Unsupervised Attenuation of Aliased Ground Roll in Undersampled 3-D Land Seismic Data

Ji Li, Dawei Liu, and Daniel Trad

***Abstract*—Random spatial undersampling and irregular acquisition geometries in land seismic surveys often produce strong, aliased ground roll whose low-velocity energy overlaps with reflections in both the time–space and frequency–wavenumber domains. Aliasing distorts the noise's apparent moveout and frequency content, making conventional FK and τ-p filters ineffective for ground-roll suppression without damaging the signal. Irregular sampling prevents the direct application of many conventional filtering techniques. Even methods that can operate on irregular grids, such as generalized Fourier transforms, fail to recover the missing spectral components, making coherent noise suppression with high fidelity particularly challenging. An unsupervised 3-D ground-roll attenuation method based on implicit neural representations (INRs) is introduced to address these limitations. The seismic wavefield is modeled as a continuous function of receiver coordinates, source coordinates, and time, enabling direct denoising from the observed measurements without requiring external training labels. Owing to the inductive bias of INRs, smooth, coherent components are reconstructed preferentially before sharper or less correlated features, which allows effective separation of reflections from ground roll. To reinforce this separation, normal moveout correction is applied before network training, and a horizontal-gradient penalty is included in the loss function to suppress residual crossline variations associated with noise. This regularization improves optimization stability and eliminates the need for manual early stopping. Validation on synthetic and field 3-D land datasets, including cases with 50% and 75% random trace decimation, demonstrates substantial ground-roll attenuation while preserving reflection continuity and amplitude, confirming the method's effectiveness under challenging acquisition scenarios.**

***Index Terms*—Ground roll, implicit neural representation (INR), seismic denoising, unsupervised deep learning.**

I. Introduction

ECONOMIC and logistical constraints in land seismic acquisition often require compromises in spatial sampling density [1], [2], [3]. In 3-D surveys, these constraints frequently lead to irregular or sparse sampling patterns in both the receiver and source domains. While such designs reduce acquisition costs and cycle time, they introduce a significant side effect: coherent low-velocity noise, particularly ground roll, becomes aliased in space [4], [5]. Aliasing folds high-wavenumber energy into lower wavenumber regions, causing ground roll to overlap with the spectral content of desired reflections in both inline and crossline directions [6], [7]. This overlap is more severe in multishot 3-D datasets than in single-shot 2-D profiles because the noise characteristics vary from shot to shot, and aliasing can occur along multiple spatial axes [3], [8], [9]. Once ground-roll energy is aliased, conventional separation criteria based on apparent velocity, slowness, or spatial frequency become unreliable [4], [10], [11].

Ground roll in land seismic data is dispersive [12] and of high amplitude and has a low apparent velocity [13], [14]. It dominates near-source traces and, if left untreated, can obscure weak primary reflections essential for accurate subsurface imaging. The problem is compounded in undersampled 3-D data cubes organized along receiver position, source position, and time, where cross-shot variability produces complex noise patterns. Failure to suppress this noise before further processing risks introducing persistent artifacts that degrade imaging and inversion results [7], [15].

Traditional attenuation methods are primarily filter-based, operating in the f–k domain [14], [16] or the τ–p domain [17], [18]. These methods rely on a clear spectral or slowness gap between signal and noise. Spatial aliasing undermines this gap: in the f–k domain, aliased ground roll appears at low wavenumbers indistinguishable from reflections [19], [20]; in the τ–p domain, aliasing distorts the apparent slowness of ground roll, pushing it into the range of primary events [18], [21]. Transform-based approaches, such as wavelet [22], [23], [24] and curvelet transforms [9], [25], seek to exploit sparsity differences between reflections and noise. Aliasing degrades this sparsity through spectral leakage, spreading coherent noise energy across scales and orientations and contaminating reflection coefficients.

Irregular sampling exacerbates spectral leakage, further reducing the effectiveness of sparsity-driven thresholding [26], [27]. Many frequency- and transform-domain approaches used for ground-roll attenuation, such as f–k filtering and most practical implementations of wavelet and curvelet transforms, rely on the fast Fourier transform (FFT) for efficiency and therefore require regularly gridded data. For irregularly sampled acquisitions, these methods cannot be applied directly

Received 1 October 2025; revised 19 March 2026 and 4 June 2026; accepted 3 July 2026. Date of publication 6 July 2026; date of current version 20 July 2026. This work was supported in part by the Key Program of the National Natural Science Foundation of China under Grant 42530802 and in part by Emissions Reduction Alberta through the ACT4-SPARSE Project. The work of Ji Li and Daniel Trad was supported in part by CREWES industrial sponsors and in part by the Natural Sciences and Engineering Research Council of Canada (NSERC) under Grant CRDPJ 543578-19. *(Corresponding author: Dawei Liu.)*

Ji Li and Daniel Trad are with the Department of Earth, Energy, and Environment, University of Calgary, Calgary, AB T2N 1N4, Canada (e-mail: li.ji1@ucalgary.ca; daniel.trad@ucalgary.ca).

Dawei Liu is with the School of Information and Communications Engineering, Xi'an Jiaotong University, Xi'an 710049, China (e-mail: liudawei2015@stu.xjtu.edu.cn).

Digital Object Identifier 10.1109/TGRS.2026.3710605

without first interpolating the data onto a regular grid, a step that may distort both the signal and the noise. Although discrete Fourier transforms (DFTs) [28], [29] and nonuniform FFTs (NFFTs) [30], [31] can process irregularly sampled data without gridding, they do not address the aliasing introduced by undersampling and, for large 3-D datasets, become computationally expensive.

In recent years, deep learning has emerged as a promising alternative for coherent noise attenuation in seismic data [32], [33], [34]. Supervised methods can perform well when extensive clean-noisy training pairs are available, but such data are rarely obtainable in practice, and model performance is ultimately constrained by label quality. Unsupervised learning can bypass the need for labels, but most implementations are based on convolutional neural networks (CNNs) that assume regularly sampled inputs and face scalability challenges for high-dimensional or irregularly sampled seismic datasets. Implicit neural representations (INRs) provide a fundamentally different approach for modeling continuous signals, offering an alternative to conventional grid-based methods [35], [36]. An INR expresses a signal as a continuous coordinate-to-amplitude mapping parameterized by a multilayer perceptron (MLP), enabling direct functional representation of the underlying field [35], [37]. Unlike CNNs, which operate on discrete, regularly sampled grids and have an inherent local receptive field bias [38], [39], INRs can capture both local and global structure and, importantly, operate directly on irregularly sampled coordinates without requiring prior interpolation [37], [40]. A notable property of INRs, observed in both theoretical and empirical studies, is their tendency to first represent patterns that vary slowly with the input coordinates before fitting those with higher coordinate variation [41]. This spectral bias means that, in datasets where low-frequency structures dominate, these components are naturally learned early in the optimization process, while higher frequency details emerge later. This progressive fitting behavior can be exploited in various applications to separate large-scale, slowly varying structures from rapidly varying components without explicitly applying filtering [35], [41].

Beyond computer vision and graphics, the combined properties of INRs make them particularly well-suited for seismic signal processing. Their continuous coordinate-based formulation naturally accommodates irregular acquisition geometries, while their progressive learning dynamics help distinguish large-scale reflection structures from high-frequency noise. As a result, INRs have been successfully applied to several seismic problems, including 5-D irregular interpolation [42], seismic deblending [43], and full waveform inversion [44], [45]. More recently, Li et al. [46] applied INRs to 2-D ground-roll attenuation on regularly sampled shot gathers, demonstrating the potential of coordinate-based representations for coherent noise suppression. However, restricting the model to 2-D information imposes important limitations. In cases where ground roll is extremely strong, reflections may be entirely masked within a single-shot gather, leaving insufficient information for the INRs to recover the underlying signal. By extending the framework to 3-D, the model can jointly exploit information from multiple spatial directions and receiver geometries, allowing reflections to be reconstructed even when they are obscured in individual 2-D gathers. This directional diversity makes 3-D INRs denoising more robust in practice and substantially improves separation performance under challenging ground-roll conditions. Furthermore, our formulation directly handles irregularly sampled 3-D seismic data without requiring prior interpolation, further highlighting the flexibility of INR-based representations for field applications.

Motivated by these considerations, we develop an unsupervised INR-based framework for attenuating aliased ground roll in undersampled 3-D land seismic data. The network operates directly on irregularly sampled spatial and temporal coordinates, learning a continuous representation of the seismic wavefield from the available measurements. This allows coherent noise attenuation to be performed without prior interpolation or resampling onto a regular grid. To enhance reflection preservation, we apply normal moveout (NMO) correction before training to flatten primaries, making them more distinct from ground roll in the coordinate domain. Additionally, we incorporate a horizontal-gradient regularization term into the loss function to penalize steep spatial variations associated with ground roll while preserving flat events. This design improves convergence stability, reduces sensitivity to hyperparameters, and removes the need for early stopping. Tests on synthetic and field datasets, including cases with 50% and 75% random trace decimation, demonstrate that the proposed approach achieves substantial ground-roll attenuation while maintaining reflection continuity and amplitude in both regularly and irregularly sampled scenarios.

II. Theory

A. Review of INR

INRs model a signal as a continuous function

$$f_\theta : \mathbf{c} \mapsto d_\mathbf{c}$$

where $\mathbf{c}$ denotes a spatial–temporal coordinate vector and $d_\mathbf{c}$ is the corresponding signal amplitude. The mapping is parameterized by an MLP whose weights θ are optimized to minimize a distortion measure, such as mean-squared error (mse) or mean absolute error (MAE), over the set of observed coordinates

$$\theta^* = \arg\min_\theta \sum_{\mathbf{c}\in C} \rho\left(f_\theta(\mathbf{c}),\, d_\mathbf{c}\right)$$

where $\rho(\cdot,\cdot)$ is the chosen loss function. Unlike grid-based representations that store a discrete array of samples, INRs store the signal implicitly in the network's weights, resulting in a fixed parameter cost independent of sampling density. This enables evaluation at arbitrary coordinates, making INRs naturally suited to irregularly sampled seismic data without requiring explicit interpolation.

We adopt the SIREN architecture [35], which uses sinusoidal activation functions to enable the representation of fine-scale variations. In SIREN, the input layer applies a frequency scaling factor ω to control the rate at which high coordinate-frequency components (i.e., rapid variation of the target function with respect to the input coordinates, rather than temporal frequency of the seismic waveform) can be represented

$$\mathbf{h}^{(1)} = \sin\left(\omega\,\mathbf{W}^{(1)}\mathbf{c} + \mathbf{b}^{(1)}\right)$$

where $\mathbf{c}$ is the coordinate vector and $\mathbf{W}^{(1)}$ and $\mathbf{b}^{(1)}$ are the weights and biases of the first layer. In this work, we set the frequency scaling parameter to $\omega = 30$, following the recommendation in [35] to enable the network to capture high-frequency details.

For the hidden layers, we use the same sinusoidal activation with the same frequency scaling

$$\mathbf{h}^{(l)} = \sin\left(\omega\,\mathbf{W}^{(l)}\mathbf{h}^{(l-1)} + \mathbf{b}^{(l)}\right), \quad l > 1.$$

Using the same scaling parameter for all layers simplifies the implementation while maintaining the ability of the network to represent both low- and high-coordinate-frequency components.

A distinctive property of INRs is their spectral bias [41], [47], which refers to the order in which features are fit during training. Here, "frequency" refers not to the Fourier frequency content of the seismic waveform, but to the rate of variation of the target function to the input coordinates. We can formalize this idea by considering the coordinate-gradient magnitude

$$\nu_{\text{coord}}(f) \;=\; \|\nabla_{\mathbf{c}} f(\mathbf{c})\|$$

where low values indicate slow variation over coordinates (low coordinate frequency) and high values indicate rapid variation (high coordinate frequency). This notion of frequency is closely related to the low-frequency bias observed in conventional CNN-based seismic networks, in that both favor the early learning of smooth structures [48]. However, in INRs, the frequency is defined by variation over input coordinates, whereas in CNNs, it is typically described in the data domain (e.g., low temporal frequency or absence of high-wavenumber spatial components) [49]. These two perspectives usually align for events like flat reflections, but can differ markedly for phenomena, such as ground roll, which is temporally smooth yet spatially rapidly varying in coordinate space.

To illustrate, consider two events in an x–t seismic section. A flat reflector after NMO correction appears as a horizontal line: moving along x changes the amplitude very little, so its spatial coordinate frequency is low. By contrast, ground roll, despite being low frequency in the Fourier temporal sense, cuts rapidly across traces in x, producing large spatial derivatives. In NMO-corrected seismic data, primary reflections thus have high self-similarity and low coordinate frequency and are fit early in INRs training. Ground roll, with its steep dips, irregular moveout, and partial incoherence, has high coordinate frequency and is therefore learned only in later optimization stages. This explains why INRs naturally reconstruct flattened reflections first and ground roll later. This inherent ordering can be exploited for noise attenuation: by stopping or regularizing optimization before high coordinate-frequency components are fully fit, primary reflections can be preserved while suppressing aliased ground roll. However, excessive bias toward low coordinate-frequency features can also delay recovery of fine-scale reflection details. The tunable ω parameter in SIREN allows explicit control over this tradeoff, enabling us to balance noise suppression with the preservation of high-resolution seismic features.

B. Horizontal-Gradient Regularization

Regularization terms are widely used in inverse problems to impose prior knowledge and stabilize optimization. In signal processing, total variation (TV) regularization is a common choice for promoting piecewise-smooth solutions while preserving edges [50]. The isotropic TV penalty in continuous form is given by

$$\mathcal{R}_{\text{TV}}\left[f_\theta\right] = \int_{\mathbf{c}\in\Omega} \left\|\nabla f_\theta(\mathbf{c})\right\|_2 d\mathbf{c} \tag{1}$$

where $\mathbf{c}$ denotes the coordinate vector, ∇ is the full gradient to all coordinate axes, and Ω is the coordinate domain. In discrete form, with $\mathcal{C}$ the set of sampled coordinates, this becomes

$$\mathcal{R}_{\text{TV}}\left[f_\theta\right] = \sum_{\mathbf{c}\in\mathcal{C}} \sqrt{\sum\nolimits_{k=1}^{D} \left(\Delta_k f_\theta(\mathbf{c})\right)^2} \tag{2}$$

where D is the number of coordinate dimensions and Δ_k is the finite difference along axis k.

Although isotropic TV preserves edges, it has two drawbacks for seismic processing.

1) It penalizes variations equally in all directions. In seismic data, however, useful reflections often have a dominant orientation (e.g., nearly horizontal after NMO correction), while noise, such as ground roll or migration artifacts, may be oriented differently (e.g., steeply dipping). Applying the same penalty to all directions can therefore blur signal events and fail to exploit these directional differences for noise suppression.
2) The ℓ_1-like penalty in (2) promotes sparsity in the gradient, which can lead to "staircasing" artifacts, i.e., piecewise-constant reconstruction of smoothly varying geological features.

TV can be generalized by weighting the gradient along different coordinate axes, leading to a family of orientation-aware penalties [51], [52]. In this form, the regularization becomes

$$\mathcal{R}_{\text{dir}}\left[f_\theta\right] = \sum_{\mathbf{c}\in\mathcal{C}} \sqrt{\sum\nolimits_{k=1}^{D} w_k^2 \left(\Delta_k f_\theta(\mathbf{c})\right)^2} \tag{3}$$

where w_k is the axis-specific weights. Setting $w_k = 0$ for certain axes effectively removes regularization along those directions. This provides a natural way to construct problem-specific penalties, for example, suppressing only spatial variations while leaving the temporal dimension unconstrained.

In our application, after NMO correction, primary reflections are nearly horizontal in the spatial domain and exhibit strong lateral self-similarity, whereas aliased ground roll remains steeply dipping and rapidly varying in space. This motivates penalizing only lateral variations, leaving vertical (temporal) variations unaffected.

Let the coordinate vector be

$$\mathbf{c} = \left(t, s_x, s_y, r_x, r_y\right)$$

where t is the time, (s_x, s_y) are source coordinates, and (r_x, r_y) are receiver coordinates. We define the horizontal-gradient operator

$$\nabla_{sr} = \left(\frac{\partial}{\partial s_x}, \frac{\partial}{\partial s_y}, \frac{\partial}{\partial r_x}, \frac{\partial}{\partial r_y}\right)^{\top}.$$

The continuous horizontal penalty is

$$\mathcal{R}_H\left[f_\theta\right] = \int_{\mathbf{c}\in\Omega} \left\|\nabla_{sr} f_\theta(\mathbf{c})\right\|_2^2 d\mathbf{c} \tag{4}$$

and its discrete form is

$$\mathcal{R}_H\left[f_\theta\right] = \sum_{\mathbf{c}\in\mathcal{C}} \sum_{k\in\{s_x,s_y,r_x,r_y\}} \left(\Delta_k f_\theta(\mathbf{c})\right)^2. \tag{5}$$

Note the use of ℓ_2 in (5) instead of ℓ_1. This promotes smooth transitions rather than sparse jumps, mitigating staircasing and preserving reflection continuity.

The physical location (x, y) of a midpoint is given by

$$x = \frac{s_x + r_x}{2}, \quad y = \frac{s_y + r_y}{2}.$$

Gradients with respect to (s_x, s_y, r_x, r_y) thus directly control lateral smoothness in the physical (x, y) plane. Enforcing smoothness in source–receiver space therefore enforces smoothness in the physical plane, even for irregularly sampled acquisition geometries. This allows the proposed method to apply the regularization directly in the native source-receiver coordinates, without transforming the data into $(\mathrm{cmp}_x, \mathrm{cmp}_y, \mathrm{offset}, \mathrm{azimuth})$ space and resampling onto a regular grid. By avoiding both the transformation and resampling steps, our approach preserves the original acquisition geometry and ensures that the regularization acts consistently across both dense and decimated acquisition scenarios.

The horizontal penalty is incorporated into the INR's loss

$$\mathcal{L}(\theta) = \sum_{c\in\mathcal{C}} \left\|f_\theta(c) - d_c\right\|_1 + \mu \sum_{c\in\mathcal{C}} \left\|\nabla_h f_\theta(c)\right\|_2^2 \tag{6}$$

where μ controls the strength of the horizontal smoothness regularization, this combination promotes early learning of horizontally aligned primaries while suppressing steeply dipping ground-roll energy. We adopt the ℓ_1 norm for the data-fidelity term rather than the conventional ℓ_2 norm. The ℓ_1 loss grows linearly with the magnitude of residuals, making it less sensitive to large deviations than the quadratic growth of the ℓ_2 loss. This property provides robustness against outliers, such as aliased ground roll bursts or erratic spikes, which could otherwise dominate the optimization if squared in an ℓ_2 formulation. By reducing the influence of these high-amplitude anomalies, the network can prioritize fitting the dominant reflection energy. In addition, the ℓ_1 norm implicitly encourages sparsity in the residuals, aligning with the expectation that, after effective ground-roll attenuation, the majority of coordinates will have small prediction errors, with only localized regions exhibiting significant mismatch.

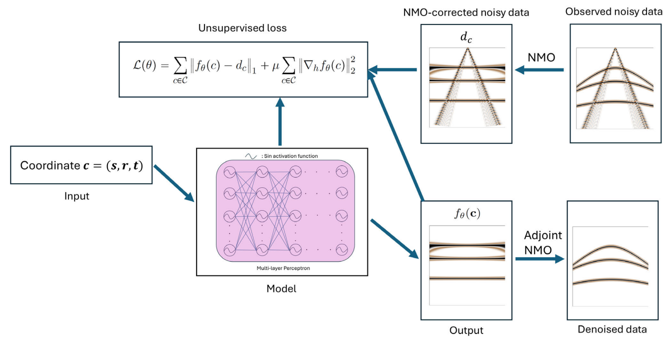

Fig. 1. Overview of the proposed framework for ground-roll separation using INR with horizontal-gradient regularization. NMO correction is applied using reflection velocities, which flatten reflection events while leaving dispersive ground roll uncorrected.

C. INRs for 3-D Ground-Roll Attenuation

We apply the horizontal-gradient regularization within the INRs training loop to exploit the spectral bias properties. After NMO correction, reflections become spatially flat and are fit early, while aliased ground roll retains steep dips and is suppressed by the lateral smoothness constraint. The INR is trained on the NMO-flattened data using the loss in (6), and the predicted reflections are restored to their original geometry by inverse NMO correction. A schematic of the complete workflow is shown in Fig. 1. Note that NMO correction is applied using reflection velocities and therefore does not flatten the dispersive ground roll, which remains largely uncorrected and only slightly modified.

The proposed INRs framework does not "eliminate" aliasing in the classical sense; rather, it reduces susceptibility to aliasing artifacts introduced by grid-based resampling and leverages strong priors to resolve ambiguities. Let S denote the sampling operator that evaluates a continuous wavefield f at the observed coordinate set $\mathcal{C}$. For regularly undersampled data with spatial content above Nyquist, there exist distinct fields $f \neq g$ such that $Sf = Sg$, i.e., the inverse problem is nonunique. In this regime, recovery necessarily depends on prior information. Our INRs formulation imposes such priors implicitly (spectral bias toward low coordinate frequency) and explicitly (horizontal-gradient regularization after NMO), yielding the variational problem

$$\min_{f\in\mathcal{F}} \sum_{\mathbf{c}\in\mathcal{C}} \left\|f(\mathbf{c}) - d_{\mathbf{c}}\right\|_1 + \mu \sum_{\mathbf{c}\in\mathcal{C}} \left\|\nabla_h f(\mathbf{c})\right\|_2^2$$

where f is parameterized by an SIREN network and ∇_h acts on source–receiver coordinates. Two practical consequences follow. First, because training is performed directly at the observed (possibly irregular) coordinates, no interpolation or binning onto a regular grid is required, avoiding the additional spectral folding and coherent aliasing that can occur during resampling. Second, when sampling is irregular or randomized, aliasing energy is distributed incoherently in coordinate space, allowing the combination of INRs' spectral bias and horizontal regularization to favor smooth, reflection-like solutions while suppressing steep, aliased ground roll. In contrast, under strictly regular decimation with severe spatial aliasing, the solution remains prior-determined: the INRs select the lowest coordinate-frequency field compatible with the samples, but cannot reconstruct high-wavenumber content that is not recoverable from the available data.

III. Experiments

In this section, we present tests on both synthetic and real land seismic datasets to evaluate the effectiveness of the proposed method. We begin with experiments on fully regularly sampled data for both synthetic and real examples. Subsequently, we apply a random undersampling operator to each original dataset to produce two additional versions, retaining only 50% and 25% of the traces with the remainder removed at random. These datasets allow us to assess ground-roll attenuation performance under both regular sampling and irregular undersampling conditions. For all 3-D cases, the data are partitioned into segments of size 1×128 samples with 50% overlap at the segment boundaries. Unless otherwise stated, a batch size of 128 and a learning rate of 1×10^{-5} are used in all experiments. For the simple synthetic example where the clean reflection signal is known, quantitative evaluation is performed using the signal-to-noise ratio (SNR). For the more complex synthetic and field datasets, the true reflection signal is unknown, and the evaluation is therefore based on qualitative analysis of reflection continuity and noise suppression.

A. 2-D Synthetic Data Example

We first evaluate the proposed algorithm on a simple 2-D synthetic dataset [Fig. 2(a)] consisting of three reflections and 300 time samples with a 4-ms sampling interval. The dataset contains 100 traces with a 20-m spacing. In addition to the reflection events, dispersive ground roll is simulated using a surface-wave model following the formulation described in [53]. In this model, ground roll is represented as a series of dispersive linear events characterized by frequency-dependent phase and group velocities

$$GR^{j,k} = \exp\left[i\left(\frac{f_0}{v_{p_j}} + \frac{f - f_0}{v_{g_j}}\right)x_k\right]$$

where x_k denotes the source–receiver offset, f is the frequency, f_0 is the central frequency of the wavelet, and v_{p_j} and v_{g_j} are the phase and group velocities, respectively. This formulation generates ground roll with dispersive characteristics similar to those observed in real land seismic data. NMO correction is first applied, after which the proposed method is used to separate reflections from ground roll. The separated reflections are then restored to their original geometry via adjoint NMO correction. Fig. 2(d) and (e) displays the final separated reflections and ground roll, respectively. For comparison, we also apply a conventional f–k filter, with its separated reflections and ground roll shown in Fig. 2(b) and (c). For quantitative evaluation, the SNR of the reconstructed reflections is computed. The f–k filtering result achieves an SNR of 4.99 dB, whereas the proposed method reaches 22.82 dB, indicating significantly improved ground-roll suppression while preserving the reflection signal.

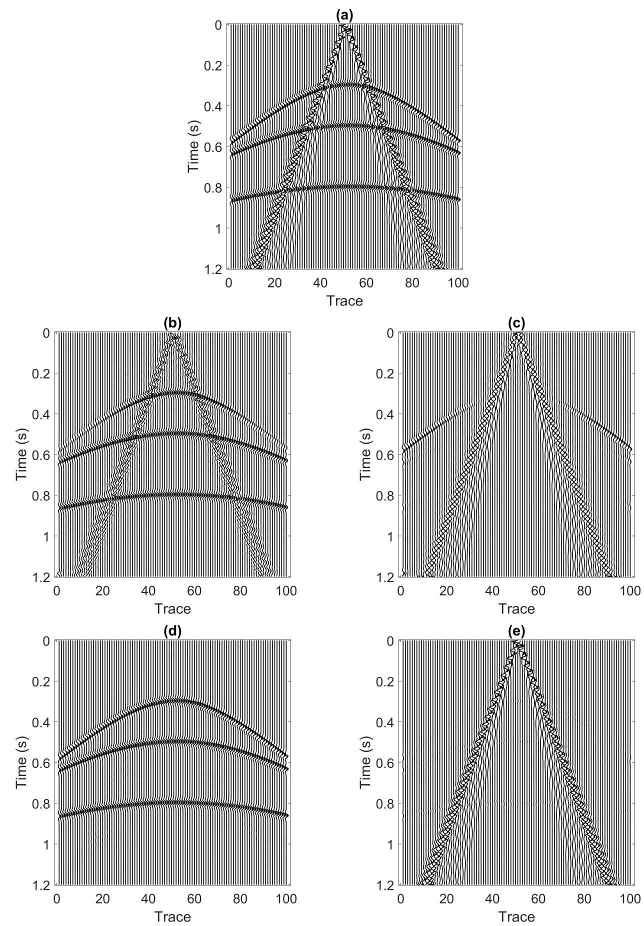

Fig. 2. 2-D synthetic data example. (a) Original data. (b) Separated reflections by f–k filter (SNR = 4.99 dB). (c) Separated ground roll by f–k. (d) Separated reflections by the proposed method (SNR = 22.82 dB). (e) Separated ground roll by the proposed method.

The f–k spectrum for this example is shown in Fig. 3. The ground roll in the original dataset is strongly aliased, causing substantial spectral overlap with the reflections. This overlap poses a significant challenge for traditional filters, such as the f–k or τ–p filter, which cannot cleanly separate the events in the frequency–wavenumber domain. As a result, the f–k filter leaves residual ground roll and induces signal leakage [Fig. 2(b)]. In contrast, the proposed method achieves a clear separation of reflections and ground roll.

To examine the effect of irregular sampling, we apply a random undersampling operator to generate two additional datasets, retaining only 50% and 25% of the original traces. The undersampled datasets and their separated results are shown in Figs. 4 and 5. Because the sampling is irregular, we omit the f–k spectra for these cases. Furthermore, since the fully sampled dataset already contains heavily aliased ground roll, undersampling further exacerbates the aliasing, making separation even more challenging for Fourier-domain filters. Despite these difficulties, the proposed method still achieves SNR of 20.31 and 18.38 dB for the 50% and 25% undersampling cases, respectively, indicating that the method remains effective even under severe irregular undersampling.

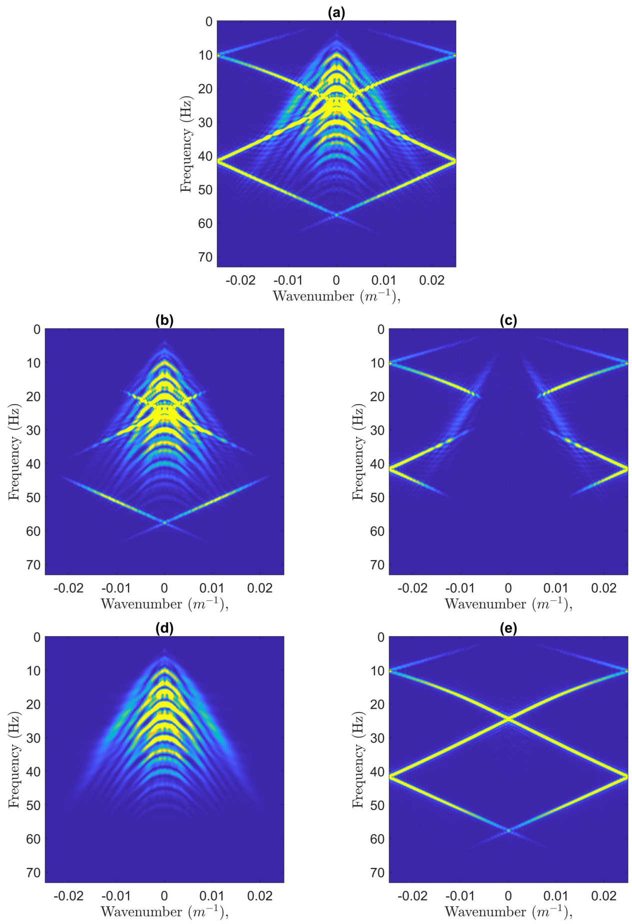

Fig. 3. Corresponding f–k spectra of the data shown in Fig. 2: (a) original data, (b) reflections separated by the f–k filter, (c) ground roll separated by the f–k filter, (d) reflections separated by the proposed method, and (e) ground roll separated by the proposed method.

B. 3-D Synthetic Data Example

We next evaluate the proposed method on a 3-D synthetic land seismic dataset generated using an elastic finite-difference scheme applied to a simple six-layer velocity model. The acquisition geometry consists of 501 receivers regularly spaced at 10-m intervals along a survey line, and 51 shots regularly spaced at 100-m intervals along the same line. Although the survey geometry is linear (2-D in the acquisition domain), the dataset is referred to as 3-D in the processing domain because each trace is indexed by two spatial coordinates (source and receiver positions) and one temporal coordinate. In this case, the three input dimensions to the network are the source x-coordinate, receiver x-coordinate, and the vertical–temporal coordinate. The method is first applied to the fully sampled dataset, with 10 out of the 51 shots randomly selected for display in Fig. 6.

Although the survey geometry is linear (2-D in the acquisition domain), the dataset is referred to as 3-D in the processing domain because each trace is indexed by two spatial coordinates (source and receiver positions) and one temporal coordinate.

Following the same procedure as in the 2-D example, we apply a random undersampling operator to generate two additional datasets by retaining only 50% and 25% of the traces.

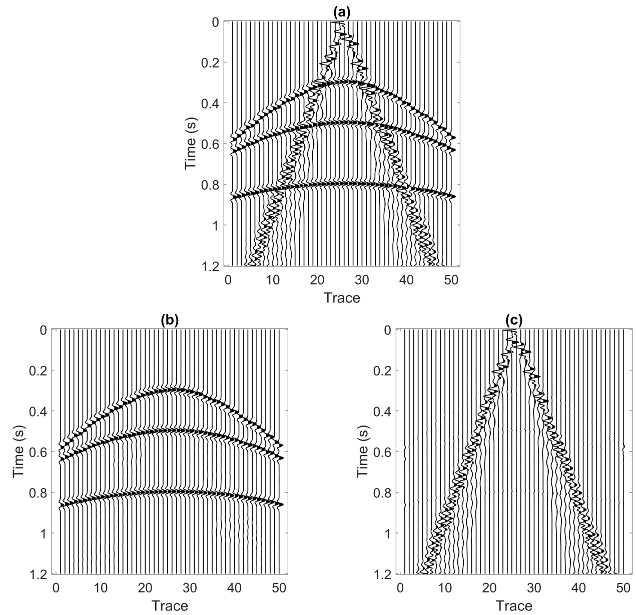

Fig. 4. 2-D irregularly undersampled data with 50% of traces randomly removed. (a) Original data. (b) Separated reflections (SNR = 20.31 dB). (c) Ground roll.

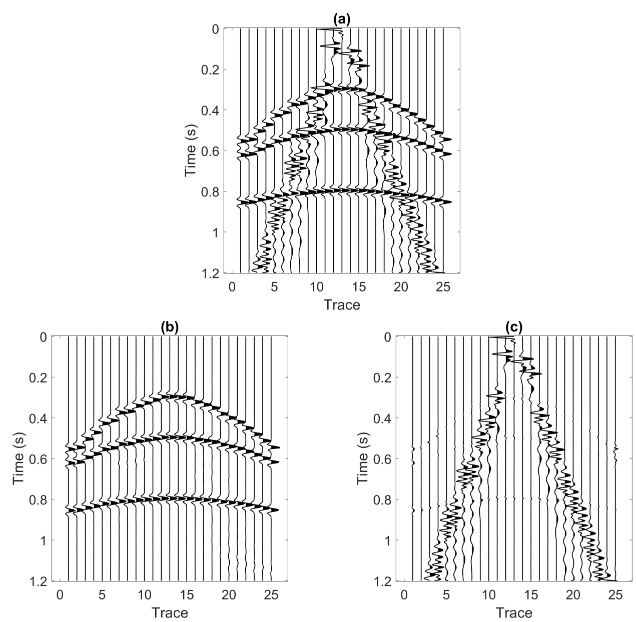

Fig. 5. 2-D irregularly undersampled data with 75% of traces randomly removed. (a) Original data. (b) Separated reflections (SNR = 18.38 dB). (c) Ground roll.

The undersampled datasets and their corresponding separated ground rolls are shown in Figs. 7 and 8. For both levels of undersampling, the proposed method successfully isolates the primary reflections while suppressing the steeply dipping ground roll. Even at 25% trace retention, the separated reflections remain laterally continuous and exhibit minimal leakage of ground-roll energy, while the extracted ground-roll sections show clear spatial coherence and negligible contamination from reflection events. This indicates that the combination of INRs representation and horizontal-gradient regularization

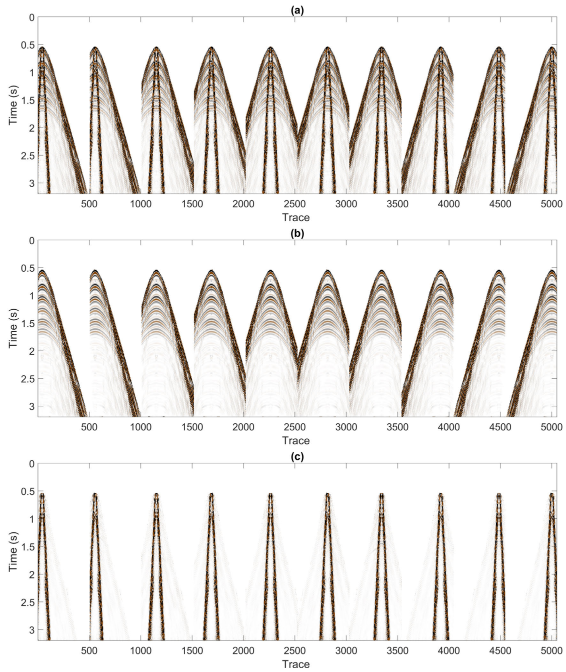

Fig. 6. Ten randomly selected shots from the 3-D synthetic land seismic data example. (a) Original data. (b) Separated reflections. (c) Ground roll.

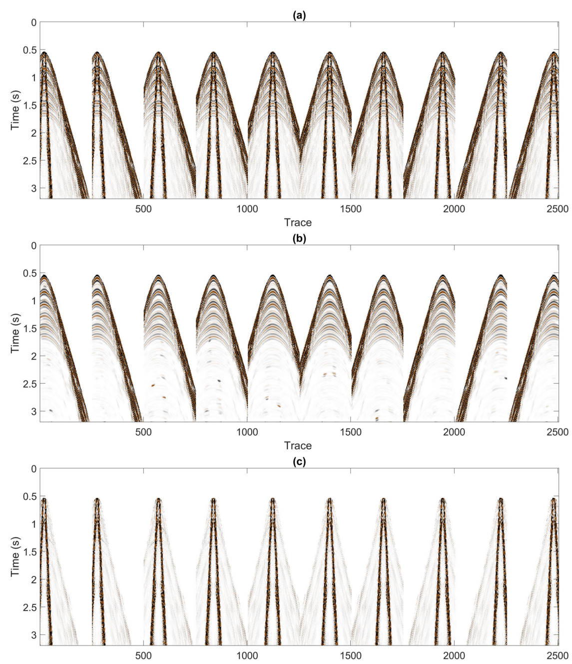

Fig. 7. Same ten shots from the 3-D synthetic land seismic data example with 50% of the data randomly removed. (a) Original data. (b) Separated reflections. (c) Ground roll.

remains robust to severe irregular undersampling in 3-D datasets, avoiding the signal smearing and noise residuals that typically arise in conventional filtering approaches under similar sampling conditions.

C. 3-D Field Data Example

Finally, we evaluate the proposed method on a real land seismic dataset. The survey geometry is illustrated in Fig. 9.

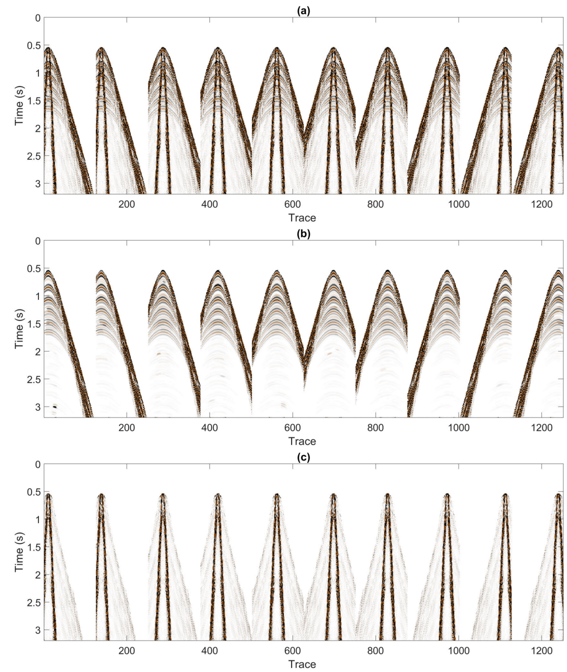

Fig. 8. Same ten shots from the 3-D synthetic land seismic data example with 75% of the data randomly removed. (a) Original data. (b) Separated reflections. (c) Ground roll.

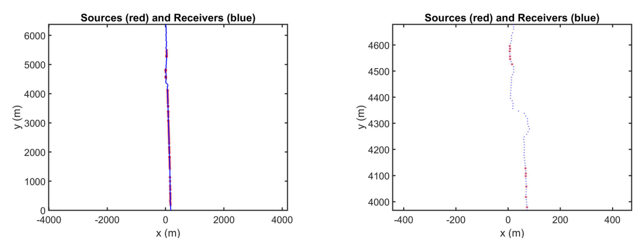

Fig. 9. Geometry of the field example. This dataset includes 639 unevenly spaced receivers and 111 unevenly spaced shots. The left panel shows the entire geometry, and the right panel shows a zoomed-in section.

The acquisition consists of 639 unevenly spaced receivers and 111 unevenly spaced shots. Compared to the crossline direction, the variation in the inline direction is minimal, resulting in a geometry that is close to a straight line along the crossline direction. Such irregular spacing, combined with the complex noise characteristics of real data, presents additional challenges for ground-roll attenuation compared to synthetic cases.

For this dataset, the input coordinates consist of four spatial dimensions: the *x*- and *y*-coordinates of the sources and the *x*- and *y*-coordinates of the receivers. As in the previous examples, we first apply the proposed method to the fully sampled dataset to establish a reference result. We then apply a random undersampling operator to remove 50% and 75% of the receivers, producing two additional irregularly undersampled datasets.

Fig. 10 presents a denoised shot gather from the real land seismic dataset, obtained by applying the proposed method both to a single 2-D shot gather and to the full 3-D dataset. The comparison shows that the 3-D approach produces reflections

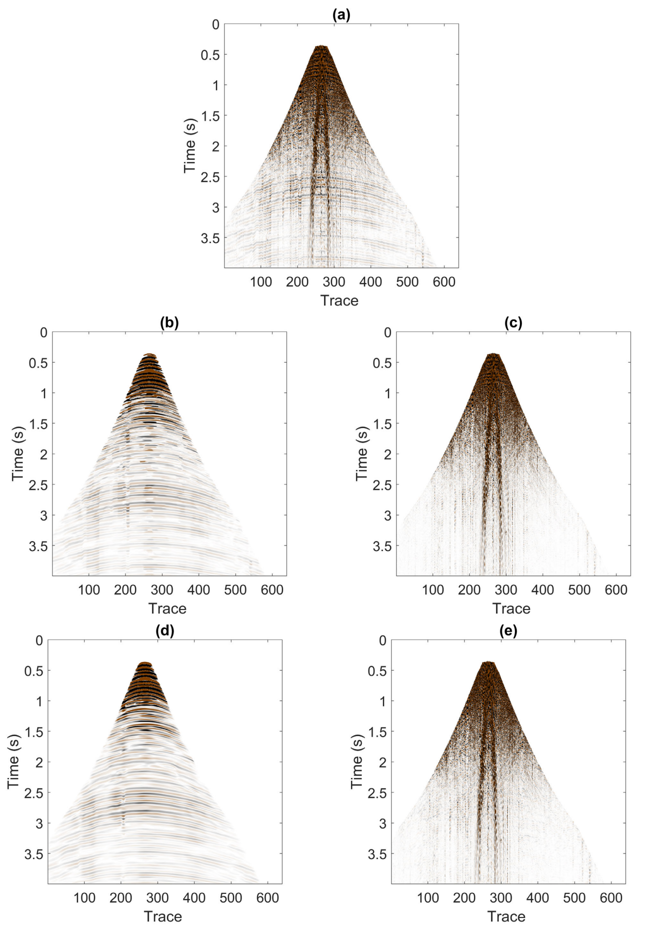

Fig. 10. Example shot from the 3-D land seismic dataset. (a) Original data. (b) Reflections recovered by applying the proposed method to a single 2-D shot gather. (c) Noise component estimated by the 2-D method. (d) Reflections recovered by applying the proposed method to the full 3-D dataset. (e) Noise component estimated by the 3-D method.

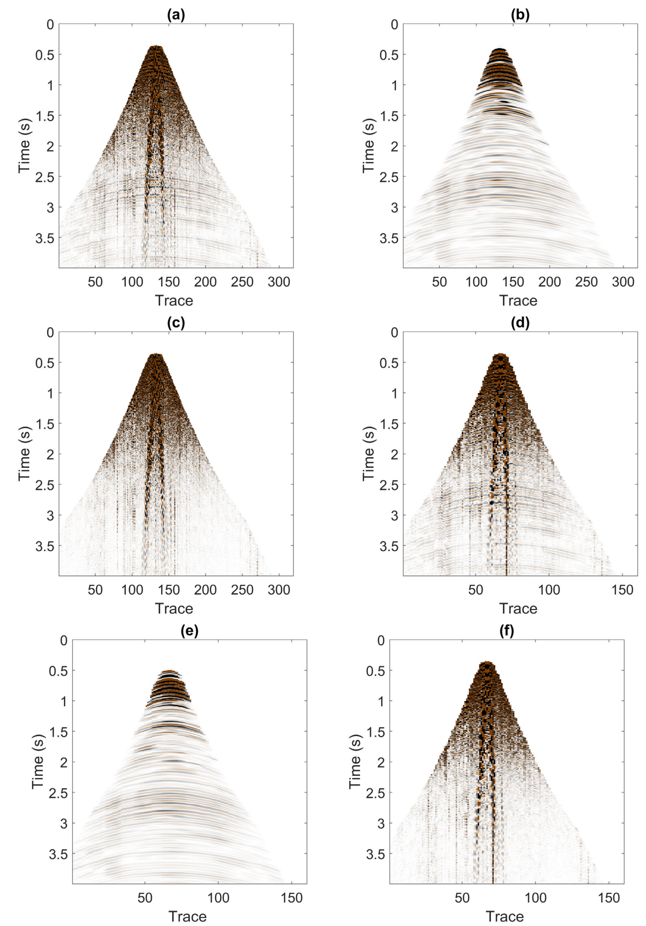

Fig. 11. Example shot from the 3-D land seismic dataset under irregular undersampling. (a) Data with 50% of receivers randomly removed. (b) Reflections recovered from (a). (c) Noise component estimated from (a). (d) Data with 75% of receivers randomly removed. (e) Reflections recovered from (d). (f) Noise component estimated from (d).

that are more continuous and smooth, particularly in the near-offset region where strong ground roll is present, whereas the 2-D method struggles to fully suppress the noise and preserve reflection continuity. This improvement can be attributed to the additional spatial information available in the 3-D formulation. When the network is trained on a single 2-D shot gather, the model relies only on information within that gather. In the near-offset region, where ground roll is strong and reflections may be heavily masked, the available information may be insufficient for the network to reliably recover the underlying signal. In contrast, the 3-D INR is trained using multiple shots and receiver positions simultaneously, allowing the network to exploit spatial correlations across the dataset. Because reflection events exhibit consistent structural patterns across neighboring shots and receivers, the network can learn these patterns from less contaminated regions and use them to reconstruct reflections in areas where they are obscured by strong ground roll. This implicit spatial prior provided by the 3-D coordinate representation improves the robustness of noise suppression and leads to more continuous reflection events in the near-offset region.

Fig. 11 presents a representative shot gather from the land seismic dataset under irregular undersampling, with 50% and 75% of receivers randomly removed. The corresponding denoised results obtained using the proposed method are also shown. To further evaluate the robustness of the method across different source positions, Figs. 12–14 display ten randomly selected shot gathers for the fully sampled and undersampled datasets, each accompanied by their denoised counterparts. These examples collectively demonstrate that the proposed method maintains effective ground-roll attenuation even under severe irregular acquisition conditions.

We also generated the final stacked section along the survey line, as shown in Fig. 9. The stacked sections corresponding to the fully sampled dataset are presented in Fig. 15. Specifically, Fig. 15(a) shows the original data containing strong ground roll, Fig. 15(b) presents the denoised result obtained using the proposed method, and Fig. 15(c) displays the separated ground roll. For improved visualization, Fig. 16 provides a zoomed-in view of the stacked sections. From the stacked sections, it is evident that the strong ground roll severely obscures some reflection events, particularly in the shallow region, making them difficult to interpret. After applying the

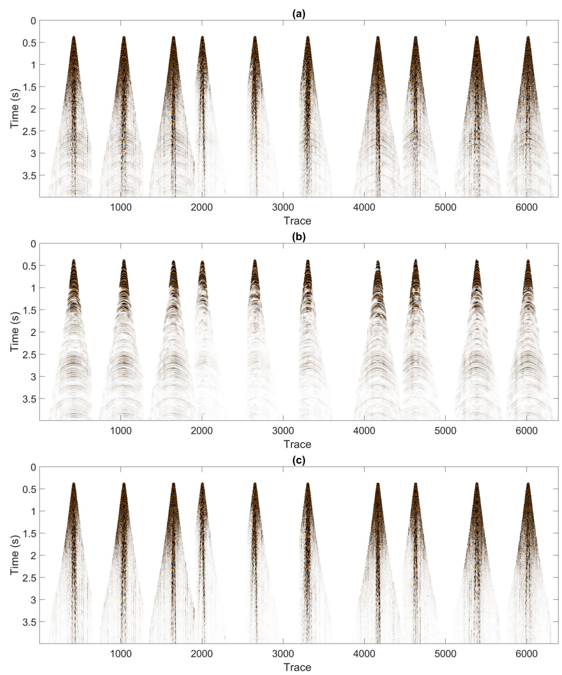

Fig. 12. Ten randomly selected shots from the 3-D field seismic data example. (a) Original data. (b) Separated reflections. (c) Ground roll.

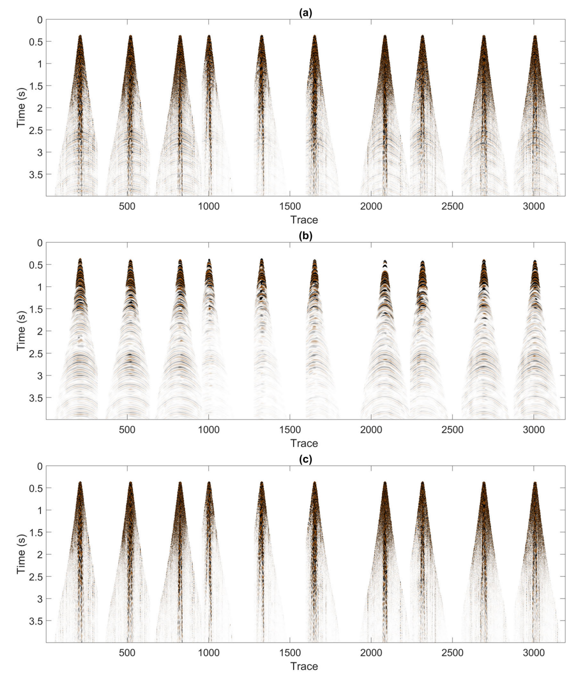

Fig. 13. Same ten shots from the 3-D field seismic data example with 50% of the data randomly removed. (a) Original data. (b) Separated reflections. (c) Ground roll.

proposed method, these events become clearly visible, indicating effective ground-roll attenuation. Finally, we compare the amplitude spectra of the original and denoised data in Fig. 17. The spectra confirm that the proposed method successfully removes the low-frequency ground roll while preserving the spectral content of the reflection signals.

IV. Discussion

A. Comparison With Existing Methods

Existing ground-roll attenuation methods include traditional transform-domain approaches, supervised deep learning methods, and unsupervised neural network-based techniques. Traditional methods, such as f–k filtering, τ–p transforms, and sparsity-based methods, are computationally efficient and effective when signal and noise are well separated in the transform domain. However, their performance can be limited in undersampled or irregularly sampled data, where aliased ground-roll energy overlaps with reflection events. Supervised CNN-based methods can achieve good denoising performance when sufficient paired noisy and clean training data are available. In practical field applications, however, reliable clean labels are usually difficult to obtain. Unsupervised CNN-based methods, such as deep image prior (DIP), avoid the need for labels, but they generally operate on regularly sampled grids. Therefore, irregularly sampled data often need to be interpolated before processing, which may introduce artifacts and affect the separation of ground roll and reflections. The proposed INR-based method provides an alternative by modeling the seismic wavefield as a continuous coordinate-based function. This allows the network to operate directly on irregularly sampled receiver and source coordinates without prior interpolation. In addition, the 3-D formulation enables

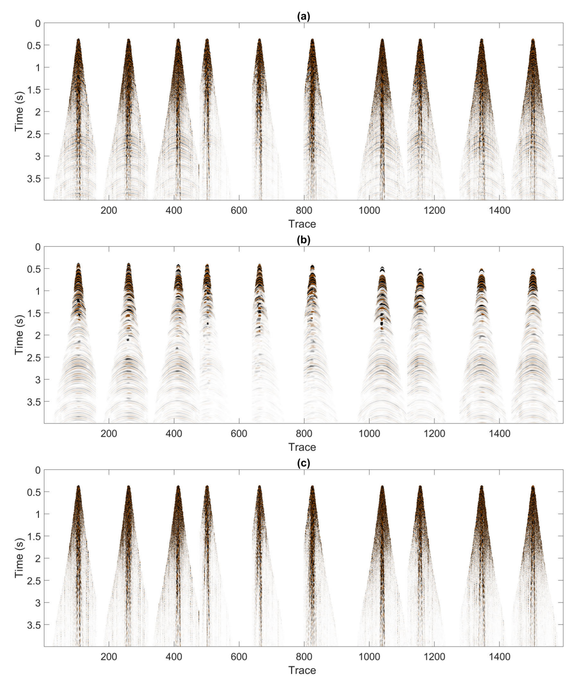

Fig. 14. Same ten shots from the 3-D field seismic data example with 75% of the data randomly removed. (a) Original data. (b) Separated reflections. (c) Ground roll.

the model to jointly use information from multiple spatial directions, improving reflection preservation when ground roll is strong or aliased. Compared with conventional filtering and CNN-based approaches, the proposed method is therefore more suitable for challenging irregularly sampled 3-D datasets.

It should be noted that INR-based methods require optimizing the network parameters for each dataset, which leads to higher computational cost than fast transform-domain filters

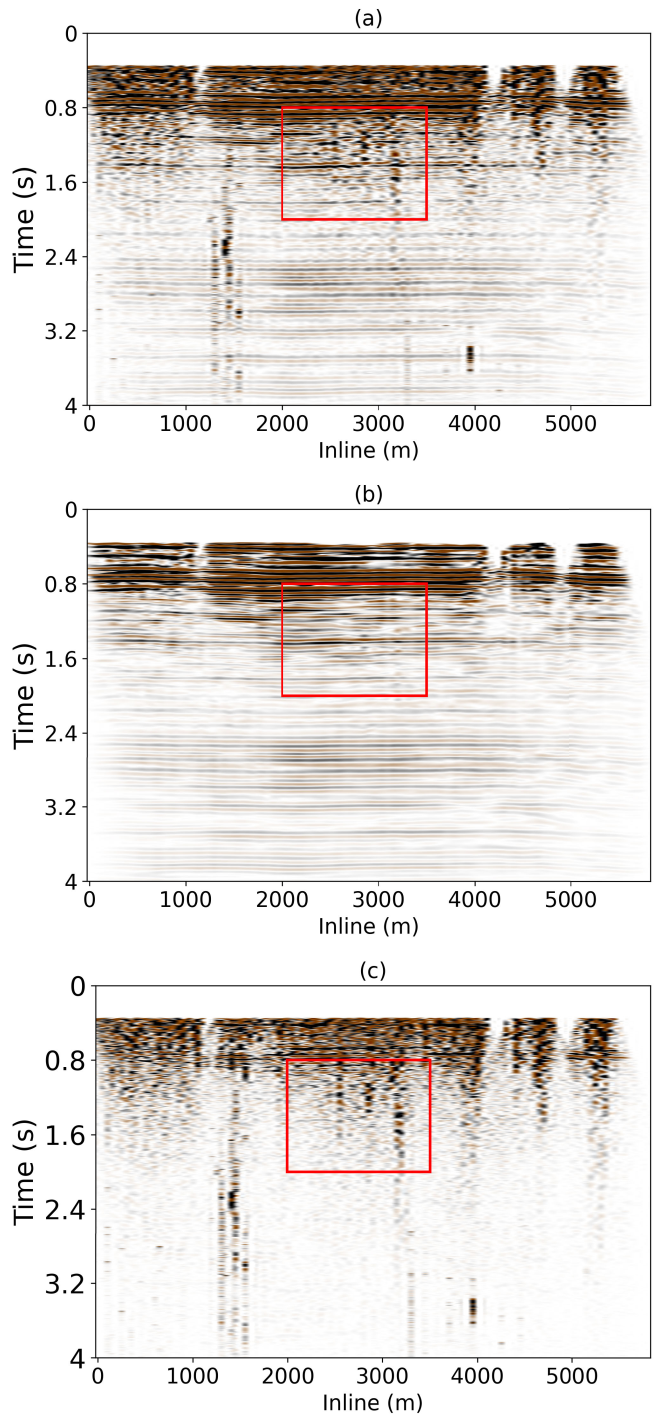

Fig. 15. Stacked section along the survey line. (a) Original data. (b) Denoised result. (c) Ground roll.

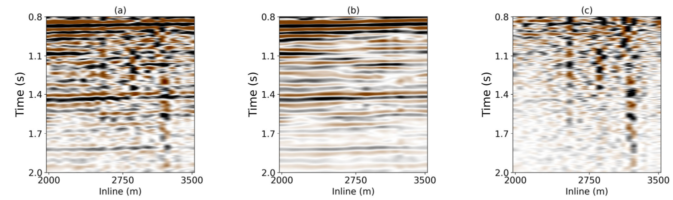

Fig. 16. Zoomed-in stacked section along the survey line. (a) Original data. (b) Denoised result. (c) Ground roll.

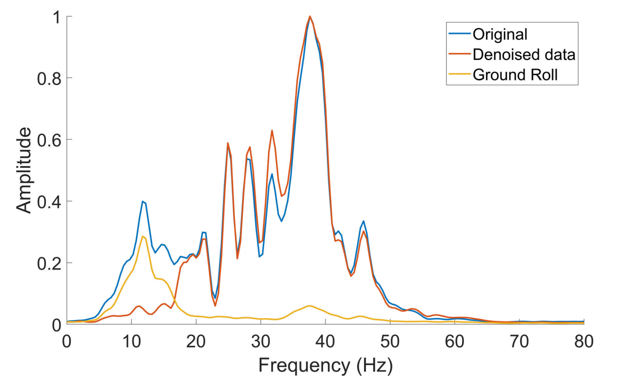

Fig. 17. Normalized frequency spectrum of the original data, denoised data, and the ground roll.

or pretrained CNN inference. However, for datasets where irregular sampling and strong aliasing limit the effectiveness of conventional methods, this additional cost can be justified by the improved robustness of the proposed framework.

B. Computational Cost and Practical Considerations

Compared with traditional transform-domain filtering approaches, such as f–k filtering, the proposed INR-based method requires iterative optimization and therefore has a higher computational cost. All experiments in this study were conducted on a workstation equipped with an NVIDIA RTX 4070 Ti GPU (12-GB memory) and an Intel Core i7-14700K CPU.

For a single 2-D shot gather, the optimization typically requires less than 1 min of training to converge. The computational cost scales approximately linearly with the number of coordinate samples used for training. For the field dataset considered in this study, which contains hundreds of thousands of traces, the complete 3-D denoising process requires several hours of computation on the above hardware.

In comparison, traditional f–k filtering can process large seismic datasets very efficiently due to its reliance on FFT-based transforms. Supervised CNN-based methods can also achieve fast inference once a model has been trained offline. However, these approaches typically assume regularly sampled data and may require interpolation or preprocessing when applied to irregular acquisition geometries. Unsupervised CNN-based approaches, such as DIP, similarly require iterative optimization for each dataset, resulting in computational costs comparable to those of the proposed INR framework.

For larger field-scale 3-D applications, processing the entire data volume at once may become computationally demanding. One possible strategy is to divide the data into several subvolumes and process them separately. However, for coordinate-based INR models, sequential fine-tuning between subvolumes is not entirely straightforward. If global coordinates are retained, an INR trained on one subvolume may not directly generalize to another subvolume located in a different coordinate range. If local coordinates are renormalized for each subvolume, warm-start fine-tuning may become more feasible, but global coordinate consistency may be weakened. Therefore, although subvolume processing with possible warm-start initialization may improve computational efficiency, its effectiveness depends on factors, such as subvolume size, overlap, coordinate normalization, and the continuity of wavefield structures between neighboring subvolumes. A systematic evaluation of these factors is beyond the scope of the present study and will be considered in future work.

In practice, the proposed method is intended for challenging datasets where strong aliasing or irregular acquisition geometries limit the effectiveness of conventional filtering techniques. In such cases, the additional computational cost can be justified by the improved robustness in separating ground roll from reflection events. Developing strategies to further improve the computational efficiency of INR-based seismic processing, for example through more efficient optimization schemes, parallel implementations, or carefully designed subvolume processing, represents an important direction for future work.

C. Robustness to Imperfect NMO Correction

The NMO velocity used for the field dataset was obtained using conventional semblance analysis followed by manual picking of the velocity spectrum. In land seismic data, velocity estimation can be challenging due to strong ground roll, statics, and irregular acquisition geometries, and therefore, the resulting velocities may contain uncertainties. As a result, reflections after NMO correction are not always perfectly flattened.

To evaluate the sensitivity of the proposed method to such velocity errors, we performed an additional experiment using the synthetic dataset. In this test, the NMO velocity was intentionally perturbed by ±5% and ±10% relative to the true velocity before applying the proposed method. These perturbations introduce residual moveout in the reflection events and therefore simulate imperfect NMO correction.

The results are shown in Fig. 18. Despite the imperfect flattening of the reflections, the proposed method continues to effectively suppress ground roll and recover coherent reflection events. Although the separation performance gradually degrades as the velocity error increases, the method remains stable for moderate velocity uncertainties. These results indicate that the proposed framework does not require perfectly flattened reflections and is reasonably robust to typical velocity-picking errors encountered in practical seismic processing.

Another practical consideration is NMO stretch, which can distort waveforms at large offsets after NMO correction. Because the proposed framework represents the seismic wavefield using an INR, the network itself can accommodate such waveform distortions without difficulty. However, the horizontal-gradient regularization introduced in this work assumes that reflections are approximately flat after NMO correction. When NMO stretch becomes significant at far offsets, reflections may deviate from perfect horizontality, which can slightly reduce the effectiveness of this regularization. In practice, this effect is limited because NMO stretch mainly affects the far-offset region, and the main reflection energy remains approximately flat in most parts of the gather.

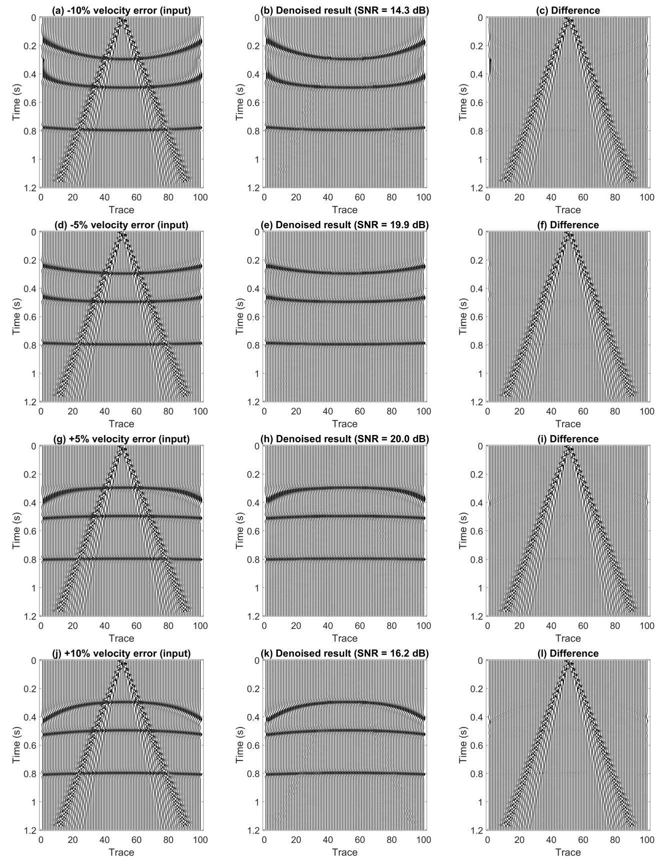

Fig. 18. Sensitivity of the proposed method to NMO velocity errors using the synthetic example. The NMO velocity used for correction is intentionally perturbed relative to the true velocity. (a) Input data after NMO correction with a −10% velocity error, (b) corresponding denoised result obtained by the proposed method, and (c) difference between the input data and the denoised result. Quantities for (d)–(f) −5%, (g)–(i) +5%, and (j)–(l) +10% velocity errors. The SNR values of the denoised results are reported in the corresponding panels.

D. Sensitivity to the Horizontal Regularization Parameter μ

The parameter μ controls the strength of the horizontal-gradient regularization in the proposed loss function. This term encourages lateral smoothness after NMO correction and helps suppress steep spatial variations associated with ground roll. Because μ balances the tradeoff between data fidelity and regularization, its value has a direct influence on the denoising performance.

To evaluate the sensitivity of the proposed method to this parameter, we performed additional experiments using the synthetic dataset. The value of μ was varied over a wide range, and the SNR of the reconstructed reflections was measured. The results are shown in Fig. 19.

When μ is small, the horizontal regularization is weak and the network primarily fits the observed data. As a result, a significant portion of the ground-roll energy remains in the reconstructed result, leading to relatively low SNR values. As μ increases, the regularization more effectively suppresses the

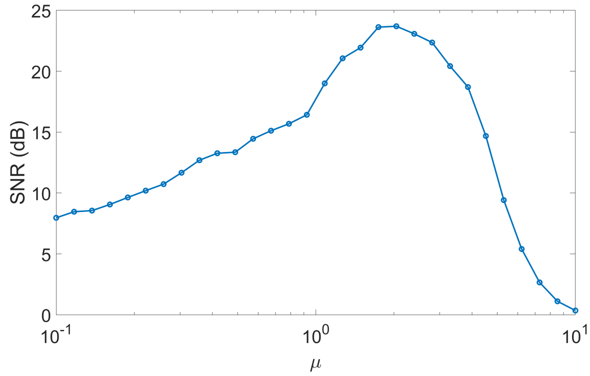

Fig. 19. Sensitivity of the proposed method to the horizontal regularization parameter μ.

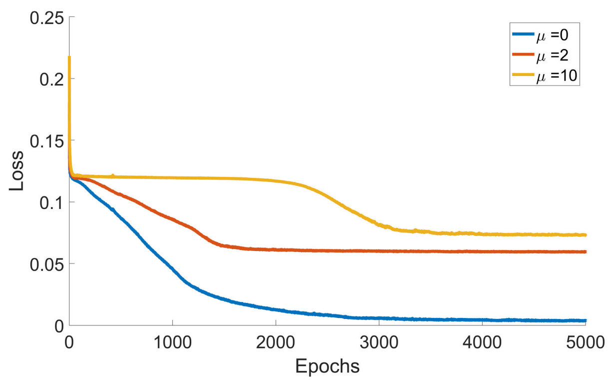

Fig. 20. Loss evolution with training iterations for different values of μ, illustrating the effect of horizontal regularization on convergence and overfitting.

spatial variations associated with ground roll, and the SNR improves accordingly.

However, when μ becomes too large, the regularization term dominates the optimization objective. In this case, the network tends to underfit the data and the convergence becomes slower. Although acceptable results can still be obtained with a larger number of training iterations, the reconstruction quality degrades when the number of epochs is fixed, resulting in reduced SNR values for large μ.

These results indicate that the proposed method is relatively stable within a moderate range of μ values. In practice, μ can be selected empirically by balancing ground-roll suppression and reflection preservation. In this work, a value near the peak of the curve in Fig. 19 was used for both synthetic and field datasets.

Fig. 20 shows the evolution of the loss function with extended training epochs for three different values of μ. When $\mu = 0$, corresponding to no regularization, the loss continues to decrease toward zero as training proceeds. In this case, the network progressively fits both the horizontal reflections and the ground roll, indicating overfitting and the need for manual early stopping.

For a suitable value of $\mu = 2$, the loss curve stabilizes after an initial decrease and does not continue to decrease with further training. This behavior indicates that the horizontal regularization effectively prevents the network from fitting the ground roll, allowing the training process to converge without the need for manual early stopping.

When $\mu = 10$, the convergence rate becomes slower due to the strong regularization. Although the network can still recover most reflection events with extended training, some signal loss is observed compared to the case with $\mu = 2$. This effect is primarily attributed to residual NMO stretch at far offsets, which leads to imperfect flattening of reflections and reduces the effectiveness of the horizontal regularization.

V. Conclusion

We present an unsupervised 3-D ground-roll attenuation method based on INRs. Unlike traditional filtering techniques or supervised deep learning approaches, the proposed method operates directly in the spatial coordinate domain and requires no labeled training data. By applying NMO correction and exploiting both the INR's spectral bias toward low coordinate-frequency components and a horizontal derivative regularization term, the method suppresses steeply dipping, aliased ground roll while preserving horizontally aligned reflections. The horizontal regularization enforces lateral smoothness, improves stability, and allows the network to be trained to convergence without overfitting to aliased noise. The method is evaluated on both fully sampled and irregularly undersampled synthetic and real 3-D seismic datasets, demonstrating consistent removal of aliased ground roll while preserving reflection events, even under irregular acquisition geometries. Compared to a 2-D implementation, the 3-D formulation achieves superior denoising performance, producing smoother and more continuous reflections in regions where strong ground roll obscures the signal.

An interesting extension of the proposed framework is to jointly perform ground-roll attenuation and interpolation using the continuous representation capability of INRs, which may further improve performance under severe undersampling and will be investigated in future work.

References

[1] D. G. Stone, *Designing Seismic Surveys in Two and Three Dimensions*. Houston, TX, USA: Society of Exploration Geophysicists, 1994.

[2] A. Cordsen, M. Galbraith, and J. Peirce, *Planning Land 3-D Seismic Surveys*. Houston, TX, USA: Society of Exploration Geophysicists, 2000.

[3] M. Bacon, R. Simm, and T. Redshaw, *3-D Seismic Interpretation*. Cambridge, U.K.: Cambridge Univ. Press, 2007.

[4] Ö. Yilmaz, *Seismic Data Analysis: Processing, Inversion, and Interpretation of Seismic Data*. Houston, TX, USA: Society of Exploration Geophysicists, 2001.

[5] D. C. Henley, "Coherent noise attenuation in the radial trace domain," *Geophysics*, vol. 68, no. 4, pp. 1408–1416, Jan. 2003.

[6] S. A. Hosseini, A. Javaherian, H. Hassani, S. Torabi, and M. Sadri, "Adaptive attenuation of aliased ground roll using the shearlet transform," *J. Appl. Geophys.*, vol. 112, pp. 190–205, Jan. 2015.

[7] S. Verma, S. Guo, T. Ha, and K. J. Marfurt, "Highly aliased ground-roll suppression using a 3D multiwindow karhunen-Loève filter: Application to a legacy Mississippi lime survey," *Geophysics*, vol. 81, no. 1, pp. V79–V88, Jan. 2016.

[8] Y. Wang, "Seismic trace interpolation in the $f-x-y$ domain," *Geophysics*, vol. 67, no. 4, pp. 1232–1239, Jan. 2002.

[9] M. Naghizadeh and M. Sacchi, "Ground-roll attenuation using curvelet downscaling," *Geophysics*, vol. 83, no. 3, pp. V185–V195, May 2018.

[10] X. Liu, "Ground roll suppression using the Karhunen-Loeve transform," *Geophysics*, vol. 64, no. 2, pp. 564–566, Apr. 1999.
[11] H. Karsli and Y. Bayrak, "Ground-roll attenuation based on Wiener filtering and benefits of time-frequency imaging," *Lead. Edge*, vol. 27, no. 2, pp. 206–209, Feb. 2008.
[12] M. I. Al-Husseini, J. B. Glover, and B. J. Barley, "Dispersion patterns of the ground roll in eastern Saudi Arabia," *Geophysics*, vol. 46, no. 2, pp. 121–137, Feb. 1981.
[13] R. Saatcilar and N. Canitez, "A method of ground-roll elimination," *Geophysics*, vol. 53, no. 7, pp. 894–902, Jul. 1988.
[14] G. Beresford-Smith and R. Rango, "Suppression of ground roll by windowing in two domains," *1st Break*, vol. 7, no. 2, pp. 55–63, 1989.
[15] S. K. Chiu, "3D attenuation of aliased ground roll on randomly undersampled data," in *Proc. SEG Int. Expo. Annu. Meeting*, Aug. 2019, p. D043S110R001.
[16] S. Treitel, J. L. Shanks, and C. W. Frasier, "Some aspects of fan filtering," *Geophysics*, vol. 32, no. 5, pp. 789–800, Oct. 1967.
[17] R. H. Tatham, J. W. Keeney, and L. Noponen, "Application of the tau-p transform (slant-stack) in processing seismic reflection data," *Explor. Geophys.*, vol. 14, no. 3, pp. 163–172, Sep. 1983.
[18] G. Turner, "Aliasing in the tau-p transform and the removal of spatially aliased coherent noise," *Geophysics*, vol. 55, no. 11, pp. 1496–1503, Nov. 1990.
[19] S. A. Hosseini, A. Javaherian, H. Hassani, S. Torabi, and M. Sadri, "Shearlet transform in aliased ground roll attenuation and its comparison with f-k filtering and curvelet transform," *J. Geophys. Eng.*, vol. 12, no. 3, pp. 351–364, Jun. 2015.
[20] Y. Hu, L. Wang, F. Cheng, Y. Luo, C. Shen, and B. Mi, "Ground-roll noise extraction and suppression using high-resolution linear radon transform," *J. Appl. Geophys.*, vol. 128, pp. 8–17, May 2016.
[21] A. Dev and G. A. McMechan, "Spatial antialias filtering in the slowness-frequency domain," *Geophysics*, vol. 74, no. 2, pp. V35–V42, Jan. 2009.
[22] A. J. Deighan and D. R. Watts, "Ground-roll suppression using the wavelet transform," *Geophysics*, vol. 62, no. 6, pp. 1896–1903, Nov. 1997.
[23] W. Wang, J. Gao, W. Chen, and J. Xu, "Data adaptive ground-roll attenuation via sparsity promotion," *J. Appl. Geophys.*, vol. 83, pp. 19–28, Aug. 2012.
[24] X. Chen, W. Chen, X. Wang, and W. Wang, "Sparsity-optimized separation of body waves and ground-roll by constructing dictionaries using tunable Q-factor wavelet transforms with different Q-factors," *Geophys. J. Int.*, vol. 211, no. 1, pp. 621–636, Oct. 2017.
[25] C. Yarham, U. Boeniger, and F. Herrmann, "Curvelet-based ground roll removal," in *Proc. SEG Tech. Program Expanded Abstracts*, 2006, pp. 2777–2782.
[26] G. Hennenfent and F. J. Herrmann, "Irregular sampling–from aliasing to noise," in *Proc. 69th EAGE Conf. Exhib. Incorporating SPE EUROPEC*, 2007.
[27] R. Sternfels, A. Prescott, G. Pignot, L. Tian, and D. Le Meur, "Irregular spatial sampling and rank-reduction–interpolation by joint low-rank and sparse inversion," in *Proc. 78th EAGE Conf. Exhib.*, 2016, pp. 1–5.
[28] S. Winograd, "On computing the discrete Fourier transform," *Math. Comput.*, vol. 32, no. 141, pp. 175–199, 1978.
[29] A. J. W. Duijndam, M. A. Schonewille, and C. O. H. Hindriks, "Reconstruction of band-limited signals, irregularly sampled along one spatial direction," *Geophysics*, vol. 64, no. 2, pp. 524–538, Apr. 1999.
[30] J. A. Fessler and B. P. Sutton, "Nonuniform fast Fourier transforms using min-max interpolation," *IEEE Trans. Signal Process.*, vol. 51, no. 2, pp. 560–574, Feb. 2003.
[31] S. Bagchi and S. K. Mitra, *The Nonuniform Discrete Fourier Transform and Its Applications in Signal Processing*, vol. 463. Cham, Switzerland: Springer, 2012.
[32] H. Li, W. Yang, and X. Yong, "Deep learning for ground-roll noise attenuation," in *Proc. SEG Tech. Program Expanded Abstr.*, Aug. 2018, pp. 1981–1985.
[33] H. Kaur, S. Fomel, and N. Pham, "Seismic ground-roll noise attenuation using deep learning," *Geophys. Prospecting*, vol. 68, no. 7, pp. 2064–2077, Sep. 2020.
[34] D. Liu, M. D. Sacchi, X. Wang, and W. Chen, "Unsupervised deep learning for ground roll and scattered noise attenuation," *IEEE Trans. Geosci. Remote Sens.*, vol. 61, 2023, Art. no. 5920613.
[35] V. Sitzmann, J. Martel, A. Bergman, D. Lindell, and G. Wetzstein, "Implicit neural representations with periodic activation functions," in *Proc. NIPS*, 2020, pp. 7462–7473.
[36] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "NeRF: Representing scenes as neural radiance fields for view synthesis," *Commun. ACM*, vol. 65, no. 1, pp. 99–106, Jan. 2022.
[37] M. Tancik et al., "Fourier features let networks learn high frequency functions in low dimensional domains," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 33, 2020, pp. 7537–7547.
[38] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proc. IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
[39] A. van den Oord et al., "WaveNet: A generative model for raw audio," 2016, *arXiv:1609.03499*.
[40] J. N. P. Martel, D. B. Lindell, C. Z. Lin, E. R. Chan, M. Monteiro, and G. Wetzstein, "ACORN: Adaptive coordinate networks for neural scene representation," 2021, *arXiv:2105.02788*.
[41] N. Rahaman et al., "On the spectral bias of neural networks," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 5301–5310.
[42] J. Li, D. Liu, D. Trad, and M. D. Sacchi, "Robust unsupervised 5D seismic data reconstruction on regular and irregular grids," *Geophysics*, vol. 89, no. 6, pp. V537–V549, Jan. 2024.
[43] W. Xu, D. Liu, X. Wang, M. D. Sacchi, and W. Chen, "Implicit neural representations for self-supervised seismic deblending," *Geophysics*, vol. 90, no. 6, pp. 1–148, 2025.
[44] B. Kang, R. Chen, K. Yang, M. Li, and B. Wu, "Implicit full waveform inversion with adaptive Fourier frequency bases learning," *Geophys. J. Int.*, vol. 244, no. 1, Nov. 2025, Art. no. ggaf404.
[45] R. Chen, B. Wu, M. Li, and Y. Luo, "Full-waveform inversion with velocity model low-rank implicit neural representation," *IEEE Trans. Geosci. Remote Sens.*, vol. 63, 2025, Art. no. 5918816.
[46] J. Li, D. Liu, and M. D. Sacchi, "Unsupervised ground-roll attenuation via implicit neural representations," *Geophysics*, vol. 90, no. 2, pp. V111–V121, Jan. 2025.
[47] Z.-Q. J. Xu, Y. Zhang, and Y. Xiao, "Training behavior of deep neural network in frequency domain," in *Proc. Int. Conf. Neural Inf. Process.*, 2019, pp. 264–274.
[48] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, "Understanding deep learning (still) requires rethinking generalization," *Commun. ACM*, vol. 64, no. 3, pp. 107–115, Mar. 2021.
[49] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The unreasonable effectiveness of deep features as a perceptual metric," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 586–595.
[50] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Phys. D, Nonlinear Phenomena*, vol. 60, nos. 1–4, pp. 259–268, Nov. 1992.
[51] I. Bayram and M. E. Kamasak, "A directional total variation," in *Proc. 20th Eur. Signal Process. Conf. (EUSIPCO)*, Aug. 2012, pp. 265–269.
[52] J. Zhang, R. Lai, and C.-C. J. Kuo, "Adaptive directional total-variation model for latent fingerprint segmentation," *IEEE Trans. Inf. Forensics Security*, vol. 8, no. 8, pp. 1261–1273, Aug. 2013.
[53] D. Le Meur, N. Benjamin, R. Cole, and M. Al Harthy, "Adaptive groundroll filtering," in *Proc. 70th EAGE Conf. Exhib. Incorporating SPE EUROPEC*, 2008.