Implicit Neural Representations for Unsupervised Seismic Data Processing

Ji Li, Daniel Trad

ABSTRACT

Implicit Neural Representation (INR) has recently gained prominence across various domains within computer vision and signal processing. INR involves learning a continuous function over a set of points, representing a significant advancement in seismic signal analysis when contrasted with conventional methods. Notably, the network's intrinsic low-frequency bias encourages the acquisition of self-similar features ahead of high-frequency, incoherent ones during training. This unique characteristic can effectively mitigate coherent and incoherent noise in seismic data. Furthermore, the capacity to represent signals continuously, unlike discrete forms, facilitates the interpolation of seismic signals on irregular grids. In this paper, we present an approach utilizing the Siren (a specific INR architecture with a sine activation function) to address various seismic processing challenges, including denoising for both coherent (e.g., ground roll) and incoherent (random and erratic) noise as well as seismic data interpolation. Notably, our method stands out because it obviates the need for paired training datasets, rendering it a zero-shot unsupervised learning approach. Evaluation results demonstrate its superior performance compared to other state-of-the-art unsupervised deep learning techniques and traditional methods.

INTRODUCTION

The reconstruction of seismic data from incomplete and noisy sources is a crucial aspect of seismic data processing. Traditional methods predominantly rely on sparse transformation techniques, where seismic data is transformed into domains such as Radon (Durrani and Bisset, 1984), Fourier (Gülünay, 2003; Xu et al., 2005), Wavelet (Jian et al., 2006), and Curvelet (Candes et al., 2006), wherein the signals exhibit sparsity. Proximal methods like FISTA (Beck and Teboulle, 2009) and ADMM (Boyd et al., 2011; Wen et al., 2016) are commonly employed for solving sparse inversion problems in traditional seismic data reconstruction. However, the landscape has evolved with the rise of deep learning, offering superior performance in various seismic data processing applications. These encompass seismic inversion (Zheng et al., 2019; Li et al., 2019), denoising of random and erratic data (Richardson and Feller, 2019; Liu et al., 2018; Wang and Hu, 2021), seismic interpolation (Wang et al., 2019; Kaur et al., 2021; Oliveira et al., 2018), and ground roll attenuation (Li et al., 2018; Kaur et al., 2020).

Deep learning techniques fall into two primary categories: supervised learning (SL) and unsupervised learning (UL). SL relies on abundant paired clean and noisy datasets for training, a challenging demand given the scarcity of such paired data in real seismic domains. Consequently, various unsupervised and self-supervised methods have emerged(Liu et al., 2020; Qian et al., 2022; Saad et al., 2021), drawing on concepts like Deep Image Prior (DIP) (Ulyanov et al., 2018) and Noise2Noise (Lehtinen et al., 2018), along with it's variants (Krull et al., 2019; Pang et al., 2021).

Implicit neural representation (INR) has emerged as a promising signal processing framework characterized by a multilayer perceptron (MLP) incorporating linear layers and element-wise nonlinear activation functions. In contrast to CNNs, INRs lack locality biases, contributing to their enhanced performance. Recent successes of INR span applications like surface representation (Sitzmann et al., 2020), volume rendering (Martin-Brualla et al., 2021; Mildenhall et al., 2021), and generative modelling (Chan et al., 2021). In the context of seismic signal processing, Goyes-Peñafiel et al. (2023) applied an MLP network with Fourier mapping for 3D irregular seismic data interpolation.

This paper proposes an exploration of the efficacy of INR, specifically leveraging Sinusoidal Representation Networks (SIREN) (Sitzmann et al., 2020), in addressing seismic data processing challenges. Our approach employs Implicit Neural Representations with Periodic Activation Functions and demonstrates its effectiveness in seismic noise denoising, encompassing both coherent and incoherent noise. Additionally, we apply the network for interpolating missing seismic data. Importantly, all methods are based on the inherent properties of INR without requiring an additional training dataset, rendering our approach zero-shot unsupervised.

THEORY

Implicit Neural Reoresentation

Implicit Neural Representation (INR) serves as a continuous learned function approximator, relying on multilayer perceptrons (MLPs). Its continuous nature proves particularly advantageous when handling irregularly sampled signals. The primary objective of an INR is to capture the mapping from input coordinates, denoted as $\mathbf{r}$, to corresponding data, such as signal values or RGB colors. The continuous parameterization of INR facilitates storing signals at a consistent memory cost relative to spatial resolution, distinguishing INRs for reconstructing high-dimensional signals.

Both INR and classical neural network architectures face a notable challenge in their pronounced spectral bias toward lower frequencies (Rahaman et al., 2019), traditionally limiting their utility in implicit representation tasks. Recent efforts have introduced various solutions to mitigate this spectral bias in neural networks. For instance, Tancik et al. (2020) proposed incorporating a Fourier mapping layer before the MLP, while Sitzmann et al. (2020) suggested utilizing an MLP with sinusoidal activations. Both approaches aim to bias the networks toward higher frequencies. In this paper, we adopt the Sinusoidal Representation Networks (SIREN) introduced by Sitzmann et al. (2020)

INR for Seismic data processing

In the context of seismic data, INR performs a mapping from coordinate systems to the signal value space. For example, in the 2D case, the function is represented as $f(t, h) = d$, where t is time, h is offset, and d corresponds to the signal value at that specific coordinate. In the 3D case, the equation becomes $f(t, x, y) = d$, with x and y representing the two horizontal spatial dimensions. This formulation extends to a 5D case as well.

What sets INR apart is its departure from conventional training methods that heavily rely on large paired datasets. Instead, INR efficiently achieves data training by parameterizing seismic signals using only a single dataset. In our specific application, the training input comprises coordinates, while the training target corresponds to the signal values at those specific coordinates. This innovative approach streamlines the training process, showcasing the adaptability and effectiveness of INR in capturing the inherent complexities of continuous signals.

Seismic denoising

Incoherent noise case

Similar to the Deep Image Prior (DIP) method proposed by Ulyanov et al. (Ulyanov et al., 2018), which leverages U-net architecture (Ronneberger et al., 2015), the INR approach also effectively suppresses noise during early training iterations and prioritizes the extraction of self-similar features. This unique characteristic allows INR to reconstruct coherent seismic data before addressing incoherent noise. It adapts to both random and erratic noise.

For the denoising of incoherent noise, the loss function is defined simply as:

$$\min_{\theta} \|f_{\theta}(r) - x_0\|_2^2. \tag{1}$$

Where $\mathbf{r}$ is the coordinate of x_0. Alternatively, the ℓ_2 norm can be replaced with the ℓ_1 norm for erratic noise case.

In Figure 1, the Peak Signal-to-Noise Ratio (PSNR) is depicted, illustrating the comparison between the training output and the original seismic data. The blue curve represents the PSNR between the output and the original clean seismic data, while the orange curve represents the PSNR between the output and the noisy data. In Part (a), a fixed learning rate is employed. The network initially reconstructs the coherent signal, reaching a peak PSNR, after which it starts reconstructing the incoherent noise, causing a decline in PSNR. However, this training process also inhibits noise learning, making it challenging for INR to precisely learn the same input noise data, as indicated by the lower peak of the orange curve compared to the blue one. The reason behind that is the multilayer structure of INR imposes a certain low-rank structure over the coefficients. similar to the sparsity assumption (Yüce et al., 2022).

To enhance training robustness against noise, a decreasing learning rate is utilized in Part (b), starting from $1e^{-4}$ and dropping by 50% after every 100 epochs. Consequently, when the network begins reconstructing the noise, the PSNR experiences a gradual decline.

Coherent noise: Ground roll

We observe that the INR process tends to prioritize the extraction of coherent signals over incoherent noise. Furthermore, it demonstrates a sequential extraction of signals rather than an all-encompassing approach. In Figure 2, a synthetic shot gather features three linear

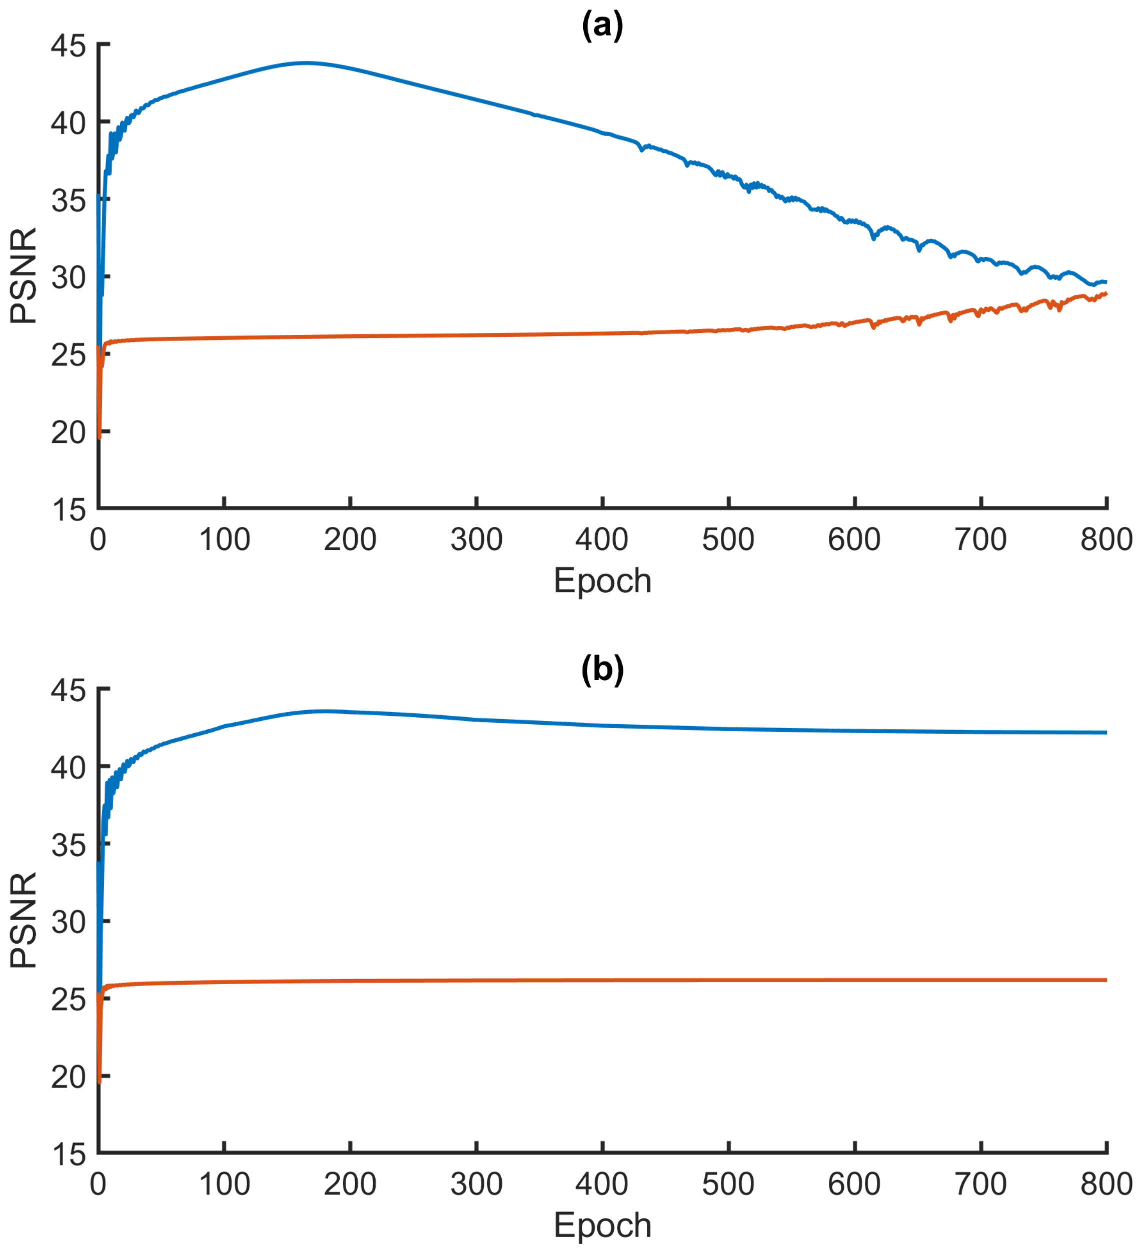

FIG. 1. The PSNR between the output and the original clean seismic data (the blue curve), and the PSNR between the output and the noisy seismic data (the orange curve).

events and two hyperbolic events. Among the linear events, two are horizontal, while one exhibits a deep dip. The hyperbolic events include one close to horizontal and another with a deep dip.

In Figure 3, the training output is depicted at intervals of 200 epochs. Notably, the network begins by reconstructing events close to horizontal, irrespective of whether they are linear or curved. Subsequently, the network systematically extracts events with deep dips with much slower rate. Leveraging this characteristic, the network can effectively eliminate coherent ground roll noise. Ground roll noise typically manifests as linear features with low velocity, resulting in deep dips in seismic gathers. To address this, we perform Normal Moveout (NMO) correction to flatten reflections. Post-NMO correction, applying the network yields immediate reconstruction of the flattened events within the initial few hundred epochs, effectively removing the low-velocity ground rolls. Once all flattened events are extracted, we apply inverse NMO to restore the original reflections.

Sesimc interpolation

The proposed network is versatile, not only for its primary purpose but also for interpolating missing traces in seismic data. During the training phase, Implicit Neural Representation (INR) imposes constraints on the output, encompassing local continuity, global consistency, and a tendency towards smoothness in the early iterations. These constraints are integral to the structure and the training process of INR. Specifically, local continuity is achieved through a continuous representation, part of the global consistency stems from the use of the sine activation function and the inclination towards smoothness results from the training approach typical of parameterized methods. These methods inherently favour achieving smooth training during the initial stages.

Based on these underlying assumptions, INR can effectively interpolate incomplete seismic data. In addressing the interpolation problem, we posit that INR represents a completed shot gather, and its optimization involves minimizing the loss function:

$$\min_{\theta} \|S f_{\theta}(c) - x_0\|_2^2. \tag{2}$$

Here, the operator S denotes the sampling operator.

EXAMPLES

In this section, we leverage the proposed network to address various seismic data processing challenges.

Denoising

Beginning with denoising, we apply the method to mitigate seismic noise in synthetic shot gathers generated through the finite difference method. Figure 4 presents an example with random noise, while Figure 5 illustrates a scenario with erratic noise. To optimize the denoising process, we employ an ℓ_2 norm loss function for random noise and an ℓ_1 norm loss function for erratic noise. The outcomes demonstrate the efficacy of the network in effectively attenuating both random and erratic noise components.

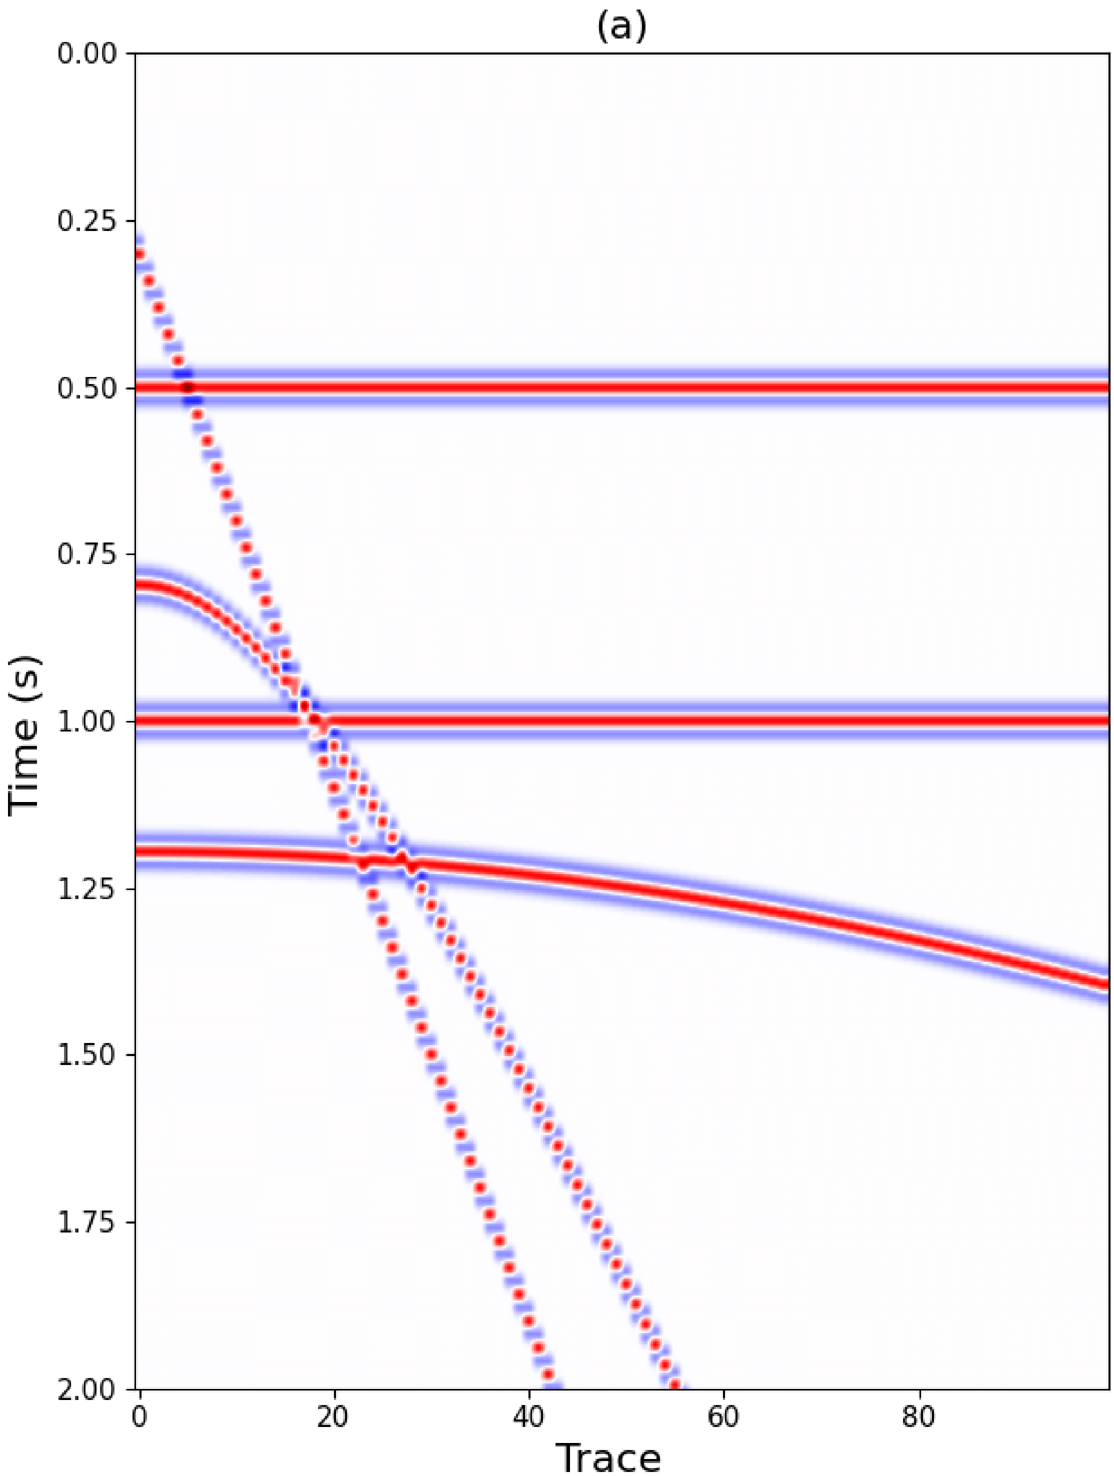

FIG. 2. A synthetic shot gather with three linear events (two horizontal and one with deep dip) and two hyperbolic events (one close to horizontal and the other has deep dip).

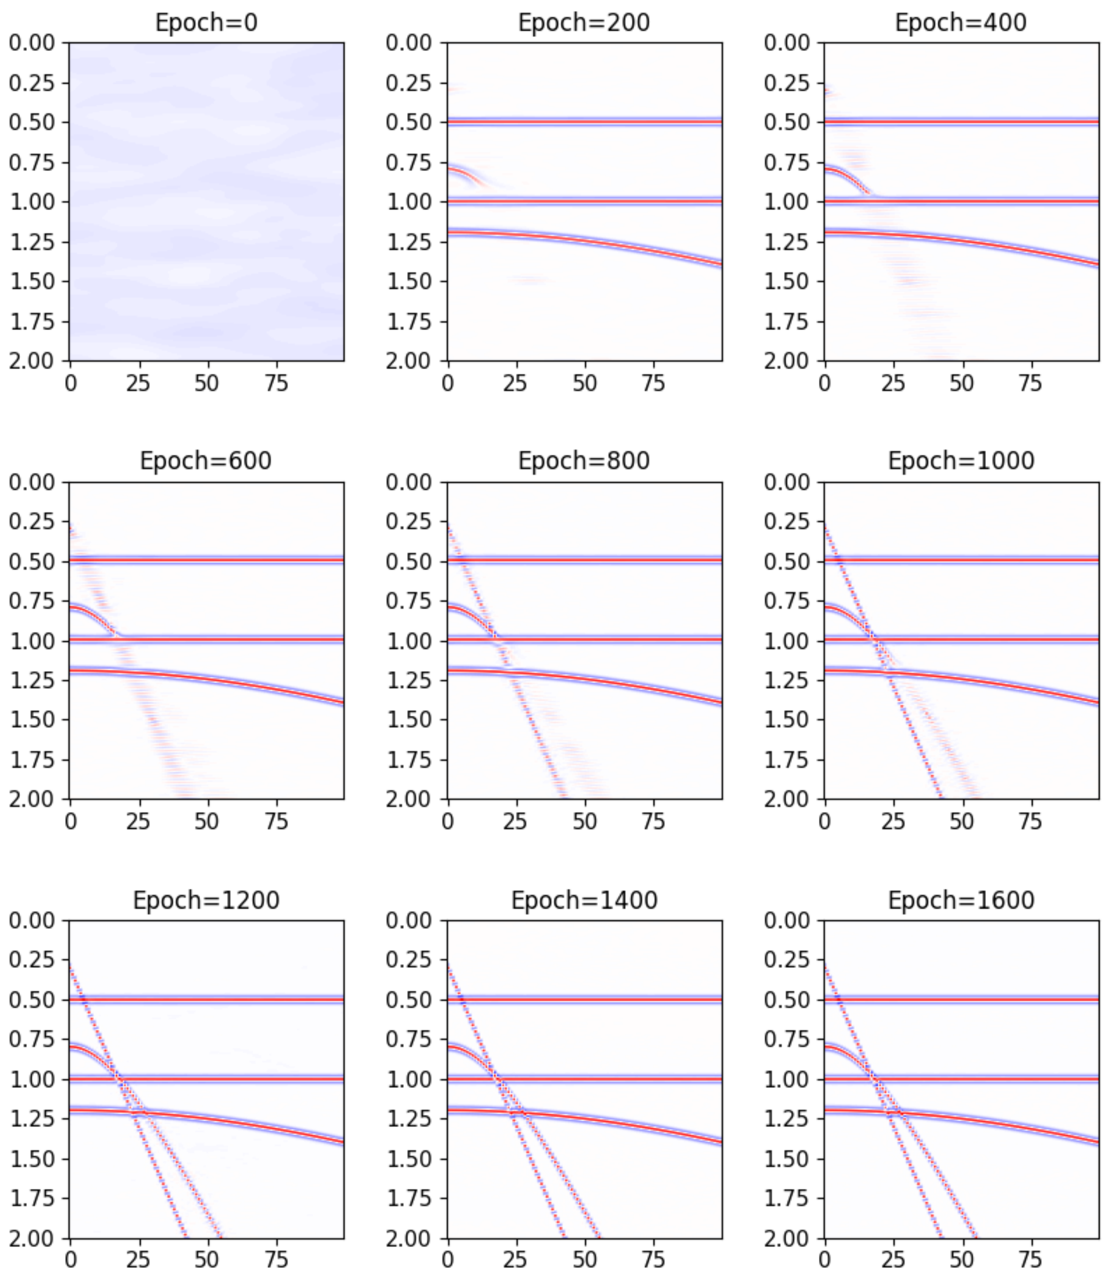

FIG. 3. Training output after every 200 epochs.

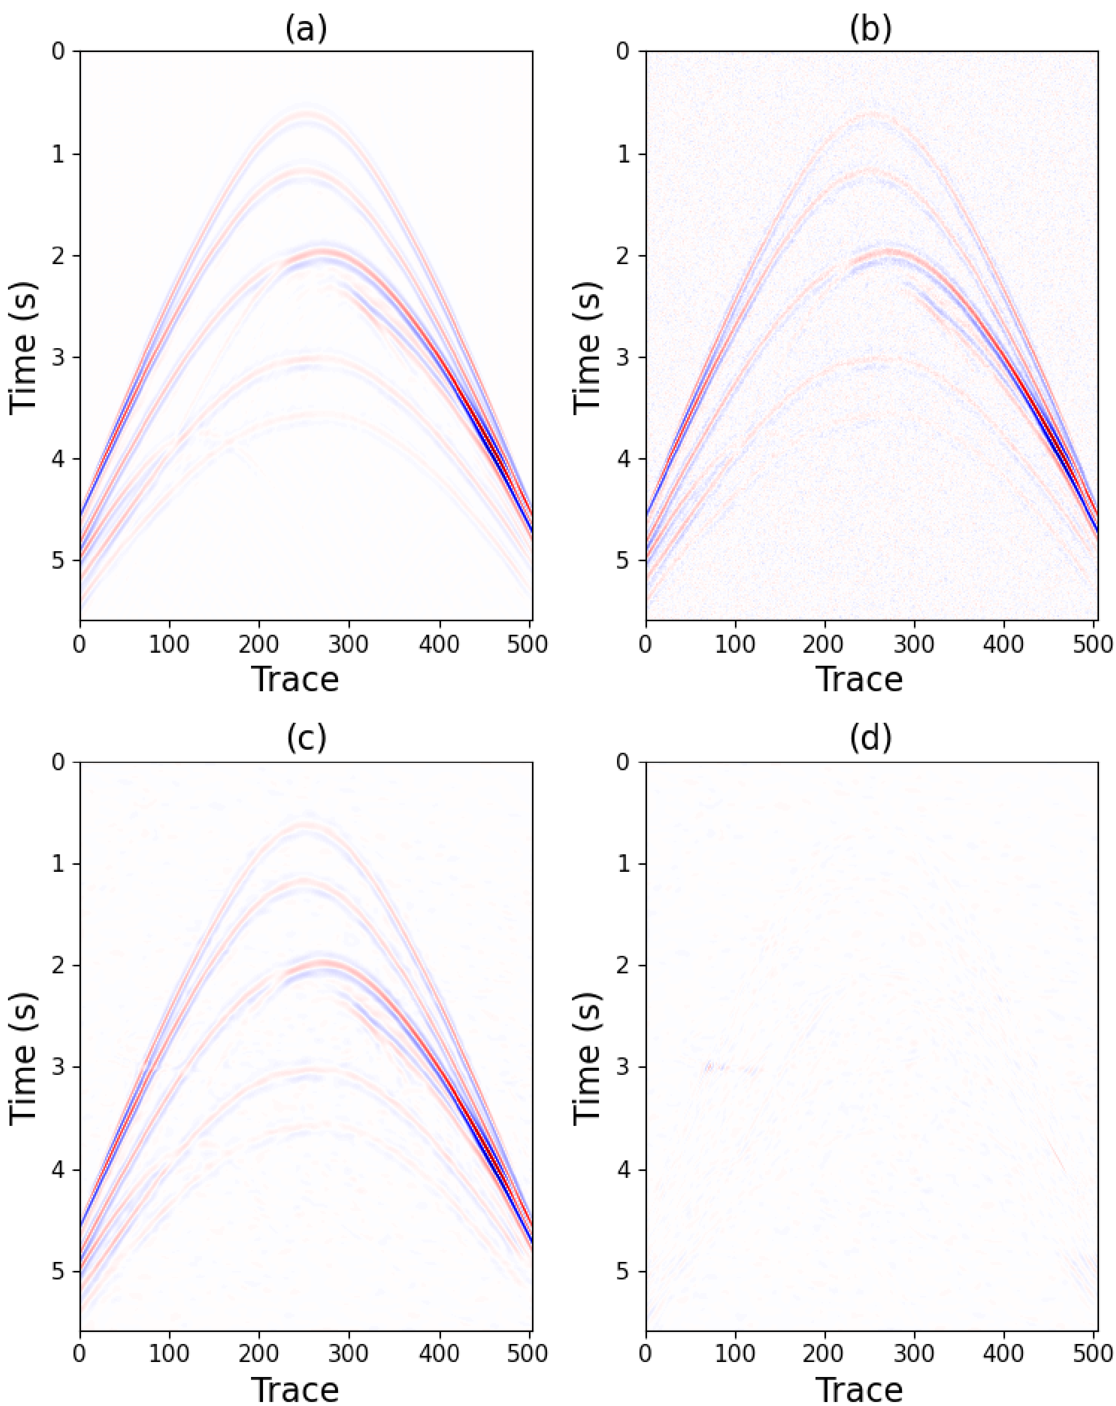

FIG. 4. Synthetic example with random noise. (a) Clean data. (b) Noisy data. (c) Denoised result. (d) Errors.

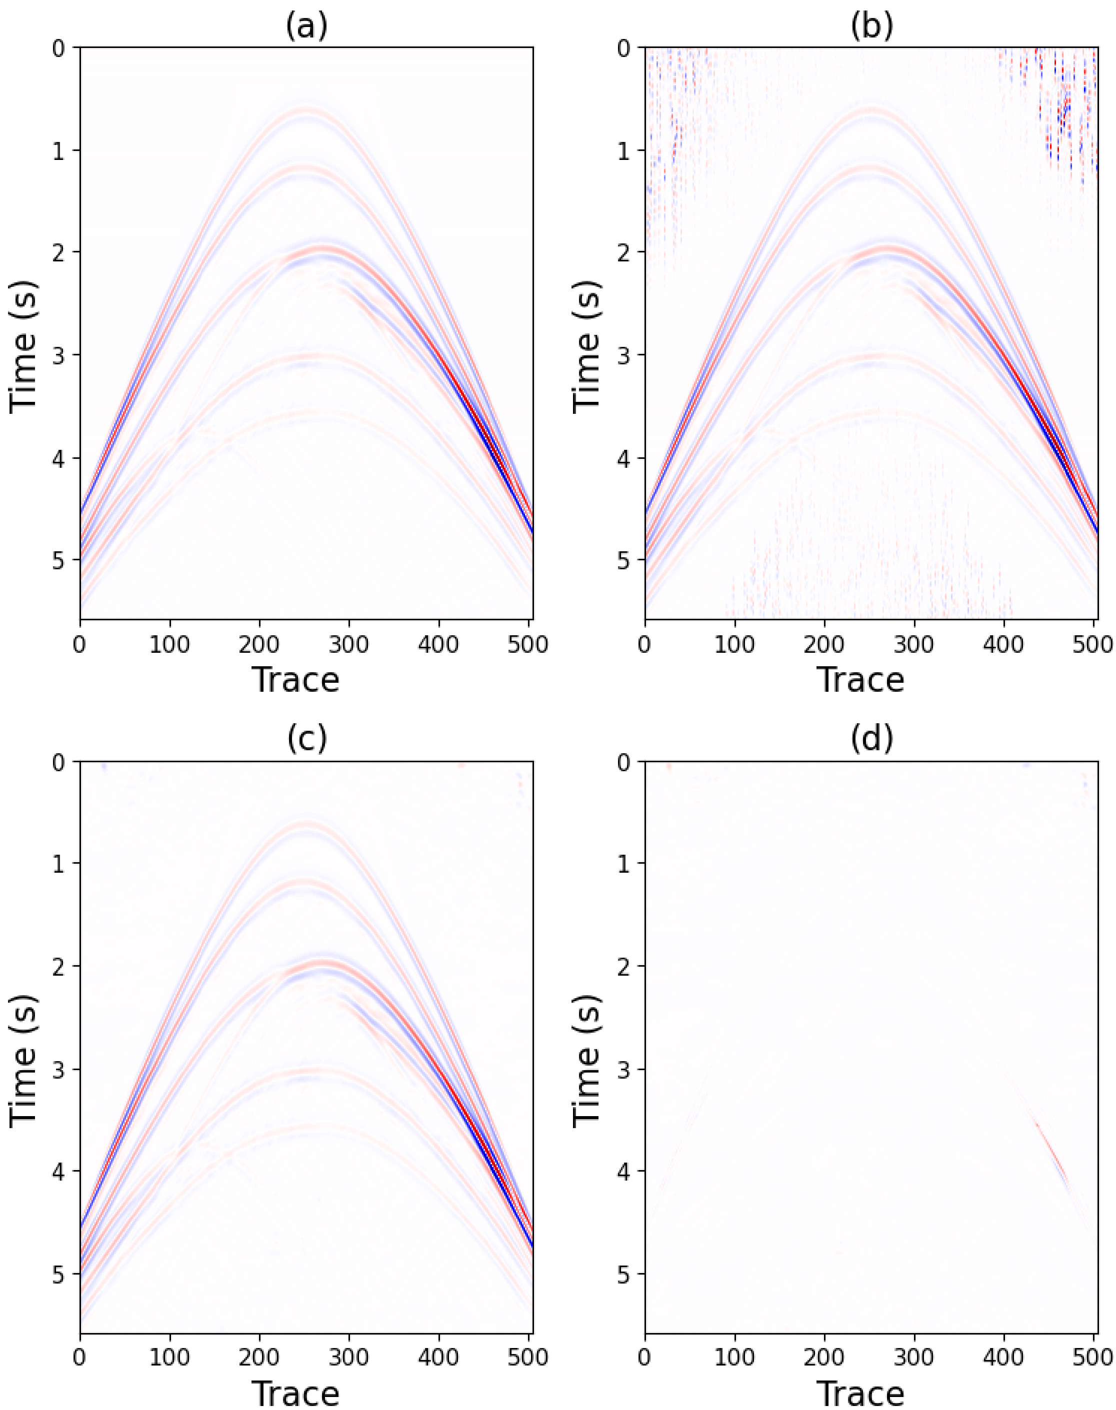

FIG. 5. Synthetic example with erratic noise. (a) Clean data. (b) Noisy data. (c) Denoised result. (d) Errors.

Interpolation

Next, we utilize the network to address the interpolation challenge within the same dataset. Initially, we randomly remove half of the traces and subsequently apply the network to solve the predefined loss function tailored for interpolation. The interpolation example is illustrated in Figure 6, providing visual evidence that the network excels in successfully interpolating the seismic data.

Ground Roll attenuation

In the end, we apply the proposed network to eliminate the ground roll from both synthetic and real land seismic data examples. The synthetic case, illustrated in Figure 7, features three reflections and ground roll. Initially, we employ NMO correction to flatten the reflections, as depicted in Figure 8(a). Subsequently, our proposed neural network extracts these flattened reflections, as shown in part (b). After successfully extracting the flattened reflections and eliminating the ground roll, we employ inverse NMO to revert the reflections to their original hyperbolic form, showcased in Figure 8(e).

This process is replicated on real data, exemplified in Figure 9, with the corresponding results presented in Figure 10. In both cases, our method effectively removes the coherent ground roll, highlighting its success in enhancing seismic data quality.

CONCLUSIONS

In this study, we reviewed the INR framework, delving into its potential applications in seismic data processing. Remarkably, the INR network exhibited the capacity to address various seismic data processing challenges without necessitating additional training data. Our investigation, focusing on seismic data denoising, interpolation, and ground roll attenuation, yielded promising results. These outcomes underscore the efficacy of the INR network in enhancing the quality of seismic data.

The implications of our findings extend beyond the immediate applications explored in this paper, paving the way for heightened research interest in this innovative direction. As a testament to the adaptability of the INR framework, we aim to extend its application to high-dimensional problems such as irregular 5D interpolation. Additionally, we plan to investigate its efficacy in addressing other coherent noise attenuation issues, including the removal of multiples. This forward-looking approach positions our research at the forefront of advancing seismic data processing methodologies.

ACKNOWLEDGMENTS

We thank the sponsors of CREWES for their continued support. This work was funded by CREWES industrial sponsors and NSERC (Natural Science and Engineering Research Council of Canada) through the grant CRDPJ 543578-19.

REFERENCES

Beck, A., and Teboulle, M., 2009, A fast iterative shrinkage-thresholding algorithm for linear inverse problems: SIAM journal on imaging sciences, **2**, No. 1, 183–202.

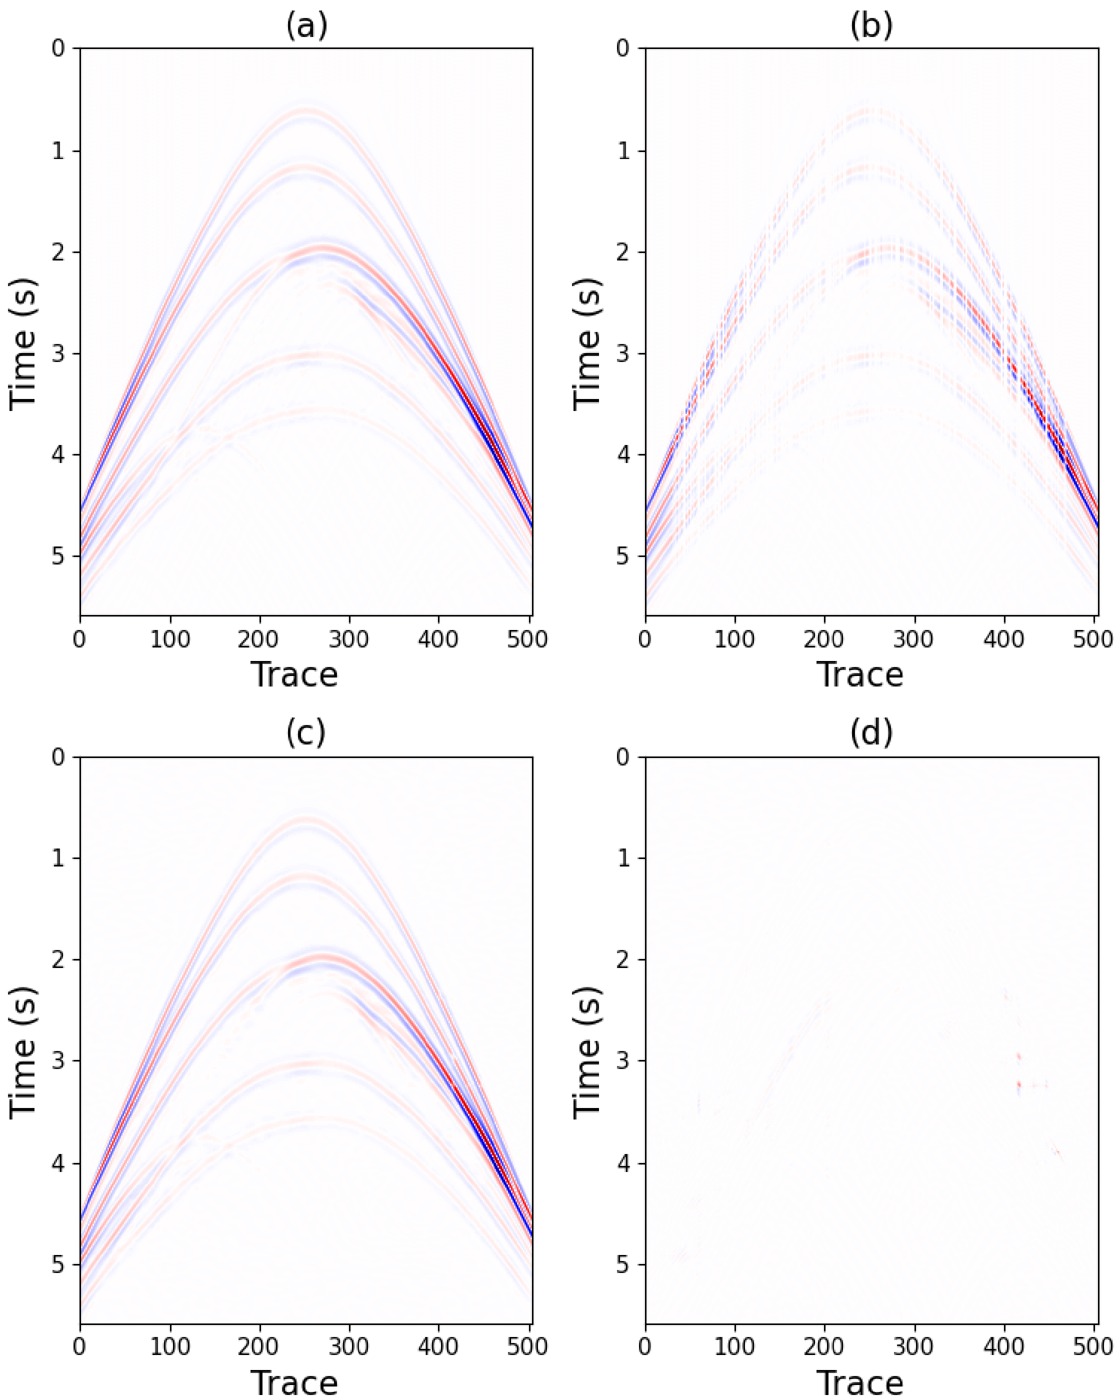

FIG. 6. Synthetic example with missing traces. (a) Completed data. (b) Decimated data. (c) Interpolated result. (d) Errors.

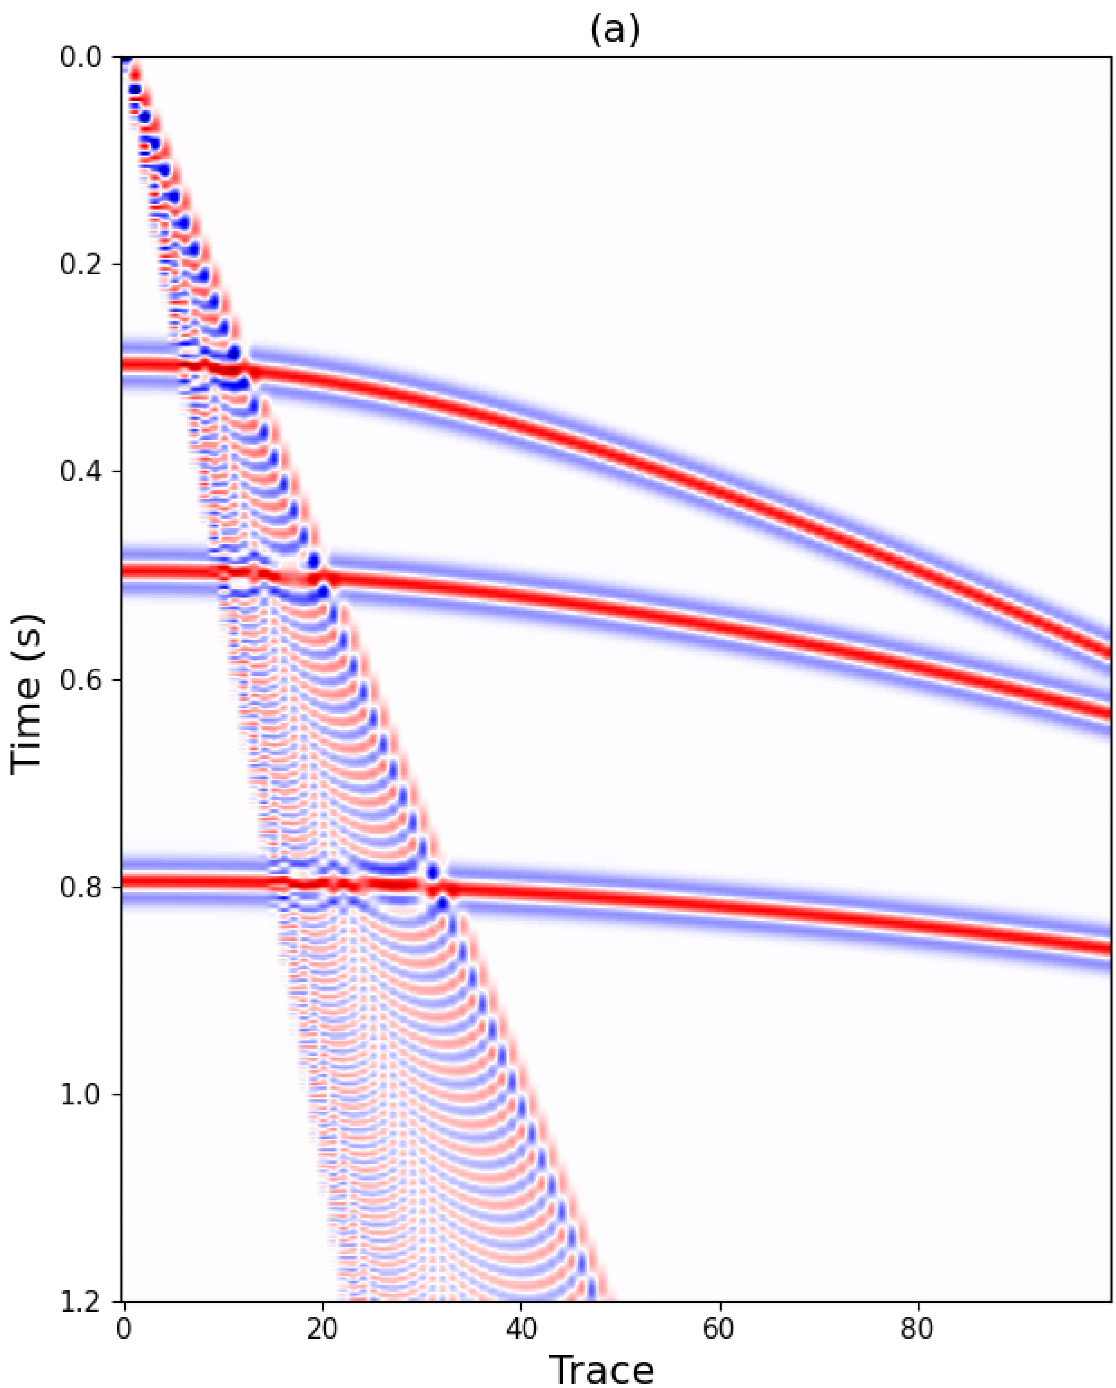

FIG. 7. Synthetic example with ground roll.

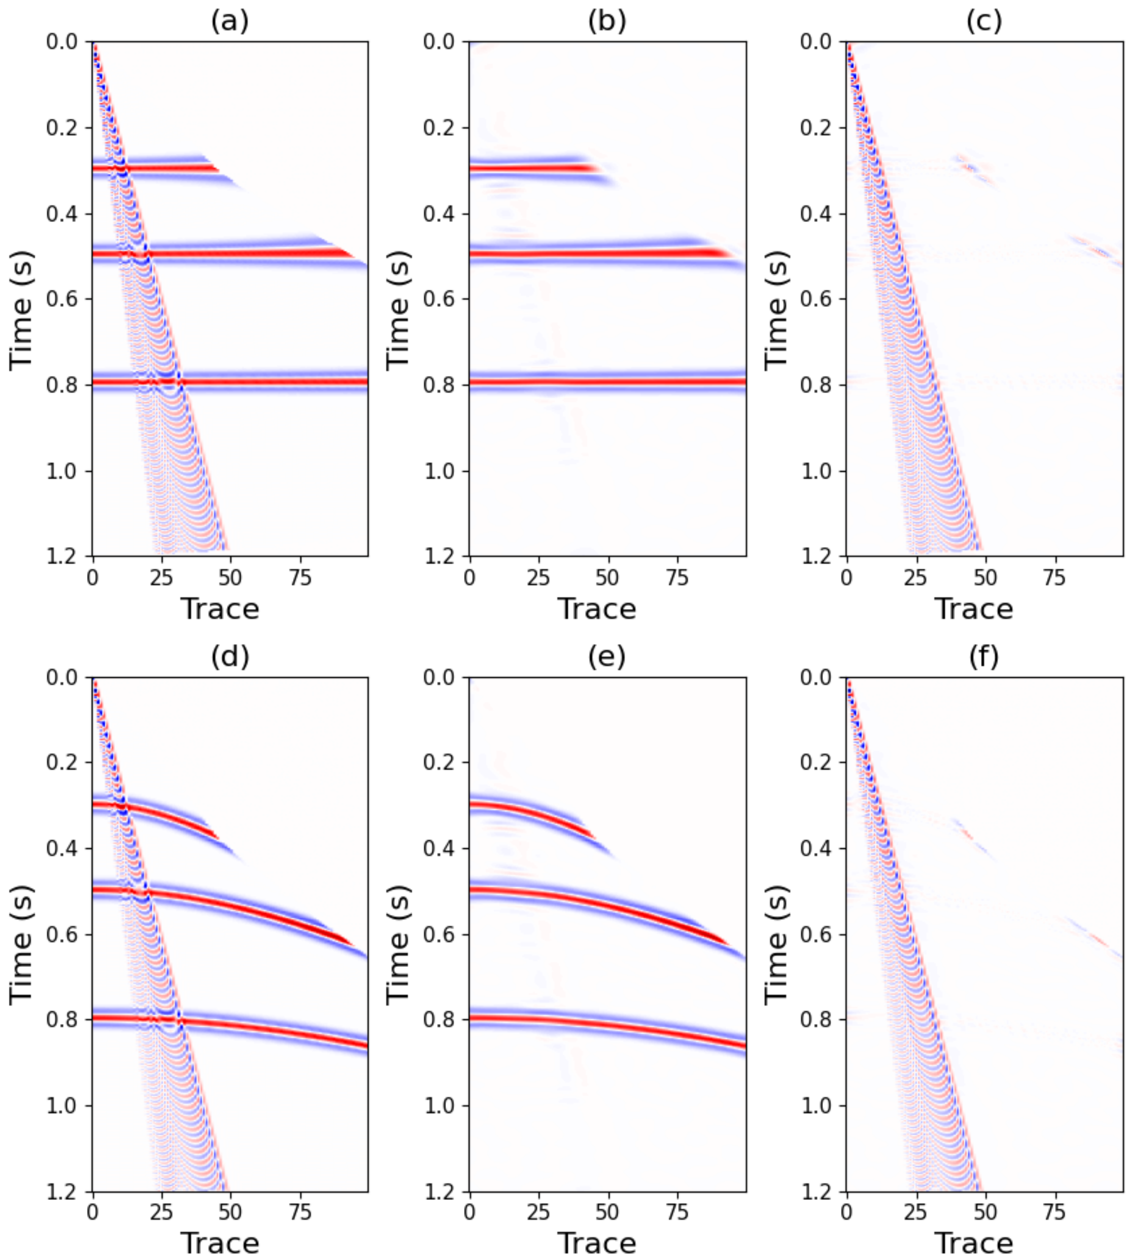

FIG. 8. Synthetic ground roll attenuation example. (a) Figure 7 after NMO correction. (b) Reconstructed result by proposed method. (c) Errors between (a) and (b). (d) Part (a) after inverse NMO correction. (e) Part (b) after inverse NMO correction. (f) Errors between (d) and (e).

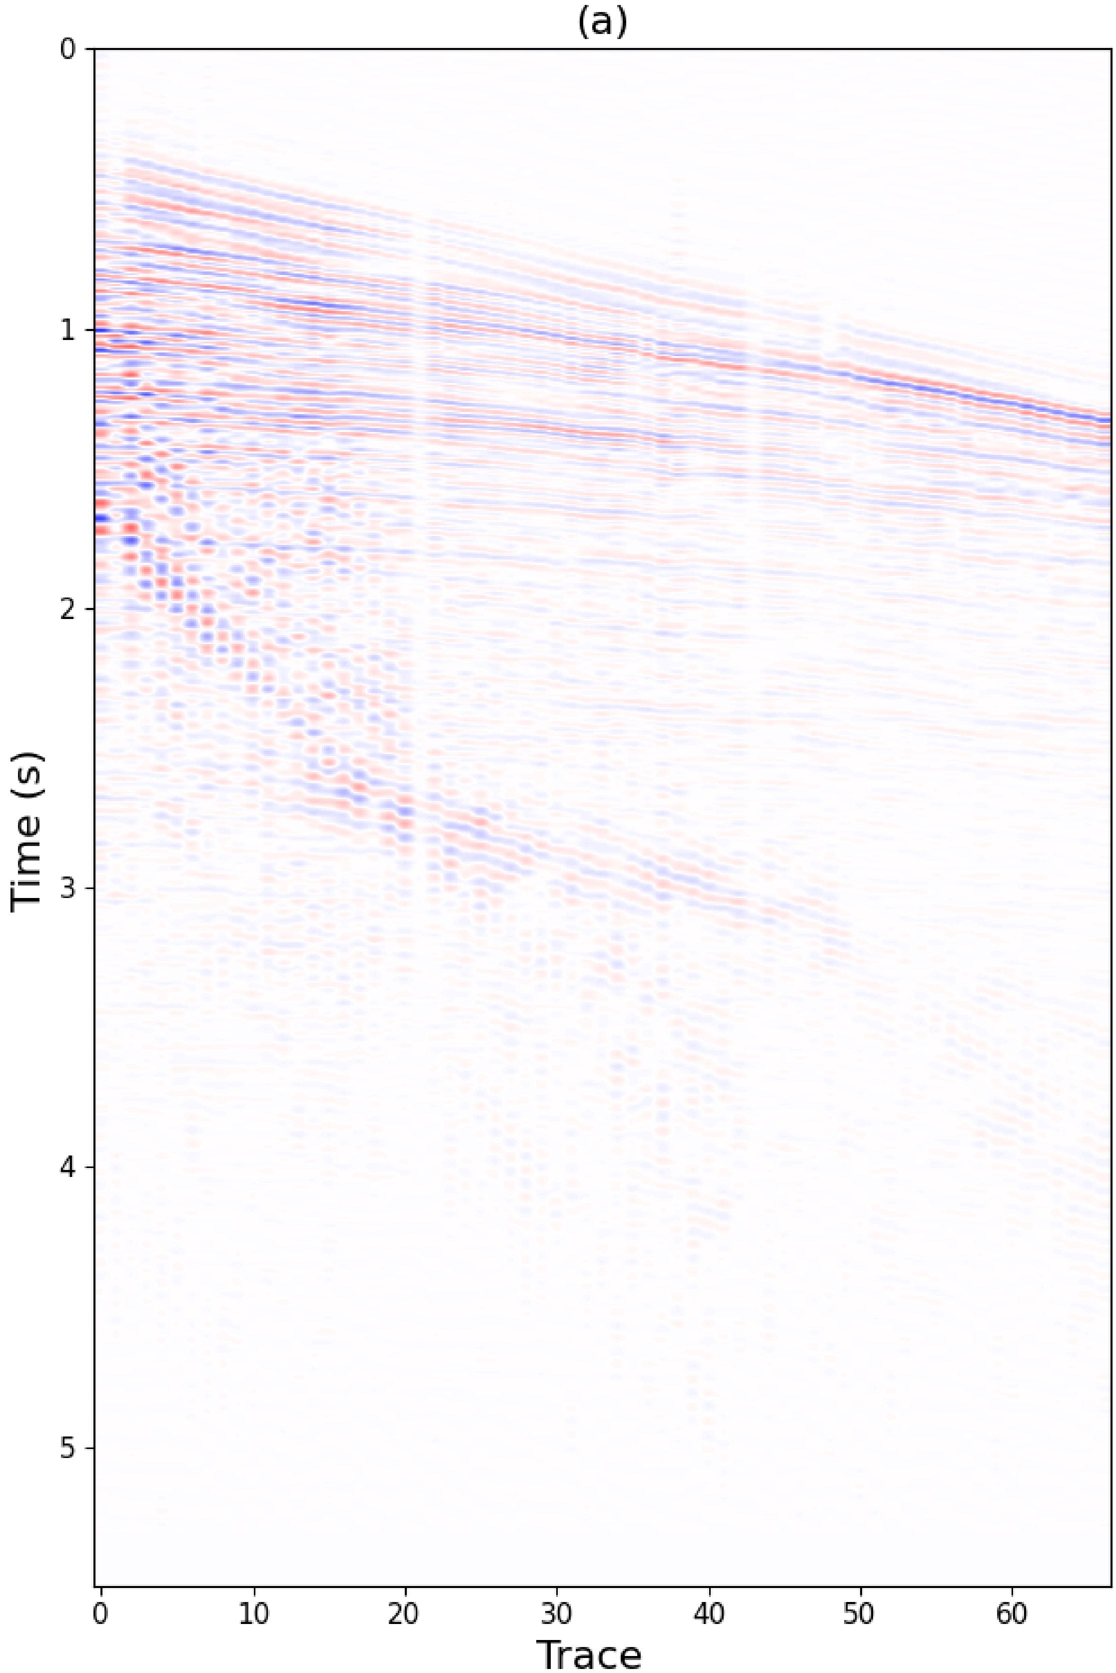

FIG. 9. Real data example with ground roll.

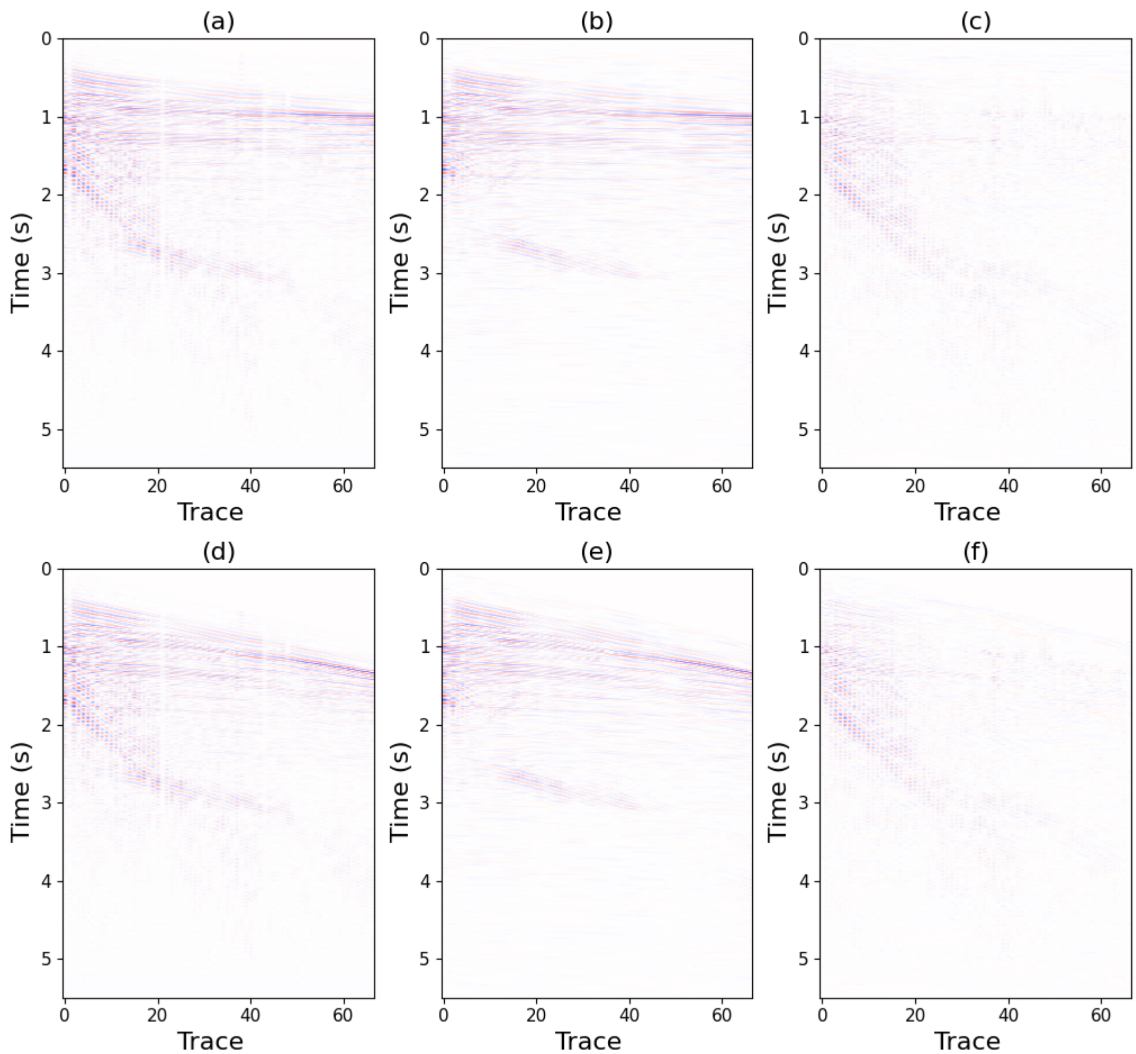

FIG. 10. Real data ground roll attenuation example. (a) Figure 9 after NMO correction. (b) Reconstructed result by proposed method. (c) Errors between (a) and (b). (d) Part (a) after inverse NMO correction. (e) Part (b) after inverse NMO correction. (f) Errors between (d) and (e).

Boyd, S., Parikh, N., Chu, E., Peleato, B., Eckstein, J. et al., 2011, Distributed optimization and statistical learning via the alternating direction method of multipliers: Foundations and Trends® in Machine learning, **3**, No. 1, 1–122.

Candes, E., Demanet, L., Donoho, D., and Ying, L., 2006, Fast discrete curvelet transforms: multiscale modeling & simulation, **5**, No. 3, 861–899.

Chan, E. R., Monteiro, M., Kellnhofer, P., Wu, J., and Wetzstein, G., 2021, pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis, *in* Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 5799–5809.

Durrani, T., and Bisset, D., 1984, The radon transform and its properties: Geophysics, **49**, No. 8, 1180–1187.

Goyes-Peñafiel, P., Vargas, E., Correa, C. V., Sun, Y., Kamilov, U. S., Wohlberg, B., and Arguello, H., 2023, Coordinate-based seismic interpolation in irregular land survey: A deep internal learning approach: IEEE Transactions on Geoscience and Remote Sensing.

Gülünay, N., 2003, Seismic trace interpolation in the fourier transform domain: Geophysics, **68**, No. 1, 355–369.

Jian, M., Jinghuai, G., and Wenchao, C., 2006, On the denoising method of prestack seismic data in wavelet domain, *in* SEG International Exposition and Annual Meeting, SEG, SEG–2006.

Kaur, H., Fomel, S., and Pham, N., 2020, Seismic ground-roll noise attenuation using deep learning: Geophysical Prospecting, **68**, No. 7, 2064–2077.

Kaur, H., Pham, N., and Fomel, S., 2021, Seismic data interpolation using deep learning with generative adversarial networks: Geophysical Prospecting, **69**, No. 2, 307–326.

Krull, A., Buchholz, T.-O., and Jug, F., 2019, Noise2void-learning denoising from single noisy images, *in* Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2129–2137.

Lehtinen, J., Munkberg, J., Hasselgren, J., Laine, S., Karras, T., Aittala, M., and Aila, T., 2018, Noise2noise: Learning image restoration without clean data: arXiv preprint arXiv:1803.04189.

Li, H., Yang, W., and Yong, X., 2018, Deep learning for ground-roll noise attenuation, *in* SEG Technical Program Expanded Abstracts 2018, Society of Exploration Geophysicists, 1981–1985.

Li, S., Liu, B., Ren, Y., Chen, Y., Yang, S., Wang, Y., and Jiang, P., 2019, Deep-learning inversion of seismic data: arXiv preprint arXiv:1901.07733.

Liu, D., Deng, Z., Wang, C., Wang, X., and Chen, W., 2020, An unsupervised deep learning method for denoising prestack random noise: IEEE Geoscience and Remote Sensing Letters, **19**, 1–5.

Liu, D., Wang, W., Chen, W., Wang, X., Zhou, Y., and Shi, Z., 2018, Random noise suppression in seismic data: What can deep learning do?, *in* SEG International Exposition and Annual Meeting, SEG, SEG–2018.

Martin-Brualla, R., Radwan, N., Sajjadi, M. S., Barron, J. T., Dosovitskiy, A., and Duckworth, D., 2021, Nerf in the wild: Neural radiance fields for unconstrained photo collections, *in* Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 7210–7219.

Mildenhall, B., Srinivasan, P. P., Tancik, M., Barron, J. T., Ramamoorthi, R., and Ng, R., 2021, Nerf: Representing scenes as neural radiance fields for view synthesis: Communications of the ACM, **65**, No. 1, 99–106.

Oliveira, D. A., Ferreira, R. S., Silva, R., and Brazil, E. V., 2018, Interpolating seismic data with conditional generative adversarial networks: IEEE Geoscience and Remote Sensing Letters, **15**, No. 12, 1952–1956.

Pang, T., Zheng, H., Quan, Y., and Ji, H., 2021, Recorrupted-to-recorrupted: Unsupervised deep learning for image denoising, *in* Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2043–2052.

Qian, F., Guo, W., Liu, Z., Yu, H., Zhang, G., and Hu, G., 2022, Unsupervised erratic seismic noise attenuation with robust deep convolutional autoencoders: IEEE Transactions on Geoscience and Remote Sensing, **60**, 1–16.

Rahaman, N., Baratin, A., Arpit, D., Draxler, F., Lin, M., Hamprecht, F., Bengio, Y., and Courville, A., 2019, On the spectral bias of neural networks, *in* International Conference on Machine Learning, PMLR, 5301–5310.

Richardson, A., and Feller, C., 2019, Seismic data denoising and deblending using deep learning: arXiv preprint arXiv:1907.01497.

Ronneberger, O., Fischer, P., and Brox, T., 2015, U-net: Convolutional networks for biomedical image segmentation, *in* Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18, Springer, 234–241.

Saad, O. M., Oboue, Y. A. S. I., Bai, M., Samy, L., Yang, L., and Chen, Y., 2021, Self-attention deep image prior network for unsupervised 3-d seismic data enhancement: IEEE Transactions on Geoscience and Remote Sensing, **60**, 1–14.

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**, 7462–7473.

Tancik, M., Srinivasan, P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J., and Ng, R., 2020, Fourier features let networks learn high frequency functions in low dimensional domains: Advances in Neural Information Processing Systems, **33**, 7537–7547.

Ulyanov, D., Vedaldi, A., and Lempitsky, V., 2018, Deep image prior, *in* Proceedings of the IEEE conference on computer vision and pattern recognition, 9446–9454.

Wang, B., Zhang, N., Lu, W., and Wang, J., 2019, Deep-learning-based seismic data interpolation: A preliminary result: Geophysics, **84**, No. 1, V11–V20.

Wang, K., and Hu, T., 2021, Deblending of seismic data based on neural network trained in the csg: IEEE Transactions on Geoscience and Remote Sensing, **60**, 1–12.

Wen, F., Liu, P., Liu, Y., Qiu, R. C., and Yu, W., 2016, Robust sparse recovery in impulsive noise via ℓ_p-ℓ_1 optimization: IEEE Transactions on Signal Processing, **65**, No. 1, 105–118.

Xu, S., Zhang, Y., Pham, D., and Lambaré, G., 2005, Antileakage fourier transform for seismic data regularization: Geophysics, **70**, No. 4, V87–V95.

Yüce, G., Ortiz-Jiménez, G., Besbinar, B., and Frossard, P., 2022, A structured dictionary perspective on implicit neural representations, *in* Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 19,228–19,238.

Zheng, Y., Zhang, Q., Yusifov, A., and Shi, Y., 2019, Applications of supervised deep learning for seismic interpretation and inversion: The Leading Edge, **38**, No. 7, 526–533.