MEOS GEO 2025

SPE-226965-MS

Transforming 3D UHR Seismic Data Processing with Machine Learning Solutions

B. Caselitz[1], M. Raafat[2] and J.Brittan[1]

[1]TGS UK
[2]TGS Egypt

Introduction

The field of geophysical exploration, particularly in seismic data acquisition and processing, is undergoing a profound transformation through the integration of machine learning (ML) technologies. While seismic data processing has traditionally relied on computationally intensive and manually tuned methods, machine learning offers the promise of accelerated workflows, improved resolution, and consistency across datasets. This work focuses on the application of ML techniques specifically to 3D ultra-high-resolution seismic (UHRS) data, highlighting the potential of artificial intelligence to both enhance quality and significantly reduce processing *time*.

Although UHRS datasets are generally smaller in size than conventional seismic datasets, the increasing scale of modern surveys, such as those exceeding 500 square kilometers as cited by Caselitz et al. (2024), has made computational efficiency and turnaround time critical considerations. Even for smaller surveys, the typical bottleneck lies in the testing and parameterization phase of processing, which can be dramatically streamlined through ML approaches.

Machine learning's core advantages in seismic data processing include eliminating complex parameterization steps, reducing turnaround time, and delivering consistent results across surveys. Currently, ML applications for UHRS data are primarily focused on denoising and deghosting. However, the potential extends further into velocity model building and even the restructuring of processing workflows into integrated, end-to-end models.

This work investigates several supervised machine learning techniques, including RIDNet deep neural networks, Fourier Neural Operators (FNO), and U-Net Convolutional Neural Networks (CNNs), and applies them to various stages of the UHRS processing sequence. These methods are explored through real-case applications, including receiver deghosting, velocity model building, and integrated pre-stack correction workflows. The broader objective is to demonstrate how ML can fundamentally change UHRS processing by merging formerly discrete steps into cohesive, real-time solutions that are more efficient and equally robust.

Machine Learning Techniques for UHRS Processing

In recent years, machine learning (ML) has become a transformative tool in seismic data processing, enabling more efficient interpretation and analysis of complex subsurface structures. Among the various ML approaches, supervised learning techniques have demonstrated superior performance over unsupervised methods in seismic processing tasks such as denoising, fault detection, and lithofacies classification. The primary advantage of supervised learning lies in its use of labeled data, which allows models to learn direct mappings between seismic attributes and desired outputs, significantly enhancing accuracy and generalization. In contrast, unsupervised techniques, which rely solely on data patterns without labeled outcomes, often struggle to capture the nuanced geophysical relationships inherent in seismic datasets. Supervised models, such as convolutional neural networks (CNNs), have shown remarkable success in reducing noise while preserving geological features, owing to their ability to be trained on synthetic or real seismic datasets. Additionally, supervised learning provides greater control over model behavior, interpretability, and quality assurance. While unsupervised methods remain useful for exploratory analysis and feature extraction, the precision and reliability of supervised approaches make them more suitable for applications in seismic processing.

Supervised learning techniques were employed throughout this study. These approaches rely heavily on access to high-quality training datasets, which are ideally derived from geophysical processing outputs of previously processed data and/or synthetic data. The ML model development workflow begins with supervised learning, where inputs and corresponding ground-truth outputs are used to generate a prediction model that can be applied to new datasets.

Three primary models were explored:

1. **RIDNet Deep Neural Network (DNN):**

   RiDNet is a sophisticated deep convolutional neural network architecture originally developed for the task of real-world image denoising. It addresses the challenge posed by real noisy images, which often contain complex, non-Gaussian noise arising from sensor imperfections, compression artifacts, and various environmental factors. Unlike synthetic noise, which is typically Gaussian and easier to model, real noise exhibits diverse characteristics that complicate traditional denoising approaches.

   The core innovation of RiDNet lies in its use of Residual-in-Residual Dense Blocks (RRDBs), which form the building blocks of the network. Each RRDB contains multiple convolutional layers arranged in a densely connected fashion, meaning each layer receives inputs from all preceding layers within the block. This dense connectivity promotes extensive feature reuse and helps the model learn a rich representation of the input image. Additionally, the RRDB incorporates residual connections at two levels: inside the dense block and wrapping the entire block itself. These nested residual connections facilitate better gradient flow during training, helping to mitigate the vanishing gradient problem common in deep networks. They also enable the model to capture both low-level details, such as edges and textures, and high-level abstract features necessary for accurate denoising.

   A unique aspect of RiDNet is the integration of attention mechanisms that enhance its ability to focus on informative features while suppressing noise. Specifically, the model employs channel attention and spatial attention modules. The channel attention mechanism operates by weighting different feature channels according to their relevance, enabling the network to emphasize important characteristics while ignoring less useful information. Spatial attention, on the other hand, applies a learned mask over spatial locations within the feature maps, highlighting regions of the image that require more focus during denoising. These attention modules are applied after the RRDBs, refining the feature maps to better preserve meaningful image content while attenuating noise patterns.

   The overall architecture of RiDNet follows an encoder-decoder-like design but does not employ downsampling, thereby maintaining the original spatial resolution of the input throughout the network. The input image first passes through an initial convolutional layer that extracts low-level features. This is followed by a series of RRDBs that progressively enhance the feature representations. Afterward, the attention modules refine these features before a final convolutional layer reconstructs the denoised output. A global residual connection adds the input image back to the network's output, a technique known as residual learning. This approach allows the network to focus on learning the noise component directly, which has been shown to improve both training stability and denoising quality.

While RiDNet was designed for natural image denoising, its architectural principles make it highly adaptable to geophysical data processing, especially seismic data denoising. Seismic data often suffer from complex noise such as random noise, coherent ground roll, and acquisition artifacts, which can obscure important subsurface features. By leveraging the residual-in-residual structure, dense connectivity, and attention modules, RiDNet can be adapted to learn both local waveform characteristics and global structural patterns found in seismic volumes. The residual learning strategy is also advantageous for seismic processing, as the goal is typically to separate noise from the true signal.

In summary, RiDNet is a powerful and elegant neural network architecture that combines residual learning, dense connectivity, and attention mechanisms to excel at removing real-world noise from images. Its principles translate well beyond natural images, offering promising avenues for improving the quality of seismic and other geophysical data through deep learning-based denoising.

This model, previously presented by Farmani et al. (2023), excels in denoising and deghosting applications. It is currently in production use across multiple UHRS projects. RIDNet models are trained individually for each survey to handle specific acquisition configurations and noise characteristics.

2. **Fourier Neural Operator (FNO):**

The Fourier Neural Operator (FNO) is a deep learning architecture designed to learn mappings between infinite-dimensional function spaces, making it particularly well-suited for solving partial differential equations (PDEs). Unlike traditional neural networks that operate in the spatial domain, FNOs leverage the Fourier transform to operate in the frequency domain. This allows them to capture global patterns and long-range dependencies more efficiently, which is especially valuable in modeling complex physical systems.

In the context of seismic applications, FNOs offer a powerful framework for modeling wave propagation, seismic inversion, and other forward and inverse problems governed by PDEs. Seismic wavefields are inherently governed by the wave equation, a PDE that describes how seismic energy propagates through the Earth. Traditional numerical solvers for these equations, such as finite-difference or finite-element methods, can be computationally expensive and time-consuming, especially when applied to large-scale 3D models.

FNOs address this challenge by learning a surrogate model that approximates the solution operator of the wave equation. Once trained, an FNO can predict the evolution of seismic wavefields across time and space with significantly reduced computational cost. This makes it highly attractive for tasks like real-time seismic simulation, full waveform inversion (FWI), and data-driven velocity model building.

Moreover, because FNOs operate in the frequency domain, they are naturally aligned with the spectral characteristics of seismic data, which are often analyzed using Fourier-based methods. This spectral perspective allows FNOs to generalize well across different geological settings and input conditions, making them robust tools for seismic interpretation and modeling.

In summary, the Fourier Neural Operator provides a scalable, efficient, and physically informed approach to seismic data processing and modeling, bridging the gap between data-driven learning and physics-based simulation. Its ability to learn complex mappings in high-dimensional spaces with fewer training samples and faster inference times positions it as a transformative tool in modern geophysics.

Introduced by Crawley et al. (2023), FNO is a physics-informed machine learning model that has demonstrated effectiveness in velocity model generation. This study represents the first application of FNO to UHRS data. Although initially trained on conventional seismic datasets, FNO shows promise for UHRS applications, particularly in its ability to replicate the resolution and geological coherence of high-frequency full waveform inversion (FWI) outputs at a fraction of the computational cost.

3. Additionally, U-Net CNNs were used for post-migration noise attenuation and reflection-diffraction separation. These secondary applications support hazard identification (e.g., detection of boulders or shallow objects) and enhance the interpretability of processed seismic images. The Convolutional Neural Network U-Net (CNN U-Net) is a specialized deep learning architecture originally developed for biomedical image segmentation, but it has since been widely adopted in geophysics, particularly for seismic data interpretation. Its strength lies in its encoder-decoder structure with skip connections, which enables it to capture both global context and fine-grained spatial details, an essential capability for processing complex seismic images.

   The encoder part of the U-Net progressively reduces the spatial dimensions of the input seismic section while increasing the depth of feature maps, allowing the network to learn abstract, high-level representations of geological structures. The decoder then upsamples these features back to the original resolution, reconstructing detailed spatial information. Crucially, the skip connections between corresponding layers in the encoder and decoder preserve fine-scale features that might otherwise be lost during downsampling, ensuring that subtle geological boundaries are accurately captured.

   One of the key advantages of CNN U-Net in seismic processing is its ability to learn directly from labeled seismic images, making it a powerful supervised learning tool. Once trained, the model can generalize to unseen data, enabling rapid and automated interpretation across large seismic volumes. This significantly reduces the manual effort required by geophysicists and enhances consistency in interpretation.

   In summary, CNN U-Net provides a robust and efficient framework for seismic image segmentation and interpretation. Its ability to combine deep feature extraction with precise localization makes it an indispensable tool in modern seismic workflows.

Application Examples

1. Receiver Deghosting

Deghosting of UHRS data is complex due to the variability in receiver and source depths. Traditional geophysical deghosting methods, like those described by Bekara et al. (2024), rely on inversion-based techniques that are accurate but computationally expensive.

In this study, RIDNet was trained using input-output pairs generated by conventional inversion-based deghosting of a single acquisition sequence. Once trained, the RIDNet model was applied to other sequences in the dataset with remarkable speed and fidelity. A comparison of results on a QC line (excluded from training data) shows that the machine learning method produced results visually and structurally equivalent to geophysical deghosting, with dramatically reduced processing times (Figure 1).

This not only validates RIDNet's ability to generalize but also underscores ML's value in accelerating early-phase processing without sacrificing output quality.

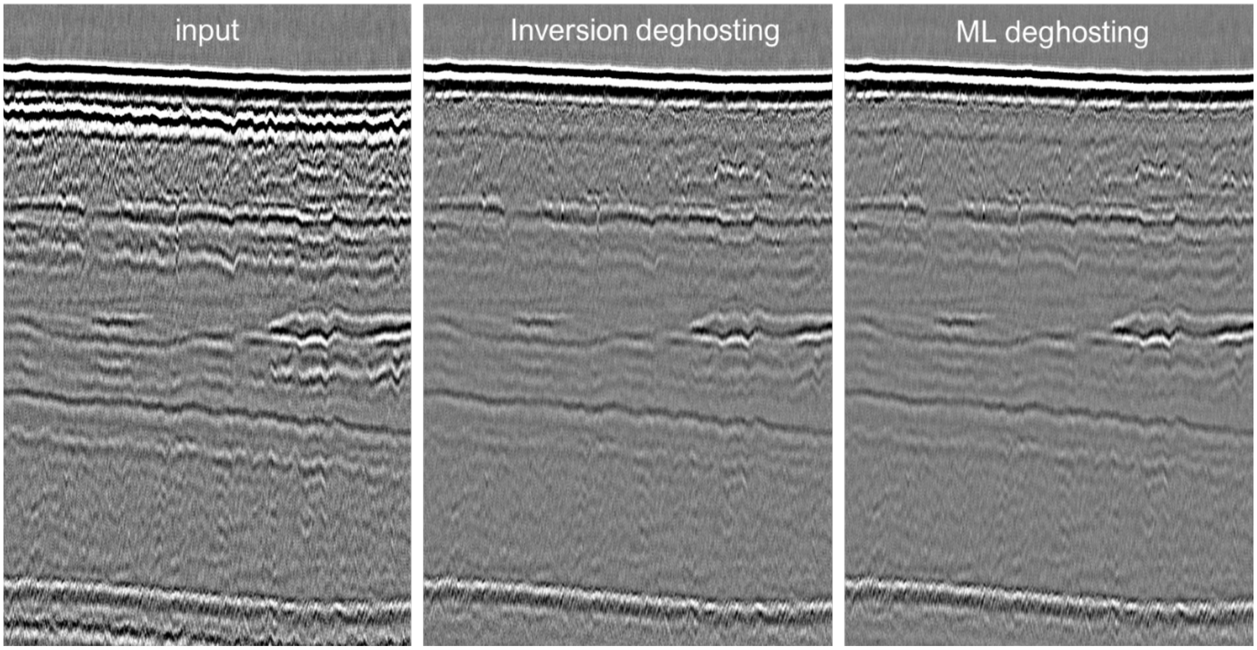

Figure 1: 2d stack comparison QC. Input to receiver deghosting (left), inversion deghosting (center), ML deghosting (right). Both methods produce a high-quality output.

2. Velocity Model Building with FNO

Velocity model building is central to seismic imaging. Traditional techniques such as tomographic inversion and full waveform inversion (FWI) are computationally demanding and require extensive manual QC. This challenge is particularly acute in the shallow subsurface, where UHRS data aims to resolve fine-scale sedimentary features.

The FNO model, trained entirely on synthetic data from conventional surveys (with 15 Hz RTM angle gathers), was applied to UHRS datasets. Despite not being trained on UHRS-specific geometries or frequency content, the resulting ML-based velocity model demonstrated excellent geological conformity. When overlaid on migrated stacks, the FNO-derived velocity model highlighted shallow channels filled with low-velocity sediments more clearly than the initial model (Figure 2).

This indicates that FNO models can deliver near-FWI resolution at a fraction of the computational cost, providing a viable alternative for high-resolution velocity modeling in production environments.

3. Integrated Post-Stack ML Workflow

One of the most innovative contributions of this study is the integration of multiple pre-stack corrections into a single ML model. Traditionally, UHRS processing includes discrete steps for statics correction, NMO correction, deghosting, denoising, and multiple suppression, often requiring iterative loops and intermediate QC.

Using RIDNet, a neural network was trained to perform all of these steps in one pass. A proof-of-concept application compared a conventional pre-stack processed stack against a raw stack that was processed post-stack using the ML model (Figure 3). The ML-based stack exhibited better signal-to-noise characteristics than the unprocessed raw stack and retained many of the qualities of the fully processed version.

Though some residual multiple energy remained, the integrated ML approach points toward a new paradigm in UHRS data processing, one that favors speed and integration over modular, sequential workflows. This capability could support the creation of ultra-fast-track seismic products within hours rather than days.

Benefits and Challenges of ML in UHRS

Machine learning is rapidly transforming seismic data processing by introducing new levels of speed, precision, and automation. One of the most significant advantages of ML is its ability to drastically reduce processing times, particularly in stages that are traditionally labor-intensive and quality-control-heavy, such as deghosting and velocity model building. By automating these steps, ML not only accelerates workflows but also frees up geophysicists to focus on higher-level analysis and decision-making.

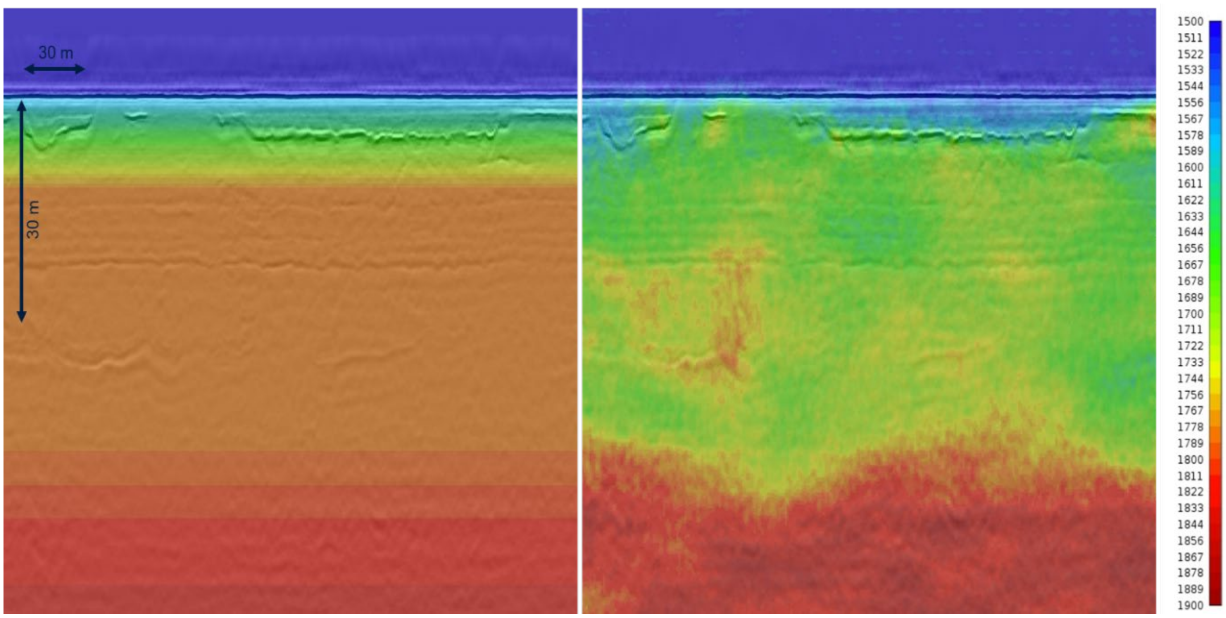

Figure 2: Migrated stack with overlaid velocity model (m/s) – Initial model and stack (left), ML model and stack (right). The ML algorithm is able to capture low velocity related to shallow channels just beneath the seabed.

Another key benefit is the resolution that ML models can achieve. Advanced architectures like the Fourier Neural Operator (FNO) have demonstrated the ability to produce results that approach the quality of full waveform inversion (FWI), especially in resolving shallow subsurface features. This level of detail is critical for accurate imaging and interpretation in complex geological settings.

ML also enhances consistency across datasets. Unlike geophysical parameter testing, which can vary between individuals or teams, ML models provide repeatable outputs that are less susceptible to human bias.

Furthermore, ML enables the simplification of seismic workflows. Traditional processing sequences often consist of discrete, sequential steps, each requiring separate tools and expertise. ML models can integrate multiple steps into a unified process, streamlining operations and reducing the potential for error.

Once trained, ML models are also highly efficient in terms of resource usage. They can be deployed across multiple datasets with minimal computational overhead, making them cost-effective for repeated application in similar geological settings.

However, the adoption of ML in seismic processing is not without its challenges. A primary limitation is the dependence on high-quality training data. ML models require labeled datasets derived from previous geophysical outputs, which are often proprietary. In ultra-high-resolution seismic surveys, for example, data cannot typically be reused across projects, necessitating retraining for each new survey.

Generalizability is another concern. While models like FNO show promise in adapting across domains, ML models tend to perform best when tailored to the specific acquisition geometry and geological context of a given survey. This limits their out-of-the-box applicability and requires careful tuning for each use case.

Interpretability also remains a challenge. Deep learning models, including those used in seismic processing, often function as "black boxes," making it difficult to understand the rationale behind their predictions. This opacity necessitates additional quality control measures to ensure reliability and trustworthiness.

In summary, while machine learning offers transformative benefits for seismic processing, ranging from speed and resolution to workflow integration, it also introduces new challenges related to data availability, model generalization, and interpretability. Addressing these challenges is essential for fully realizing the potential of ML in geophysical applications.

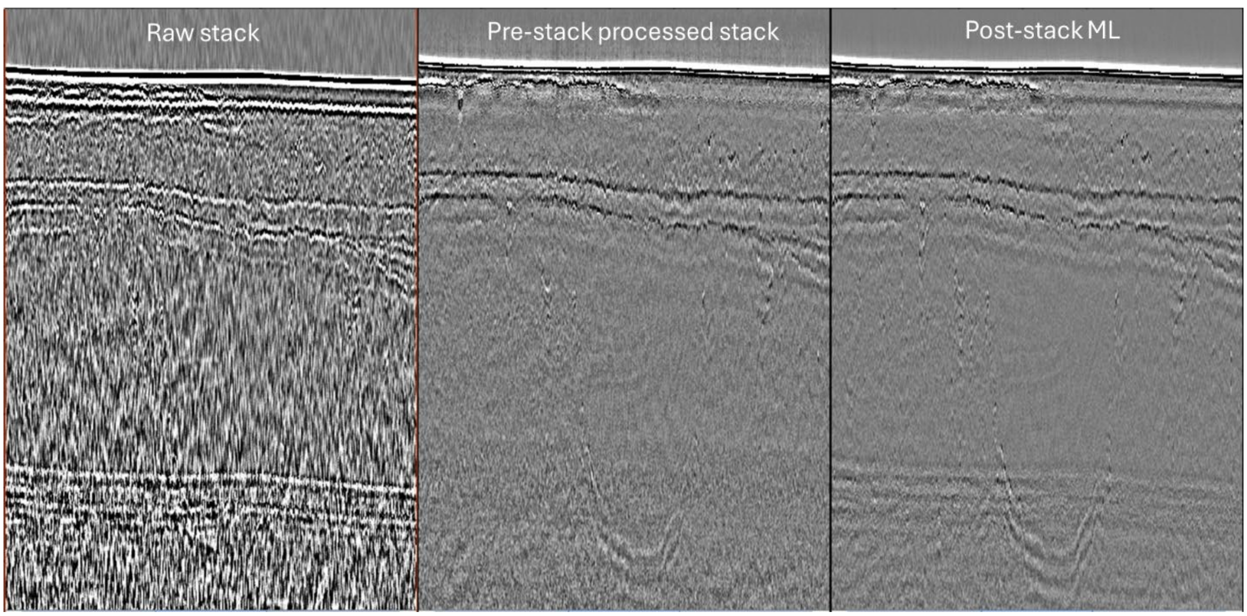

Figure 3: Unmigrated stack comparison QC. Raw stack (left), pre-stack processed stack (center), post-stack ML (right). The ML solution provides a better signal-to-noise ratio but leaves more multiples energy.

Conclusion

Machine learning stands as a transformative technology in the processing of 3D ultra-high-resolution seismic data. From denoising and deghosting to velocity model building and multi-step integration, ML not only accelerates traditional workflows but also redefines what is possible in terms of resolution, efficiency, and scalability.

While certain barriers, especially data privacy and model generalizability, remain, the demonstrated potential of supervised learning techniques such as RIDNet and FNO suggests that ML will increasingly become central to the seismic processing toolkit. The key lies in building high-quality training datasets, adapting models to specific geological and acquisition conditions, and rethinking workflows to fully exploit the integrative capabilities of artificial intelligence.

As the energy sector continues its digital transformation, the insights gained from studies like this will pave the way for smarter, faster, and more accurate subsurface imaging providing vital data for safer operations, better exploration, and ultimately more efficient resource development.

Acknowledgements

We thank TGS management for their support in publishing and our TGS colleagues who have been involved in the 3D UHRS projects.

References

Bekara M., Davison C., Lange M., Limonta L. (2024). Deghosting of UHR seismic data via dynamic ghost tracking, EAGE 2024.

Caselitz B., Limonta L., Oukili J., Tegnander J., Catterall V. (2023). 3D Ultra High resolution seismic processing – A case study from Offshore USA. IMAGE 2024.

Crawley S., Huang G., Djebbi R., Ramos J., Chemingui N. (2023). High resolution angle gather tomography with Fourier neural operators. IMAGE 2023.

Farmani B., Lesnes M., Pal Y. (2023). Multisensor noise attenuation with RIDNet. EAGE 2023.