consequently, image-to-image mapping aligns naturally with the target velocity representation. Operating in the image domain significantly reduces the dimensionality of the input data, thereby simplifying the learning task faced by the network.

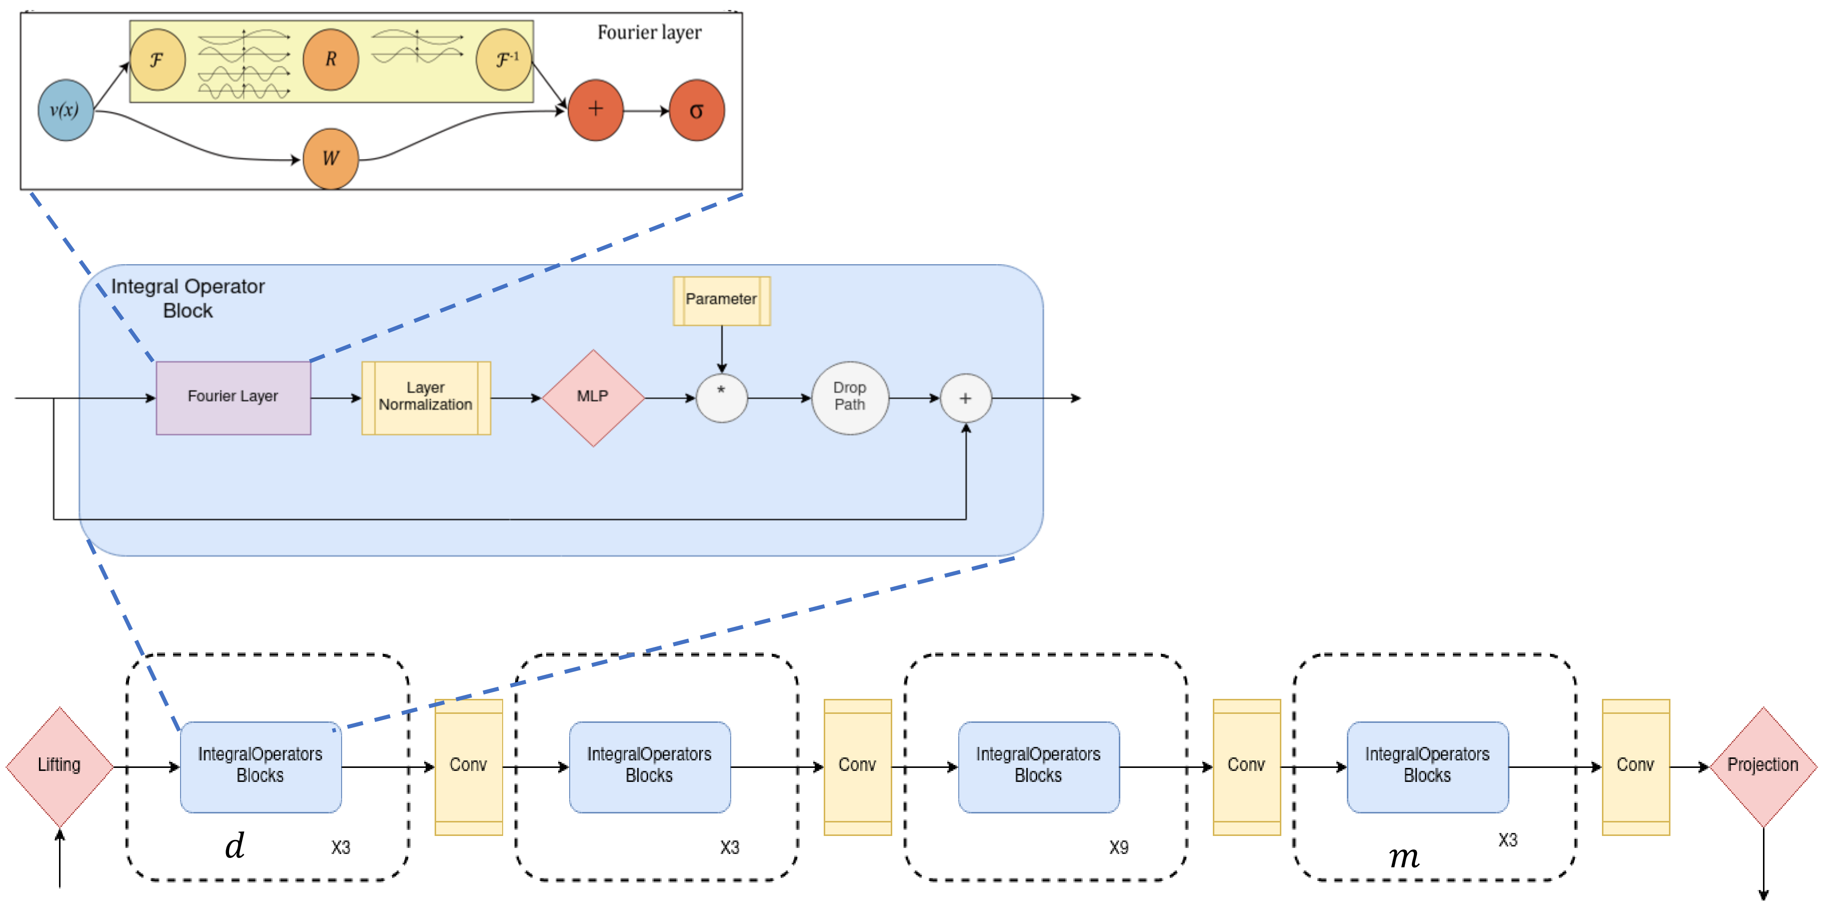

Figure 1—The network architecture, consisting of a series of IOBs and CNNs. The FNOs facilitate the learning of long-range dependencies in the data (Li *et al.*, 2020).

Secondly, the migration process acts as a form of regularization, suppressing noise and stabilizing the input features. The image domain mitigates many of the inconsistencies and irregularities associated with acquisition geometry, reducing the need for the network to explicitly learn such effects. From a generalization perspective, the disparity between noisy field data and comparatively pristine synthetic data poses a challenge for data-driven methods; migration and stacking in the image domain help alleviate this issue by attenuating incoherent noise.

Despite these advantages, the image-domain formulation introduces a substantial null space in the inversion. As in conventional seismic tomography, the solution is primarily constrained by the flatness of migrated gathers, and the resulting velocity model inherits the same fundamental limitations. In particular, because amplitudes are not matched in the data domain, the achievable accuracy cannot rival that of FWI. However, it is well established that the strongly nonlinear nature of FWI requires a sufficiently accurate initial velocity model to converge reliably. The FNO-based approach presented here is therefore intended to provide such a starting model, enabling subsequent high-fidelity inversion methods to be applied more effectively.

The network is trained using synthetic two-dimensional datasets generated from approximately 50,000 distinct velocity models that span a wide range of geological scenarios. These include variations in water depth and background sediment velocity trends, as well as the presence of salt bodies and low-velocity anomalies with diverse geometries and scales. Data augmentation is achieved by introducing varying levels of velocity perturbations relative to each true model.

The proposed methodology is implemented as an iterative and interactive workflow. Following an initial prediction, the updated velocity model is used to perform a new migration, and the resulting gathers serve as input for the subsequent prediction cycle. The total number of iterations required depends on the complexity of the problem; however, practical experience indicates that approximately two to five iterations are generally sufficient to recover the macro velocity model.

RESULTS

In the first example, we evaluate the ability of the trained network to construct an initial FWI model starting from a simple one-dimensional velocity gradient. Figures 2a and 2c present the initial velocity model, overlaid with the single-shot FWI gradient, along with a comparison between the observed seismic data and synthetics generated using a reflectivity-based wave-equation formulation (Whitmore *et al.*, 2020).

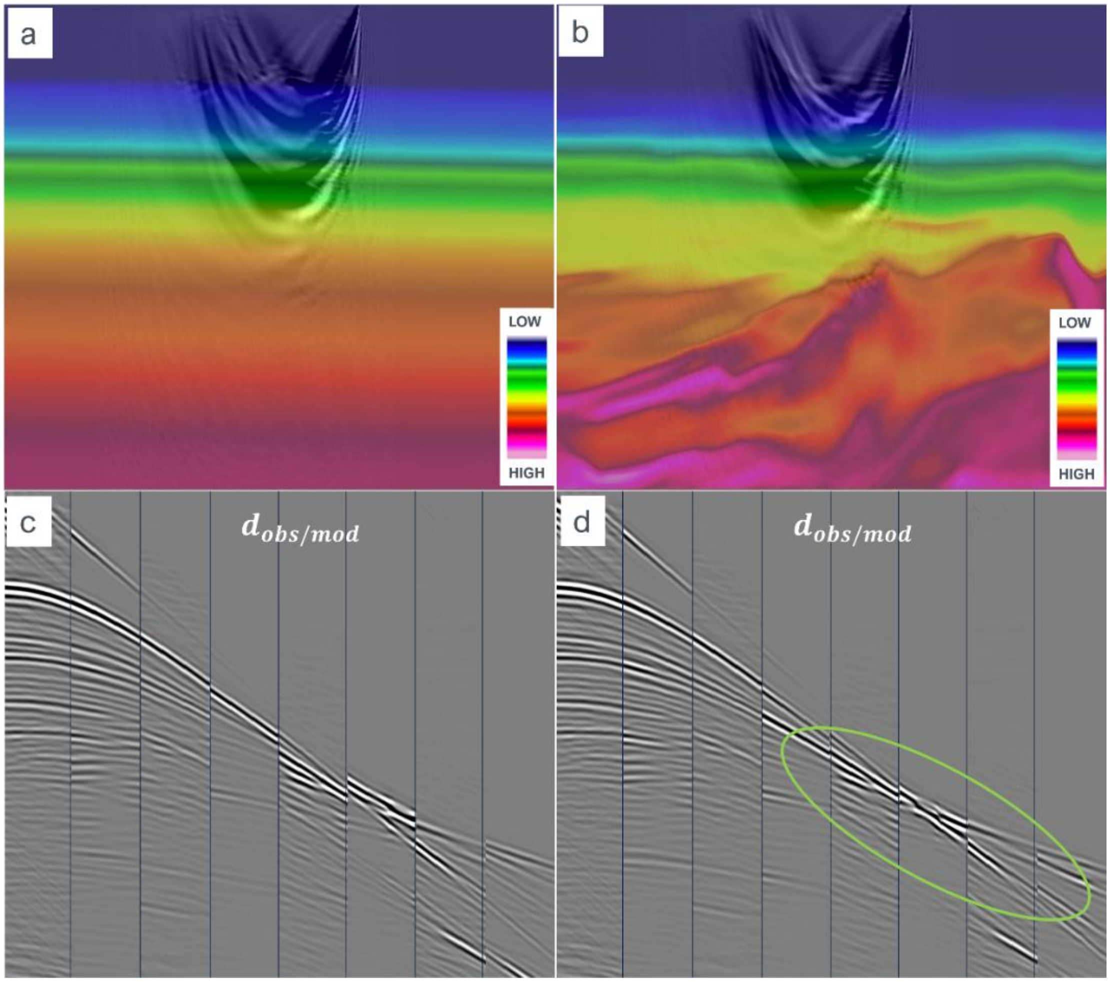

Figure 2 — Single shot FWI kernel overlaid on the initial velocity model a) and after two iterations of FNO-based network prediction b). The corresponding interleaved observed and modeled data are shown in c) and d) respectively. The forward modeling was done with 9 Hz maximum frequency to show more of the complexity in the data. Notice the significant improvements in the data-alignment after the FNO-network predictions (green ellipse) (Korsmo et al., 2025).

Figures 2b and 2d show the corresponding results after two iterations of predictions generated by the FNO-based network. The resulting updated velocity model is geologically plausible and produces a marked improvement in data alignment, particularly for far-offset refractions (highlighted by the yellow ellipse). Forward modeling was performed at 9 Hz, which is higher than typically (<5Hz) used for an initial FWI stage; this choice was made to better illustrate the complexity of the underlying velocity structure in this dataset

In the second example, the FNO-based network predictions are used to replace conventional reflection tomography for the deepest velocity model–building interval. The shallow and intermediate depth sections had previously undergone extensive FWI and tomographic updates and were therefore considered adequately resolved. The dataset consists of a multi-component streamer survey acquired in the Norwegian Sea, covering an area of approximately 6,000 square km. Figures 3a and 3d show an inline section and a depth slice at 5 km depth for the velocity model prior to application of the FNO-based updates. Figures 3b, 3c, 3e, and 3f present the corresponding predicted velocity models, shown with and without seismic image overlays. The updated velocity model exhibits strong structural consistency with the seismic image, and the magnitude of the predicted velocity changes is substantially larger than what is typically achieved using conventional gradient-based updates, reaching values of up to 1000 m/s. In the deeper sections, the network predicts a series of high-velocity volcanic intrusions and sills, with velocities comparable to, or exceeding, those of salt.

The final example is taken from the Agung area, north of Bali, a geologically complex setting shaped by the subduction of the Indo-Australian Plate. This area lies within the Sunda volcanic arc and is characterized by pronounced tectonic deformation, including compressional thrust faulting and extensive volcanism. The basin is highly deformed and faulted, with carbonate formations constituting the primary reservoir targets. Reliable discrimination between volcanic rocks and carbonates in seismic data is therefore critical for exploration in this region.

In 2023, a newly acquired marine streamer dataset with offsets extending to 10 km provided high-quality seismic data for improved imaging of carbonate reservoirs and for distinguishing them from volcanic units. The input velocity model supplied to the FNO network was partially derived from a regional multi-client velocity model–building workflow, incorporating reflection tomography, low-frequency (8 Hz) FWI, and an initial interpretation of the basement structure.

Figures 4a and 4b present cross-sections of the initial and updated velocity models after two iterations of FNO-based predictions, with the corresponding RTM images overlaid. Equivalent comparisons for a depth slice at 3 km are shown in Figures 4e and 4f. The cumulative velocity updates are displayed in Figures 4c and 4d, while the RTM images alone are shown in Figures 4g and 4h. As highlighted by the yellow arrows, the FNO network introduces several geologically consistent high-velocity geobodies. In the initial model, regions beneath the interpreted volcanic units exhibit significant imaging distortions; following the FNO updates, the basement structure in these areas is noticeably better focused, as indicated by the orange ellipse.

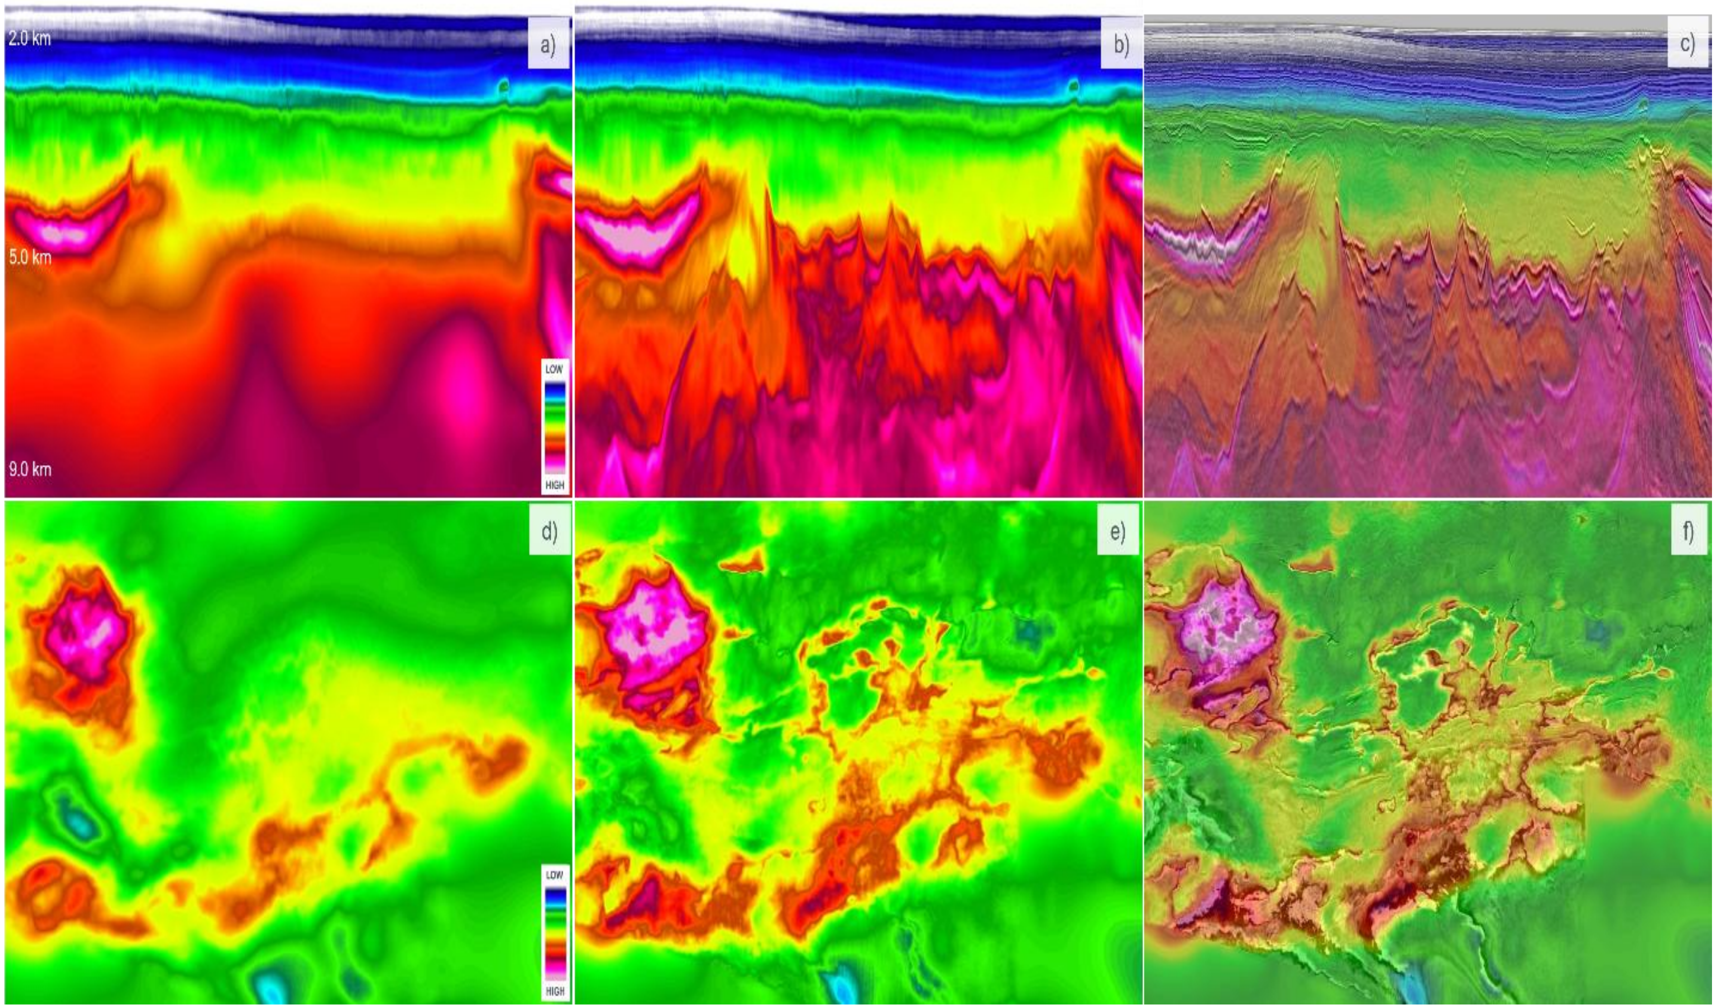

Figure 3 — Inline of the initial a) and FNO-predicted velocity model b) and with the stack overlay in c). Depth slice at 5km of the initial d), FNO-predicted e) and with the stack overlay f). The color legend shows the velocity range from 1500m/s (blue) to 5000m/s (pink). Notice the level of details and the magnitude of the update achieved after a single iteration of prediction and the structural consistency of the results. The FNO-network has detected several volcanic intrusions and sills at different levels in the deep section (Korsmo et al., 2025).

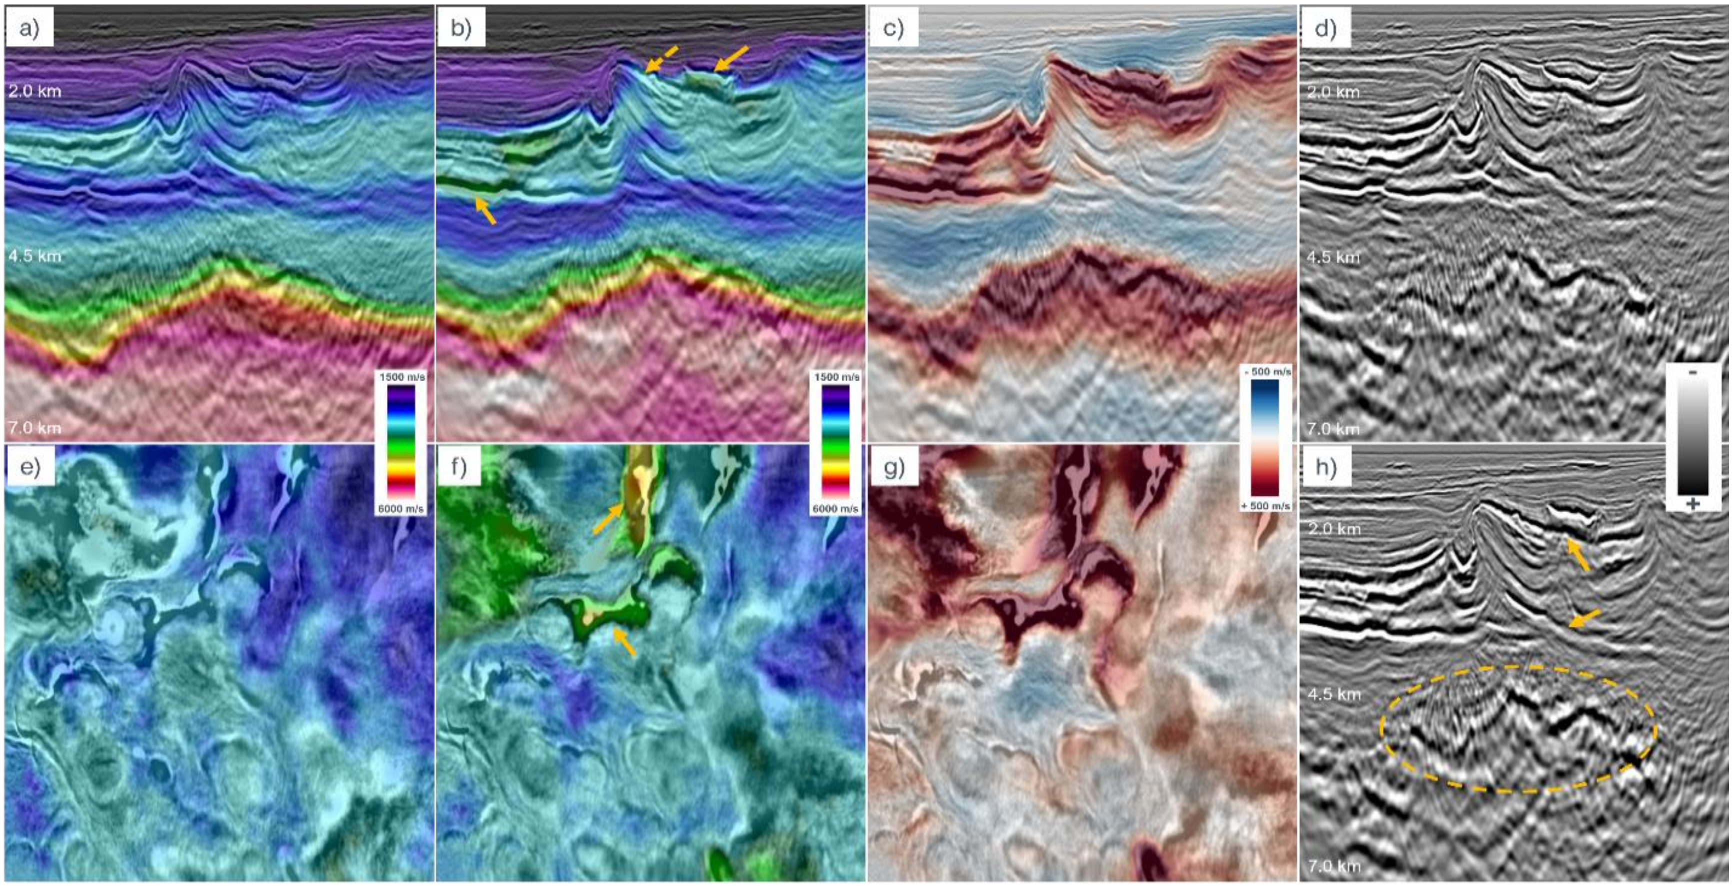

Figure 4 — Cross-section of the initial a) and FNO-predicted model b) with image overlay. Accumulated velocity update after two iterations of FNO-predictions in c) and the corresponding RTM image in d) and h). Figure e), f) and g) shows a depth slice at 3km of the initial, FNO-predicted and the accumulated velocity update. Color legend in figures a), b), e) and f) represents velocity range from 1500 m/s (blue) and 6000m/s (pink), whereas in figure c) and g) it represents accumulated velocity difference from -500m/s (blue) to +500m/s (red). Notice the structural consistency of the predicted model and how it maps out the volcanic structures. The RTM image is significantly improved in the deepest section, revealing more of the basement structure, after solving the shallow complexity with the FNO network predictions (Korsmo et al., 2025).

CONCLUSIONS

In this study, we have demonstrated the application of deep learning to VMB as a means of conditioning the initial model for FWI, thereby reducing cycle-skipping and producing geologically consistent velocity fields. We further show that even a limited number of iterations of FNO-based network predictions can augment conventional tomographic updates and lead to substantial improvements in seismic imaging quality. The proposed approach removes the need for RMO picking, manual interpretation, or specialized treatment of sharp high-velocity contrasts. Central to this capability is the use of FNOs, whose inherently non-local behavior is well suited to image-domain velocity updating.

Despite the encouraging results, it is important that the prediction process not be treated as a black box. Standard VMB quality control procedures should be applied to the outputs, with particular attention when predictions are generated outside the training model distribution. At minimum, QC should include (i) Modelled shot gathers compared to recorded shot gather, (ii) cross correlation map between modelled shot and recorded shot gather, (iii) Imaged CDP gathers to judge flatness, and (iv) gamma map at key horizons. Rigorous quality control of the resulting migrated images remains essential, as does careful evaluation of downstream FWI performance, particularly with respect to convergence speed and stability when initialized from FNO-derived velocity models. Ongoing efforts are focused on improving the realism of the synthetic training datasets used to train the FNO, with the objective of developing a single, generalized model. Achieving this goal requires that the training data adequately represent the range of geological settings in which the network will be deployed. In parallel, we are exploring alternative network architectures, loss functions, and constraint formulations to further enhance robustness and usability for operational seismic imaging workflows.

ACKNOWLEDGEMENTS

We would like to thank TGS Multiclient for the permission to present this work and our colleagues at TGS for their valuable contributions and discussions.

REFERENCES

Brittan, J. and Yuan, J., 2005, Dense multi-offset reflection tomography: SEG Technical Program Expanded Abstracts, 2534-2537.

Crawley, S., Huang, G., Djebbi, R., Ramos-Martinez, J. and Chemingui, N., 2024. Shortening Turnaround Time for High-Resolution Velocity Model Building with Deep Learning: 85th EAGE Annual Conference & Exhibition, European Association of Geoscientists & Engineers. 1-5.

Huang, G., Crawley, S., Djebbi, R., Ramos-Martinez, J., and Chemingui, N.,2023, Deep learning velocity model building using Fourier neural operators: 84th Annual International Conference and Exhibition, EAGE, Extended Abstracts.

Korsmo, O., Crawley, S., and Chemingui, N.,2025, Applications of deep neural networks for velocity model building: 5th International Meeting for Applied Geoscience and Energy. Society of Exploration Geophysicists, pp. 671-675.

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. https://arxiv.org/abs/2010.08895.

Li, Z., Huang, D.Z., Liu, B. and Anandkumar, A., 2023, Fourier neural operator with learned deformations for PDEs in general geometry: *Journal of Machine Learning Research* 24: 1-26

Luo, Z., Brittan, J., Fan, D., Mecham, B., Farmer, P. and Martin, G., 2014, Imaging complexity in the earth — Case studies with optimized ray tomography: *The Leading Edge* 33: 1016–1018, 1020, 1022.

Luo, Z., Fan, D., Farmer, P., Martin, G., 2017, Obtaining geologically conformable tomographic models through anisotropic diffusion preconditioning: *SEG Technical Program Expanded Abstracts.* Society of Exploration Geophysicists. 5961-5965.

Whitmore, N. D., Ramos-Martinez, J., Yang, Y. and Valenciano, A.A., 2020, Seismic modeling with vector reflectivity: *SEG Technical Program Expanded Abstracts.* Society of Exploration Geophysicists. 2709-2713.