

Implicit FWI using sinusoidal representation networks: application to marine field data

Introduction

Among the most advanced techniques for high-resolution subsurface characterization, Full Waveform Inversion (FWI) stands out for its ability to extract quantitative information from seismic recordings. The method casts the imaging problem as an iterative minimization of the discrepancy between synthetic and recorded waveforms, yielding detailed estimates of subsurface properties. Its strength lies in leveraging the complete information encoded in seismic traces to achieve high resolution imaging. Yet, the same richness that makes FWI attractive also underlies its well-known limitations: the inverse problem is inherently non-linear and ill-conditioned, making it prone to cycle skipping and highly sensitive to the choice of starting model. Convergence to spurious local minima of the cost function remains a persistent risk unless the initial velocity model is already a reasonable approximation of the true subsurface. Various strategies have been developed to address these issues, ranging from the design of more convex misfit functionals and the enrichment of low-frequency components to global search algorithms and Bayesian formulations that improve exploration of the model space.

Here we investigate the use of Implicit Neural Representations (INRs) as an alternative parameterization for FWI. Rather than discretizing the velocity field on a computational grid, we encode it as a continuous mapping from spatial coordinates to velocity values, learned by a compact multilayer perceptron (MLP) whose weights constitute the inversion unknowns. Building on the findings of Sun et al. (2023), who showed that this reparameterization can substantially reduce sensitivity to starting model enabling convergence even from random weight initialization, we explore INRs to regularize the inversion implicitly. This behavior is linked to the spectral bias inherent in gradient-based deep-learning optimization, which naturally favors the recovery of smooth, low-frequency structures before progressively fitting finer details. At the same time, the continuous nature of the representation allows fine-scale features to be retained without the memory overhead of dense grid discretizations.

While implicit FWI (IFWI) is gaining attention in recent years, successful applications to relatively large-scale field data remain limited. We test IFWI on field data from the Viking Graben in the North Sea. Seismic wave propagation is modelled under the acoustic approximation.

Method

Unlike purely data-driven approaches, which are often unable to generalize beyond the examples seen during training, FWI is rooted in the physics of wave propagation: the forward problem is governed by well-established physical equations, which guarantee predictive consistency and make the results physically interpretable. This naturally motivates combining deep learning with the physical framework of FWI, as doing so has the potential to yield inversion schemes that are both more robust and generalizable (Sun et al., 2023). That said, most deep-learning-based FWI methods still inherit the classical pitfalls of the technique, particularly the sensitivity to the choice of starting model.

In the field of shape and scene representation, coordinate-based neural networks proved to learn continuous functions mapping spatial coordinates to physical or visual quantities of interest (Sitzmann et al., 2020). Once trained, these networks serve as implicit, resolution-independent descriptions of the objects or fields they encode and can be queried at arbitrary locations to reconstruct images or 3D surfaces with high fidelity. As highlighted by Sun et al. (2023), the same principle applies to subsurface parametrization: rather than storing model information on a fixed computational grid, an INR encodes it as a continuous function within the network weights. The resulting representation is both memory-efficient and capable of capturing fine-scale features regardless of the chosen spatial discretization.

In seismic inversion, INR represents the subsurface parameter field $\mathbf{m}(\mathbf{x})$ within the constraints to satisfy the wave equation. This leads to IFWI, where the forward operator $\mathcal{F}$ satisfies:

$$\mathbf{R}\mathcal{F}(\mathbf{m}, s, \mathbf{x}, t) - \mathbf{d}_{\text{obs}} = 0$$

where $\mathbf{R}$ is the matrix of receiver layouts, $\mathbf{d}_{\text{obs}}$ the observed data, s the seismic source, $\mathbf{x}$ and t are spatial coordinates and time, respectively. Substituting the neural network $\mathcal{N}_{\Theta}(\mathbf{x})$ for the model parameters $\mathbf{m}$, the optimization problem becomes:

$$\underset{\Theta}{\operatorname{argmin}} \|\mathbf{R}\mathcal{F}(\mathcal{N}_{\Theta}(\mathbf{x}), s, \mathbf{x}, t) - \mathbf{d}_{\text{obs}}\|^2$$

so that the inversion reduces to updating the network weights Θ rather than a grid of velocity values.

The network corresponds to an MLP, in which information propagates through a sequence of fully connected layers. At each layer l , the output of neuron i is a weighted sum of its input followed by a nonlinear activation f :

$$a_i^{(l)} = f\left(W_i^{(l)} a_i^{(l-1)} + b_i^{(l)}\right)$$

To effectively represent the high-frequency spatial variability characteristic of realistic velocity fields, we adopt the Sinusoidal Representation Network (SIREN) architecture (Sitzmann et al., 2020). The previous equation becomes:

$$a_i^{(l)} = \sin\left(\omega_0 W_i^{(l)} a_i^{(l-1)} + b_i^{(l)}\right)$$

The parameter ω_0 governs the range of spatial frequencies the network can express and is treated as a tunable hyperparameter. Compared to conventional activations, sinusoidal units provide accurate representations of high-order derivatives.

Rather than pre-training the network to match a reference starting model, we opt for random weight initialization following the scheme of Sitzmann et al. (2020) and Sun et al. (2023). Weights at each layer are sampled from a uniform distribution:

$$W_i^{(l)} \sim \mathcal{U}(-\sqrt{6/n}, \sqrt{6/n})$$

where n is the number of input features fed into the layer. This choice is designed to ensure that the inputs to each sinusoidal unit follow approximately a standard normal distribution.

Results

To test the IFWI approach, we consider a 2D seismic line from the Viking Graben (North Sea). The dataset, employed by several studies, is publicly available in the SEG database (https://wiki.seg.org/wiki/Mobil_AVO_viking_graben_line_12). For further information see Madiba and McMechan (2003). From the original 1001 source positions, we selected a subset of 50 shots, each containing 120 channels acquired with a typical marine off-end layout. We used the first 3.2 s of the available 6 s recordings, sampled at 4 ms intervals. The source and receiver spacings are 500 m and 25 m, respectively, with minimum and maximum offsets of 262 m and 3237 m. The main challenges posed by this dataset are the relatively short maximum offset and shallow water depth (350 m), which result in strong water-bottom multiples dominating the records. Two well logs ("Well 4" and "Well 5") are available and used to validate the results, as well as to derive a prior model to regularize the inversion.

For forward modelling, the domain is discretized on a grid of $n_z = 500$ and $n_x = 1500$ nodes, with spacings of 7.5 m in depth and 20 m in the horizontal direction, yielding a model 3.75 km deep and 30 km wide. Wavefield propagation is simulated using Deepwave (Richardson, 2022), a finite-difference solver that leverages automatic differentiation for gradient computation. The source wavelet is estimated by stacking short-offset traces after flattening the first arrivals. Data pre-processing is kept minimal, consisting of a 8 Hz low-pass filter followed by trace-by-trace normalization.

The IFWI network is an MLP with two layers of 64 neurons each and sinusoidal activation frequency $\omega_0 = 15$. We run 700 iterations with the Adam optimizer, completing in approximately one hour on a

server equipped with an NVIDIA RTX 6000 GPU. A notable advantage of the INR parametrization is the dramatic reduction in inversion unknowns: a standard grid-based discretization on the same mesh employed for the forward modelling would require $n_z \cdot n_x = 7.5 \cdot 10^5$ parameters, whereas the implicit representation reduces this to just 8577 weights and biases.

The inversion is initialized from a random velocity field (Figure 1(a)). To improve convergence, we include a model regularization term in the loss function, which is now expressed as:

$$\mathcal{L}(\Theta) = \|\mathcal{F}(\mathcal{N}_\Theta(\mathbf{x})) - \mathbf{d}_{\text{obs}}\|^2 + \lambda \|\mathcal{N}_\Theta(\mathbf{x}) - \mathbf{m}_{\text{prior}}\|^2$$

The prior model $\mathbf{m}_{\text{prior}}$ is obtained by interpolating V_p from the two well logs across the full lateral extent of the domain (Figure 1(b)). This prior captures fine layering with localized velocity inversions, as well as a lateral variation below 2 km depth, with higher velocities on the left side of the model. The regularization helps the IFWI to preserve reasonable velocity values in poorly illuminated regions, while the scaling factor λ ensures that the optimization remains primarily driven by the data fitting. At the end of the inversion, $\mathcal{L}(\Theta)$ reduces to about 8.5% of its initial value. The final inverted model (Figure 1(c)) is consistent with independent studies such as Taufik and Alkhalifah (2026). Figure 1(d) compares the well log curves (blue) with velocity profiles from the initial (orange) and predicted (green) models at the corresponding lateral positions, showing that IFWI successfully recovers the general V_p trend at both well locations despite starting from a random model.




Figure 1 The random starting model (a); the prior model used to regularize the IFWI (b); the predicted model (c), with the locations of the two wells; comparison of the available well logs with the corresponding columns in the starting and predicted models (d).

Data fit is examined in Figure 2 for the shot at the second source position. Panels (a–c) show the observed, initial and predicted data, respectively, while panels (d) and (e) display the sample-by-sample residuals between observed and initial, and observed and predicted gathers. The large initial misfits are substantially reduced after inversion: some amplitude discrepancies remain, but cycle skipping, prominent in the initial data, is effectively mitigated. This is further confirmed by the single-trace comparisons at far and mid offset in Figures 2(f) and 2(g), where the predicted traces (green) agree well with the observed ones (blue): at far offset the main phases are correctly recovered, while at mid offset only minor phase and amplitude differences persist.

Conclusions

We have presented an application of acoustic IFWI, based on a SIREN network, to marine field data. Combining FWI with INR allows the inversion to benefit from the advantages of machine-learning-based optimization, most notably the implicit regularization arising from the frequency bias property of gradient descent. The use of sinusoidal activation functions further enhances the approach by enabling accurate representation of fine-scale features, contributing to the mitigation of common FWI issues




Figure 2 The observed (a), initial (b) and predicted data (c); data residual between observed and initial (d); data residual between observed and predicted (e); comparison of two selected traces (f, g).

such as cycle skipping and strong dependence on the starting model. In the test case considered, IFWI converges from a random initial model to a velocity field consistent with independent studies and supported by well log comparisons and the quality of the data fit achieved after inversion.

Further improvements of these results could be obtained by incorporating a larger number of shots and extending the frequency bandwidth to recover finer structural details. Promising directions for future work include strategies to accelerate convergence and the integration of probabilistic frameworks to provide uncertainty quantification.

Acknowledgements

The authors thank Nicola Bienati, Matteo Caporal and Lorenzo Russo from ENI for the insightful discussions about the IFWI and its application to the Viking Graben dataset.

References

- Madiba G.B. and McMechan G.A. [2003] Processing, inversion and interpretation of a 2D seismic dataset from the North Viking Graben, North Sea. *Geophysics*, 68(3), 837-848.
- Richardson A. [2022] Deepwave. <https://ausargeo.com/deepwave/>.
- Sitzmann V., Martel J., Bergman A., Lindell D. and Wetzstein G. [2020] Implicit neural representations with periodic activation functions. *Advances in Neural Information Processing Systems*, 33.
- Sun J., Innanen K., Zjang T. and Trad D. [2023] Implicit seismic full waveform inversion with deep neural representation. *Journal of Geophysical Research: Solid Earth*, 128.
- Taufik M.H. and Alkhalifah T. [2026] Accelerating Bayesian full waveform inversion using reconstruction-guided diffusion sampling. *Geophysical Journal International*, 245(2).