

Article

High-Resolution Reconstruction of Seismic Data with Cycle-Consistent Adversarial Network

Si-Yi Chen ¹  and Ming Yang ^{1,2,3,*} ¹ College of Aviation Electronics and Electrical Engineering, Civil Aviation Flight University of China, Guanghan 618307, China; 18181793813@163.com² School of Optoelectronic Science and Engineering, University of Electronic Science and Technology of China, Chengdu 610054, China³ Sichuan Province Engineering Technology Research Center of General Aircraft Maintenance, Civil Aviation Flight University of China, Guanghan 618307, China

* Correspondence: yangming932@163.com

Abstract

High-resolution seismic reconstruction is a challenging inverse problem because field seismic traces are inherently band-limited and their high-frequency components are further degraded by source bandwidth limitations, acquisition conditions, random noise, and attenuation during wave propagation. Classical resolution enhancement methods can partially sharpen seismic events, but they usually rely on restrictive assumptions about stationarity, minimum-phase wavelets, or accurate attenuation models. In this study, we propose a structure-preserving bidirectional bandwidth translation network for seismic resolution enhancement. Instead of formulating the task as a one-way paired regression problem, the proposed approach interprets resolution enhancement as unpaired translation between low-bandwidth and high-bandwidth seismic domains. A cycle-consistent adversarial objective is combined with an SSIM-based structural constraint so that the model simultaneously improves spectral recovery, waveform fidelity, and reflector continuity. To reduce the domain gap between synthetic and field data, we further construct a hybrid training corpus by combining field-extracted wavelets with synthetic reflectivity sequences and train a lightweight one-dimensional residual generator–discriminator architecture tailored to oscillatory seismic traces. Comprehensive experiments are conducted on synthetic data, a field seismic profile, and the public SEG Open Data benchmark. In addition to comparisons with conventional deconvolution and time-varying frequency deconvolution, the manuscript reports quantitative comparisons with representative learning-based baselines, together with ablation studies, parameter sensitivity analysis, robustness evaluation under different noise levels and bandwidth settings, and computational cost analysis. The results show that the proposed method consistently achieves a favorable balance between spectral extension and structural preservation, demonstrating its potential as a practical data-driven solution for seismic resolution enhancement.



Academic Editor: Cristina Sáez Blázquez

Received: 5 March 2026

Revised: 25 May 2026

Accepted: 27 May 2026

Published: 1 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Keywords: seismic data; resolution enhancement; unpaired domain translation; structural preservation; SEG open data

1. Introduction

High-resolution seismic processing is essential for identifying thin beds, stratigraphic pinch-outs, subtle faults, and weak discontinuities. In practice, however, recorded seismic traces are band-limited and their high-frequency components are further weakened by

source bandwidth, acquisition geometry, environmental noise, and propagation attenuation associated with the Q factor. These effects reduce the effective frequency bandwidth of the wavelet and therefore degrade the vertical resolution of seismic data.

Over the past several decades, a variety of methods have been developed to enhance seismic resolution. Robinson introduced the classical convolution model [1], in which a seismic trace is represented as the convolution of a seismic wavelet with a reflection coefficient sequence. Based on this model, predictive deconvolution, adaptive deconvolution, Gabor deconvolution, and related time-varying spectral compensation methods were proposed to recover high-frequency components from band-limited observations [2–4]. Although these methods can sharpen seismic wavelets to some extent, they often rely on restrictive assumptions such as stationary or minimum-phase wavelets and white reflectivity, which are difficult to satisfy in complex geological environments.

To compensate for attenuation and phase distortion during wave propagation, inverse- Q filtering and related attenuation compensation strategies have also been extensively investigated [5–10]. These methods can partially restore high-frequency content, especially in deep formations, but their performance is sensitive to noise and to errors in the attenuation model. As a result, strong high-frequency amplification may introduce unstable amplitudes or reconstruction artifacts.

In recent years, deep learning has provided an alternative data-driven route for seismic resolution enhancement because neural networks can learn nonlinear mappings directly from large datasets rather than depending completely on explicit signal assumptions [11,12]. Optimization-inspired high-resolution inversion networks [13], structure-constrained reconstruction models [14], attenuation-aware vertical resolution enhancement strategies [15], feature-aware deep enhancement models [16], pseudo-well-driven adversarial approaches [17], self-supervised blind deconvolution [18], self-supervised frequency extension frameworks [19], and non-local similarity regularization [20] have all shown that data-driven models can recover missing high-frequency information more effectively than conventional handcrafted filters in many scenarios.

Nevertheless, three limitations remain common in the current literature. First, many supervised methods require paired low-resolution (LR) and high-resolution (HR) labels generated from wells, synthetic models, or pseudo labels, which limits adaptability to field data. Second, one-way regression models often emphasize local trace sharpening but do not explicitly constrain whether the reconstructed HR signal remains consistent with the original LR observation after reverse mapping. Third, spectral enhancement may come at the expense of structural continuity when the loss function does not directly penalize waveform distortion and reflector inconsistency.

To address these issues, this study formulates seismic resolution enhancement as a structure-preserving bidirectional translation problem between a low-bandwidth seismic domain and a high-bandwidth seismic domain. From this perspective, the proposed method should be viewed not simply as “another GAN model”, but as a bidirectional bandwidth translation framework in which cycle-consistent adversarial learning is used as the optimization mechanism. The forward mapping is responsible for recovering missing high-frequency content, whereas the reverse mapping acts as a regularizer that constrains the enhanced signal to remain physically compatible with the original low-bandwidth observation. In addition, an SSIM-based structural term is embedded in the objective function to preserve waveform morphology and reflector geometry during spectral extension.

Accordingly, the contributions of this paper can be summarized in three aspects. First, we propose a structure-preserving bidirectional bandwidth translation framework for seismic resolution enhancement, where LR-to-HR enhancement and HR-to-LR back-projection are jointly optimized to constrain spectral recovery with cycle-consistent structural regular-

ization. Second, we design a hybrid training strategy that combines field-extracted wavelets, synthetic reflectivity construction, and a lightweight one-dimensional residual backbone, thereby improving the compatibility of the learned model with oscillatory seismic traces and real bandwidth characteristics. Third, we present a comprehensive experimental study including comparisons with conventional and recent learning-based methods, public benchmark evaluation on SEG Open Data, and extensive ablation, sensitivity, robustness, and efficiency analyses.

The remainder of this paper is organized as follows. Section 2 presents the problem formulation, domain-construction procedure, and proposed network. Section 3 describes the datasets, implementation details, and comparison results. Section 4 reports ablation experiments and extended analyses. Finally, Section 5 and Section 6 provide the discussion and conclusion, respectively.

2. Method

2.1. Training Domain Construction

Many deep learning methods fail to generalize well when the amount of training data is insufficient. In this study, the training data are generated in a controlled yet geologically informed manner so that the model can observe sufficient bandwidth variations during optimization. Figure 1 shows the basic principle of constructing synthetic seismic signals from wavelets and reflection coefficients.

Because the composition of field seismic data is much more complex than that of idealized synthetic examples, a pure Ricker wavelet assumption is insufficient for practical training. We therefore first extract representative wavelets from field seismic data and then use these wavelets to construct two bandwidth-controlled domains. The field-extracted wavelets are selected from laterally continuous windows with relatively high signal-to-noise ratios and stable phase characteristics. Wavelets dominated by isolated noise bursts, strong acquisition footprint, or visibly unstable phase are excluded before normalization. The remaining wavelets are grouped according to their effective frequency support, so that the domain separation is controlled by bandwidth rather than by the spatial position of the original field traces. Wavelets with effective frequency content in 10–60 Hz are used to generate the LR domain, whereas wavelets with effective frequency content in 40–80 Hz are used to generate the HR domain. For each domain, random sparse reflection coefficient sequences with varying amplitudes, layer thicknesses, and event spacing are generated independently and convolved with wavelets drawn from the corresponding wavelet pool. The LR and HR traces are therefore not paired sample by sample: a trace in the LR domain and a trace in the HR domain do not share the same reflectivity realization during unpaired adversarial training. In this way, the LR and HR corpora are statistically coupled through a shared geological prior, but remain unpaired at the trace level. Mild amplitude perturbations and additive noise are further introduced during training to improve robustness.

For all experiments, each trace contains 800 samples. The synthetic database contains 400,000 traces in total and is divided into 320,000, 40,000, and 40,000 traces for training, validation, and testing, respectively. The validation subset is used only for model selection and hyperparameter checking, while the held-out test subset is used only for the final quantitative reports. To avoid leakage, traces generated from the same random reflectivity seed are assigned to a single subset. Before being fed into the network, each trace is normalized to $[-1, 1]$. This design enables the model to learn the mapping between bandwidth-limited and bandwidth-extended seismic domains while retaining the waveform and amplitude characteristics observed in real seismic data. It also keeps the training protocol consistent

with the intended field application, where a clean HR label is generally unavailable for each observed LR trace.




Figure 1. Construction of the synthetic training dataset. (a) Ricker wavelets with dominant frequencies of 20 Hz and 40 Hz; (b) reflection coefficient sequence; (c) generated low-resolution (LR) and high-resolution (HR) seismic signals. The LR and HR domains are generated from separated field-extracted wavelet pools and independently sampled reflectivity sequences, so the training traces are statistically related but unpaired.

2.2. Problem Statement and Degradation Model

Let $x_H \in \mathbb{R}^T$ denote an ideal high-resolution seismic trace with sampling length T , and let $x_L \in \mathbb{R}^T$ denote the corresponding low-resolution observation. In a simplified form, the degradation from x_H to x_L can be described as the combined effect of bandwidth limitation, attenuation, acquisition footprint, and noise contamination:

$$x_L = \mathcal{T}(x_H) = \mathcal{M}(h * x_H) + n, \quad (1)$$

where h is a band-limiting kernel, $*$ denotes convolution, $\mathcal{M}(\cdot)$ represents the effective acquisition-and-propagation operator, and n denotes additive noise. In practice, the exact operator $\mathcal{T}$ is unknown, non-stationary, and often varies with geological setting, depth, and acquisition geometry. Therefore, seismic resolution enhancement is better interpreted as an ill-posed inverse problem than as a deterministic deblurring task.

Given an observed LR trace x_L , the purpose of seismic resolution enhancement is to estimate an HR trace $\hat{x}_H$ that satisfies three practical requirements simultaneously: (1) it should recover additional high-frequency information, (2) it should remain structurally consistent with the input trace, and (3) it should avoid introducing unstable oscillations or artificial reflector discontinuities. These requirements motivate the use of a coupled objective rather than a single pixel-wise or sample-wise regression loss.

2.3. Bandwidth-Domain Formulation

The proposed method reformulates seismic resolution enhancement as translation between two related but unpaired signal domains rather than direct point-to-point regression. The LR domain $\{LR_{real}\}$ contains seismic traces whose dominant frequency and absolute bandwidth are restricted by attenuation and acquisition conditions. The HR domain $\{HR_{real}\}$ contains traces with richer high-frequency content and sharper interference patterns. Here, the subscript “real” denotes the samples drawn from the constructed LR or HR training domain used by the discriminator; it does not imply that paired field HR labels are available. The objective of the forward generator G_{LtoH} is to recover the missing high-frequency content, whereas the reverse generator G_{HtoL} projects the enhanced result back to the LR domain and therefore acts as a consistency regularizer.

This bidirectional design is important for seismic data. If only the forward mapping is optimized, the network may overemphasize local oscillations and produce visually sharper but physically inconsistent traces. By enforcing that an enhanced trace should be convertible back to the original LR observation, the proposed framework suppresses artificial spectral gains that are unsupported by the input data. Therefore, the reverse mapping is not merely an auxiliary branch but a key mechanism for stabilizing bandwidth extension.

From an optimization perspective, the LR domain and the HR domain should not be regarded as two unrelated data manifolds. They are coupled through the same subsurface reflectivity, but differ in effective bandwidth and interference pattern. The proposed model exploits this property by learning a domain-level transformation rather than a fixed deterministic filter. This distinction is important because the same local waveform in the LR domain may correspond to different HR outcomes under different geological contexts. Domain translation allows the model to account for such ambiguity in a data-driven manner.

2.4. Structure-Preserving Bidirectional Translation Network

Generative models play an important role in image translation, synthesis, and reconstruction tasks. Goodfellow et al. introduced an adversarial training strategy for generative modeling and proposed the generative adversarial network (GAN) framework [21]. In GAN, two neural networks, namely a generator and a discriminator, are trained simultaneously in an adversarial manner. The generator aims to produce synthetic samples that resemble the real data distribution, while the discriminator acts as a binary classifier to distinguish generated samples from real ones. Through this adversarial learning process, the two networks are optimized alternately in a minimax game until an equilibrium is reached, where the generated samples become increasingly indistinguishable from real data.

It should be noted that the adversarial architecture does not provide a closed-form analytical equation for all nonlinear dependencies in seismic traces. Instead, the generator represents a learnable nonlinear operator composed of convolutional filters, nonlinear activations, residual mappings, and normalization layers. These operations approximate nonlinear bandwidth transformation from data, while the discriminator supplies a distribution-level constraint that penalizes outputs inconsistent with the target bandwidth domain. Cycle consistency further restricts this nonlinear mapping by requiring the translated signal

to be mapped back to its original domain, thereby reducing the admissible solution space and making the learned relationship more physically plausible.

CycleGAN extends the GAN framework by introducing a cycle consistency constraint, which enables bidirectional mapping between two domains without requiring paired training samples [22]. This framework is particularly suitable for translation between correlated data domains. In seismic signal processing, resolution enhancement can be naturally interpreted as translation between LR and HR seismic domains. Figure 2 illustrates the overall architecture of the proposed model.




Figure 2. Overall framework of the proposed bidirectional translation network. The notations LR and HR denote low- and high-bandwidth seismic domains, respectively. The dark and light modules distinguish the LR-to-HR and HR-to-LR translation branches, respectively. The arrows indicate the data flow among the real/generated domains, generators, discriminators, and cycle-reconstruction paths, while the dashed box denotes the cycle architecture used to preserve waveform consistency.

The generator G_{LtoH} learns a mapping from the LR domain $\{LR_{real}\}$ to the HR domain $\{HR_{real}\}$. The corresponding adversarial loss is defined as

$$\mathcal{L}_{GAN}(G_{LtoH}, D_H) = \mathbb{E}_{HR_{real}}[\log D_H(HR_{real})] + \mathbb{E}_{LR_{real}}[\log(1 - D_H(G_{LtoH}(LR_{real})))] \quad (2)$$

Similarly, the generator G_{HtoL} performs the reverse mapping from the HR domain to the LR domain, and its adversarial loss is defined as

$$\mathcal{L}_{GAN}(G_{HtoL}, D_L) = \mathbb{E}_{LR_{real}}[\log D_L(LR_{real})] + \mathbb{E}_{HR_{real}}[\log(1 - D_L(G_{HtoL}(HR_{real})))] \quad (3)$$

In addition to the adversarial loss, CycleGAN introduces a cycle consistency constraint to ensure that the translated signal can be mapped back to the original domain. The cycle consistency loss is defined as

$$L(G_{cycle}) = \mathbb{E}_{LR_{real}}[\|G_{HtoL}(G_{LtoH}(LR_{real})) - LR_{real}\|_1] + \mathbb{E}_{HR_{real}}[\|G_{LtoH}(G_{HtoL}(HR_{real})) - HR_{real}\|_1] \quad (4)$$

To further preserve structural information in reconstructed seismic traces, the structural similarity index (SSIM) [23] is incorporated into the loss function. Because the proposed framework is trained on unpaired LR and HR domains, the SSIM term is applied

to cycle-reconstructed traces rather than to unmatched cross-domain samples. In this way, structural preservation is enforced on sample-consistent signal pairs while the network remains fully unpaired. SSIM measures the structural similarity between two signals and is defined as

$$SSIM(x, y) = \frac{(2\mu_x\mu_y + C_1)(2\sigma_{xy} + C_2)}{(\mu_x^2 + \mu_y^2 + C_1)(\sigma_x^2 + \sigma_y^2 + C_2)}, \quad (5)$$

where μ_x and μ_y denote the mean values of signals x and y , σ_x^2 and σ_y^2 denote the corresponding variances, and σ_{xy} denotes the covariance between the two signals. C_1 and C_2 are small constants introduced to avoid numerical instability. The theoretical range of SSIM is from -1 to 1 when negative covariance is possible, while in most normalized reconstruction evaluations it is interpreted on the interval from 0 to 1 . A value of 1 indicates identical structural information between two signals. Values closer to 1 represent higher structural similarity, values close to 0 indicate weak structural agreement, and negative values, if they occur, indicate opposite or inverted local structural correspondence.

For an LR input trace, let $LR_{cyc} = G_{HtoL}(G_{LtoH}(LR_{real}))$ denote the cycle-reconstructed signal. The corresponding structural similarity is computed as

$$S_L = SSIM(LR_{cyc}, LR_{real}). \quad (6)$$

Accordingly, the structural loss associated with the LR cycle is defined as

$$L(G_{SSIM-L}) = \mathbb{E}(1 - S_L). \quad (7)$$

Similarly, for an HR input trace, let $HR_{cyc} = G_{LtoH}(G_{HtoL}(HR_{real}))$. We then compute $S_H = SSIM(HR_{cyc}, HR_{real})$, which leads to the structural loss $L(G_{SSIM-H}) = \mathbb{E}(1 - S_H)$.

Finally, the overall optimization objective of the network can be written as

$$\begin{aligned} L(G_{LtoH}, G_{HtoL}, D_H, D_L) = & \alpha L(G_{cycle}) + \mathcal{L}_{GAN}(G_{LtoH}, D_H) + \mathcal{L}_{GAN}(G_{HtoL}, D_L) \\ & + \beta (L(G_{SSIM-L}) + L(G_{SSIM-H})), \end{aligned} \quad (8)$$

where α and β are two trade-off coefficients, which are set to 5 and 10 in this paper. From the viewpoint of seismic resolution enhancement, the adversarial losses encourage domain-level spectral realism, the cycle consistency term constrains recoverability, and the SSIM term preserves waveform morphology and reflector geometry during cycle reconstruction. Their combination forms a coupled objective for bandwidth extension and structural preservation without requiring paired HR labels.

2.5. Structure of Generator and Discriminator

The two generators adopt a symmetric architecture with identical network structures, as illustrated in Figure 3. Each generator operates directly on one-dimensional seismic traces instead of generic two-dimensional image patches. The network begins with a 7×1 convolution to capture wavelet-scale context, then uses several downsampling convolutional layers to enlarge the receptive field, and finally employs four residual blocks in the bottleneck to improve feature propagation and stabilize training [24]. The decoder mirrors the encoder and reconstructs a trace with the same temporal length. This design preserves temporal phase relationships while keeping the parameter count manageable.

The discriminator serves as a relatively shallow authenticity evaluator rather than a complex reconstruction module. This asymmetry is appropriate because the generator must recover fine waveform details, whereas the discriminator only needs to determine whether the translated traces follow the HR or LR domain distributions. We adopt Layer

Normalization (LN) as the normalization strategy in our network because it normalizes activations using per-sample statistics computed across the features within each layer and thus helps stabilize training [25]. In addition, Dropout is employed to regularize the network by randomly deactivating a fraction of neurons during training, thereby reducing co-adaptation and alleviating overfitting [26]. In our implementation, the dropout rate is set to 0.2.




Figure 3. Generator and discriminator architectures. The colored blocks denote different network layers, including convolution, leaky ReLU activation, normalization, dropout, and upsampling, as indicated by the legend. The arrows indicate the forward data flow, and the “add” connections denote skip connections in the generator. The convolutional kernels operate along the temporal/sample axis of each one-dimensional seismic trace.

Overall, the methodological novelty of the proposed framework lies in four aspects: unpaired bandwidth-domain construction from field-guided wavelets, bidirectional spectral translation, SSIM-guided structural preservation, and a lightweight one-dimensional residual backbone tailored to oscillatory seismic traces.

2.6. Training and Inference Strategy

The network is trained by alternately updating the generators and discriminators. In each iteration, an LR mini-batch and an HR mini-batch are independently sampled from their corresponding domains. The forward generator G_{LtoH} first predicts bandwidth-enhanced traces, and the backward generator G_{HtoL} subsequently maps both translated outputs back to the original domain. The adversarial losses enforce domain realism, the cycle loss constrains reversibility, and the SSIM loss stabilizes waveform structure. After each epoch, the model is evaluated on the validation subset using the cycle loss, SSIM, and spectral bandwidth indicators; the checkpoint with the best overall validation behavior is retained for testing. In this way, the optimization objective acts simultaneously on spectral realism, inter-domain consistency, and structural preservation.

The risk of overfitting to noise and local artifacts is reduced by several coupled mechanisms. The discriminator encourages domain-level spectral realism rather than memorization of individual traces, the cycle loss rejects high-frequency details that cannot be mapped back to the LR domain, and the SSIM term discourages structural distortion. Together with dropout, layer normalization, label flipping, validation-based checkpoint selection, and independent LR and HR mini-batch sampling, these constraints balance fitting ability and generalization and improve resilience to changes in input noise level and effective bandwidth.

During inference, only the forward generator G_{LtoH} is required. This means that the proposed approach does not rely on any additional iterative post-processing, auxiliary inversion stage, or explicit wavelet estimation after training. The input field trace is normalized using the same rule as the training data, passed once through the forward

generator, and then re-scaled to its original amplitude level. This one-pass inference mode is attractive for practical deployment because it keeps the online computational cost low.

2.7. Model Complexity and Practical Considerations

Because the proposed network operates on one-dimensional traces, its computational complexity increases approximately linearly with the trace length T for a fixed channel configuration. Compared with two-dimensional image super-resolution backbones, this design sacrifices part of the explicit spatial context but substantially reduces the parameter count and memory footprint. Such a trade-off is reasonable for seismic resolution enhancement because the primary target of the network is the temporal waveform and local interference pattern within each trace.

Another practical consideration is that amplitude enhancement and structural preservation must remain balanced. Excessive adversarial training may over-sharpen local waveforms, whereas excessive structural regularization may suppress recoverable high-frequency components. For this reason, the proposed framework is intentionally designed as a moderate-capacity architecture with explicit trade-off coefficients, and the influence of these coefficients is further analyzed through sensitivity experiments in Section 4.

3. Experiments and Results

3.1. Experimental Protocol and Implementation Details

To keep the experimental setting consistent with Section 2, the training corpus is not built from a single analytic Ricker wavelet. Instead, representative wavelets extracted from field seismic data are separated into an LR pool with an effective bandwidth of 10–60 Hz and an HR pool with an effective bandwidth of 40–80 Hz. The extraction windows are chosen from comparatively clean and laterally coherent parts of the field profile, and the selected wavelets are amplitude-normalized before being assigned to the two bandwidth pools. This procedure prevents the network from learning a trivial spatial correspondence between LR and HR samples and keeps the two domains separated by their bandwidth characteristics. Independent random reflectivity sequences are then convolved with wavelets from the corresponding pool to form two statistically related but sample-unpaired domains. The complete synthetic database contains 400,000 traces, which are divided into 320,000, 40,000, and 40,000 traces for training, validation, and testing, respectively. Each trace contains 800 samples and is normalized to $[-1, 1]$ before being fed into the network. During optimization, LR and HR mini-batches are sampled independently. Only for the one-way-regression ablation in Section 4 are pseudo-pairs formed by applying two bandwidth settings to the same reflectivity realization.

The main comparison experiments include three data scenarios. The first is the synthetic benchmark constructed from the above wavelet reflectivity procedure. The second is a field subset containing 100 traces extracted from a real two-dimensional seismic profile, with each trace resampled or windowed to 800 samples for consistency with the network input. The third is the public SEAM Phase I Elastic 2DEW Classic line released through SEG Open Data [27,28]. For the public benchmark, the stacked section is windowed into 800-sample traces, and the LR inputs are produced by controlled bandwidth limitation and additive noise so that all compared methods are evaluated under the same degradation setting. Unless otherwise stated, the public benchmark is divided into 70%/10%/20% training, validation, and testing subsets. Adjacent windows from the same local spatial interval are kept in the same subset wherever possible, so that the validation and test results are not inflated by near-duplicate traces.

The proposed network is implemented in PyTorch 2.0.1 and trained on a workstation equipped with an NVIDIA RTX 4090 GPU and an Intel i9 CPU. We use the Adam

optimizer [29] with $\beta_1 = 0.5$ and $\beta_2 = 0.999$, an initial learning rate of 1×10^{-3} , and a batch size of 200. The model is trained for 120 epochs; the learning rate is kept constant during the first 60 epochs and then linearly decayed to zero during the remaining epochs. The cycle-consistency and structural trade-off coefficients are fixed at $\alpha = 5$ and $\beta = 10$, respectively. During training, label flipping is randomly applied to the discriminator outputs after each mini-batch update with probability 0.3 to prevent the discriminator from becoming overly dominant. All random partitions and initialization seeds are fixed before training, and the same preprocessing, normalization, and evaluation scripts are used for all compared learning-based methods. The public-data processing protocol and split indices can therefore be reproduced from the reported settings, while the proprietary field subset is used only for qualitative and continuity-oriented validation. Figure 4 illustrates the evolution of the loss function and shows that the proposed optimization converges steadily.




Figure 4. Evolution of the training losses. The generator, discriminator, cycle-consistency, and SSIM-related losses are used to monitor optimization stability, while the final checkpoint is selected according to validation-set cycle loss, SSIM, and spectral-bandwidth indicators.

3.2. Evaluation Metrics and Compared Methods

The synthetic and public benchmark experiments are quantitatively evaluated using the structural similarity index (SSIM), peak signal-to-noise ratio (PSNR), signal-to-noise ratio (SNR), Pearson correlation coefficient (PCC), dominant frequency, and absolute frequency bandwidth. For the synthetic and public benchmarks, the reference traces are the corresponding broad-band signals generated under the same controlled protocol, so the reported similarity metrics directly reflect waveform recovery quality. For the field profile, where no strict HR reference is available, we report dominant frequency, absolute bandwidth, and adjacent trace correlation (ATC). ATC is computed as the mean Pearson correlation between each trace and its immediate neighboring trace after local amplitude normalization and is used here to characterize lateral continuity. Higher SSIM, PSNR, SNR, PCC, and ATC values indicate better reconstruction quality. In particular, a higher SSIM

means that the reconstructed trace better preserves the local waveform structure, amplitude contrast, and phase-consistent morphology of the reference trace.

In addition to conventional deconvolution and time-varying frequency deconvolution (TVFD) [30], four representative learning-based baselines are introduced for comparison: structure-constrained machine-learning high-resolution reconstruction (SC-MLHR) [14], feature-aware vertical resolution enhancement (F-VRE) [16], pseudo-well CycleGAN enhancement (PW-CycleGAN) [17], and self-supervised seismic resolution enhancement (SS-SRE) [19]. For visual clarity, Figures 5 and 6 retain only the most representative classical baselines, whereas the complete quantitative results of all methods are reported in Tables 1–3, together with the public benchmark plots in Figures 7 and 8.

For fairness, all learning-based baselines are re-trained or reproduced according to their published settings on the same training, validation, and testing partitions, and all quantitative comparisons are computed on identical test windows after amplitude normalization. The conventional baselines are tuned according to standard recommendations in the literature, with their regularization factors selected on the validation subset. Table 4 summarizes the main implementation settings used in the comparative experiments.




Figure 5. Synthetic data comparison test. The horizontal axis denotes trace number and the vertical axis denotes sample/time index. (a) Synthesized LR section. (b) Synthesized HR reference. (c) Result obtained by deconvolution. (d) Result processed by TVFD. (e) Result processed by the proposed method. Red boxes mark areas with clear visual differences.

Table 1. Quantitative comparison on the synthetic benchmark.

Method	SSIM	PSNR (dB)	PCC	Bandwidth (Hz)
LR input	0.694	18.7	0.751	26
Deconvolution	0.781	22.4	0.823	37
TVFD	0.809	23.5	0.846	45
SC-MLHR	0.842	25.0	0.874	49
F-VRE	0.855	25.7	0.881	52
PW-CycleGAN	0.868	26.2	0.893	54
SS-SRE	0.875	26.8	0.901	55
Proposed	0.892	27.6	0.917	57




Figure 6. Field-profile processing results obtained by different methods. The horizontal axis denotes trace number and the vertical axis denotes sample/time index; the right panels show enlarged views of the boxed regions. (a) Original field profile. (b) Result obtained by deconvolution. (c) Result processed by TVFD. (d) Result processed by the proposed method.

Table 2. Field-profile frequency and continuity comparison of different methods.

Method	Dominant Frequency (Hz)	Absolute Bandwidth (Hz)	ATC
Original profile	21	23	0.84
Deconvolution	30	37	0.79
TVFD	32	46	0.81
SC-MLHR	36	50	0.84
F-VRE	38	52	0.85
PW-CycleGAN	39	54	0.83
SS-SRE	41	56	0.86
Proposed	43	58	0.88

Table 3. Quantitative comparison on the SEG Open Data benchmark.

Method	SSIM	PSNR (dB)	SNR (dB)	Bandwidth (Hz)
LR input	0.691	19.8	9.2	26
Deconvolution	0.742	21.9	12.7	38
TVFD	0.761	22.6	13.4	44
SC-MLHR	0.789	23.8	14.8	49
F-VRE	0.804	24.5	15.4	52
PW-CycleGAN	0.817	24.9	15.9	54
SS-SRE	0.826	25.2	16.3	56
Proposed	0.842	26.0	17.1	59



**Figure 7.** Normalized quantitative comparison on the SEG Open Data benchmark. Larger values indicate better performance after metric-wise normalization.**Figure 8.** Average normalized amplitude spectra on the SEG Open Data benchmark. The frequency axis is given in Hz, and the legend distinguishes the LR input, baseline methods, and the proposed method. The yellow shaded region indicates the effective enhancement band.

Table 4. Main implementation settings of the compared methods.

Method	Main Implementation Setting
Deconvolution	Predictive/Wiener-type deconvolution with a 24 ms effective operator, a stabilization factor of 0.01, and validation-based tuning of the prediction lag.
TVFD	Improved time-frequency spectral modelling with a 64-sample analysis window, 75% overlap, and regularization selected on the validation subset.
SC-MLHR	Structure-constrained learning model trained on the same synthetic bandwidth-conversion corpus; structural weight tuned to maximize validation SSIM.
F-VRE	Feature-aware vertical resolution enhancement network with four convolutional stages and the same training/validation/test split as the proposed method.
PW-CycleGAN	CycleGAN-based enhancement model trained using pseudo-well supervision and the same bandwidth settings used for the proposed synthetic domain construction.
SS-SRE	Self-supervised resolution enhancement model trained on the same LR inputs with masking/reconstruction consistency and no explicit paired HR labels.
Proposed	Bidirectional bandwidth translation with cycle consistency and SSIM constraints; $\alpha = 5$, $\beta = 10$, four residual blocks, and LN regularization.

3.3. Synthetic Data Comparison Test

In this section, a synthetic 2D seismic profile is constructed to evaluate the effectiveness of the proposed method. The qualitative comparison results are illustrated in Figure 5. The HR results produced by conventional deconvolution and TVFD are shown in Figure 5c,d, respectively, while the HR result obtained by the proposed method is presented in Figure 5e. The regions with noticeable differences among the methods are highlighted by red boxes. It can be observed that the seismic signals reconstructed by the proposed approach exhibit a narrower waveform envelope and maintain stable reconstruction performance even in low-amplitude regions. In addition, the proposed method reveals more high-resolution structural details than the conventional baselines.

To further examine the local characteristics, a representative seismic trace is extracted from Figure 5 and plotted in the same panel for comparison. The traces obtained from different methods are displayed using different colors, as shown in Figure 9. The HR trace generated by the proposed method (blue curve) shows strong agreement with the synthesized HR reference trace (black curve). In particular, the region highlighted by the box demonstrates a clearer match between the reconstructed and reference signals, whereas the other two methods exhibit noticeably poorer agreement in this area. The complete numerical results of all compared methods are summarized in Table 1, where the proposed method achieves the highest SSIM, PSNR, PCC, and effective bandwidth.

The synthetic benchmark also reveals an important methodological trend. Classical methods increase bandwidth to some extent, but the gain is accompanied by either residual waveform broadening or insufficient reflector separation. By contrast, the learning-based methods improve all reported metrics more uniformly, and the proposed bidirectional model yields the strongest overall result. The monotonic improvement from LR input to the proposed method in Table 1 is consistent with the visual evidence in Figures 5 and 9, suggesting that the network is not merely amplifying high frequencies, but is recovering waveforms that remain closer to the HR reference in both phase and morphology.




Figure 9. Trace-amplitude comparison obtained by different methods. The horizontal axis denotes the sample/time index and the vertical axis denotes normalized amplitude. Different colors identify the reference and reconstructed traces, as indicated in the legend. The inset provides a zoomed view of the selected local waveform segment, and the red arrows highlight local amplitude and waveform differences among the compared methods.

3.4. Field Seismic Data Processing

The structure and distribution of real seismic data are generally much more complex than those of synthetic datasets. To further evaluate the effectiveness of the proposed approach, 100 traces were extracted from a real two-dimensional seismic profile, each containing 800 sampling points. The results obtained using different processing methods are presented in Figure 6.

It can be observed that the proposed method produces seismic signals with higher resolution and reveals clearer structural details compared with the conventional approaches. In particular, the reconstructed seismic traces exhibit improved continuity of reflection events and enhanced high-frequency components, indicating that the proposed network is capable of effectively recovering fine-scale geological information from real seismic data. When the additional learning-based baselines are considered, the proposed method still achieves the best overall balance between bandwidth gain and lateral continuity, as summarized in Table 2.

Figure 10 shows the average amplitude spectra of the field-profile subset. It can be seen that compared with the deconvolution and TVFD methods, the proposed method recovers the spectral components above about 30 Hz more effectively. Table 2 further reports the complete frequency and continuity statistics for all methods on the field subset.

An additional observation from the field test is that the proposed method improves continuity while increasing the effective bandwidth. This point is important because excessive sharpening often increases the dominant frequency at the expense of lateral stability. In Table 2, however, the proposed method achieves both the highest dominant frequency and the highest ATC, which indicates that the recovered high-frequency components are not dominated by random oscillatory artifacts. This behavior is consistent with the design goal of jointly improving spectral detail and structural preservation.




Figure 10. Average amplitude spectra of the field-profile subset. The frequency axis is given in Hz and amplitudes are normalized for comparison among methods. Different colors denote the actual data and the results obtained by deconvolution, TVF-division, and the proposed method, as indicated in the legend.

3.5. Public SEG Open Data Benchmark

To improve reproducibility, we further evaluated the method on the SEAM Phase I Elastic 2DEW Classic line released through SEG Open Data [27,28]. According to the official description, this public two-dimensional east-west dataset contains 151 shots, 901 receivers per shot, and 2001 samples per trace with an 8 ms sampling interval. For compatibility with the proposed trace-wise network, the benchmark line was stacked and windowed to 800-sample traces. The broader-band stacked traces were used as pseudo-HR references, and the LR inputs were generated by applying a 10–35 Hz band limitation together with 5 dB Gaussian noise. We used a 70%/10%/20% split for training, validation, and testing, and all comparison methods were evaluated on the same test windows.

The quantitative results are summarized in Table 3. The proposed method achieves the best SSIM, PSNR, SNR, and effective bandwidth, which indicates that the bidirectional domain constraint improves both waveform fidelity and spectral recovery. Figure 7 presents a normalized comparison of the main metrics, and Figure 8 shows the averaged amplitude spectra. Compared with PW-CycleGAN, the proposed method recovers a broader high-frequency band while maintaining better structural consistency. Compared with SS-SRE, the proposed method produces a slightly higher bandwidth gain and a smoother spectral roll-off in the 35–60 Hz range.

Compared with the synthetic experiment, the public benchmark is more challenging because its waveforms contain richer structural variability and stronger mismatch between local signal patterns. The fact that the proposed method remains superior under this setting suggests that the benefit of bidirectional bandwidth translation is not limited to one curated synthetic generator. Instead, the cycle constraint and the field-guided domain construction appear to improve transferability when the degradation characteristics vary across traces and windows.

4. Ablation and Extended Analysis

4.1. Ablation on Objective Terms and Domain Construction

To better understand which part of the proposed framework contributes most to the final performance, we conduct ablation experiments on the synthetic benchmark, the public SEG Open Data benchmark, and the field subset. The tested variants include a one-way paired regression model, a model without the cycle-consistency term, a model without the SSIM term, a model trained without field-guided wavelet extraction, and a model trained without label flipping. The corresponding results are reported in Table 5.

Several trends can be observed. First, replacing the bidirectional objective with a one-way regression strategy leads to the largest performance drop on both synthetic and public benchmarks, indicating that reverse-domain consistency is essential for suppressing unstable spectral amplification. Second, removing the SSIM term causes only a moderate decrease in synthetic SSIM but a more obvious loss in field ATC, which confirms that the structural term is especially important for preserving lateral continuity in real data. Third, training without field-guided wavelets substantially weakens transfer performance on SEG Open Data, suggesting that training-domain realism is necessary for improving generalization to public benchmark data. Finally, label flipping contributes mainly to training stability and therefore produces a smaller but still measurable benefit.

The normalized comparison shown in Figure 11 provides a more intuitive view of the contribution of each module. In particular, the full model is not the highest on every single metric by a large margin, but it is the only configuration that remains consistently strong across synthetic accuracy, public-benchmark transfer, and field continuity. This observation supports the main claim of the paper: the advantage of the proposed framework lies in balanced enhancement rather than in aggressively maximizing one specific metric.

Table 5. Ablation study on loss terms and domain construction strategy.

Variant	Synthetic SSIM	SEG SSIM	Field Bandwidth (Hz)	Field ATC
One-way paired regression	0.847	0.801	51	0.81
w/o cycle consistency	0.861	0.814	56	0.80
w/o SSIM constraint	0.873	0.821	59	0.83
w/o field-guided wavelets	0.858	0.808	53	0.82
w/o label flipping	0.880	0.829	57	0.86
Full model	0.892	0.842	59	0.88




Figure 11. Normalized comparison of the ablation variants. Each axis corresponds to one evaluation metric, and larger normalized values indicate stronger overall performance.

4.2. Influence of Network Architecture

The effect of the network backbone is further analyzed in Table 6. We compare a plain encoder–decoder, residual generators with different numbers of bottleneck blocks, and different normalization strategies under the same training configuration. The results indicate that increasing the number of residual blocks from two to four leads to a substantial improvement, whereas the gain from four to six blocks is relatively small compared with the increase in model size and inference time. In addition, LN yields better performance than BN and instance normalization in our trace-wise setting, likely because LN is less sensitive to mini-batch statistics and therefore better suited to the adversarial optimization of one-dimensional seismic traces.

Table 6. Architecture ablation and efficiency comparison.

Variant	Params (M)	Synthetic SSIM	SEG SSIM	Inference Time (ms/trace)
Plain encoder–decoder	1.72	0.857	0.807	0.41
2 residual blocks + LN	2.11	0.873	0.823	0.48
4 residual blocks + LN	2.64	0.892	0.842	0.56
6 residual blocks + LN	3.18	0.894	0.844	0.68
4 residual blocks + BN	2.64	0.881	0.831	0.55
4 residual blocks + IN	2.64	0.886	0.836	0.55

This result justifies the use of four residual blocks in the final architecture. From a practical perspective, the selected configuration offers a favorable compromise between reconstruction accuracy and computational efficiency. The inference latency remains well below 1 ms per trace on the test workstation, which is sufficient for offline section processing and can be further improved through batched inference.

4.3. Sensitivity to Trade-Off Coefficients

The hyperparameters α and β control the relative contribution of cycle consistency and structural preservation. To verify that the final choice is not arbitrary, we vary α from 1 to 10 and β from 0 to 15 while keeping all other settings unchanged. Figure 12 reports the synthetic SSIM and field ATC under these settings.




Figure 12. Sensitivity analysis of the trade-off coefficients α and β . The plotted values report validation/test trends for structural fidelity and field continuity under otherwise identical settings.

The sensitivity curves reveal two useful trends. When α is too small, the reverse-domain constraint is insufficient and the generated traces become spectrally sharper but less stable, which reduces the field ATC. When α is too large, the model becomes conservative and tends to preserve the original LR bandwidth. A similar trade-off is observed for β : introducing an SSIM term clearly improves structural continuity, but an excessively large

β slightly suppresses recoverable spectral details. In our experiments, $\alpha = 5$ and $\beta = 10$ provide the best overall balance.

4.4. Robustness to Noise and Bandwidth Mismatch

In practical seismic processing, LR inputs may differ substantially in both noise level and effective bandwidth. We therefore evaluate the robustness of the proposed method under progressively more difficult input conditions. Table 7 and Figure 13 show the SSIM obtained on the SEG Open Data benchmark when additional Gaussian noise is injected into the LR traces.

Table 7. Robustness comparison under different input noise levels on the SEG Open Data benchmark.

Input SNR (dB)	Deconv.	TVFD	SS-SRE	Proposed
20	0.776	0.791	0.843	0.857
10	0.748	0.769	0.824	0.840
5	0.716	0.742	0.806	0.827
0	0.683	0.708	0.782	0.809





Figure 13. SSIM degradation under different input noise levels on the SEG Open Data benchmark. The horizontal axis denotes input SNR in dB and the vertical axis denotes SSIM.

Although all methods degrade as the input SNR decreases, the proposed approach degrades more gracefully than the classical baselines and remains consistently superior to the strongest learning-based baseline shown here. This behavior can be explained by the dual regularization mechanism: the cycle loss limits unsupported enhancement, while the SSIM term discourages noise-induced waveform distortion.

To further examine cross-condition generalization, we evaluate the models under several LR bandwidth settings that differ from the main training configuration. The results are summarized in Table 8, and the same trend is visualized in Figure 14. The proposed model remains the best-performing method under all tested bandwidth settings, which suggests that the learned representation is not restricted to one specific degradation pattern.

Table 8. Cross-bandwidth generalization on the SEG Open Data benchmark (SSIM).

LR Bandwidth Setting (Hz)	F-VRE	PW-CycleGAN	SS-SRE	Proposed
10–30	0.781	0.793	0.802	0.816
10–35	0.804	0.817	0.826	0.842
10–40	0.798	0.809	0.819	0.834
15–45	0.775	0.789	0.797	0.811

**Figure 14.** Heatmap of cross-bandwidth generalization performance. Rows correspond to LR bandwidth settings in Hz and colors represent SSIM values on the SEG Open Data benchmark.

4.5. Practical Implications of the Extended Experiments

The extended analyses have two practical implications. First, the proposed method should not be interpreted as a purely black box sharpening model, because the ablation experiments demonstrate that each component plays a distinct and physically meaningful role in balancing enhancement and stability. Second, the robustness experiments indicate that the framework can tolerate moderate mismatches in input noise level and bandwidth, which is important for real deployment where the degradation operator is seldom known exactly.

Taken together, the ablation, sensitivity, and robustness results strengthen the conclusion drawn from the main comparison experiments. The proposed framework improves not only peak quantitative accuracy under a single curated setting, but also the stability and transferability that are required for realistic seismic-resolution-enhancement workflows.

5. Discussion

The experiments on synthetic data, the field example, and the SEG Open Data benchmark consistently show that the proposed framework improves seismic resolution without severely sacrificing structural continuity. Compared with classical deconvolution and TVFD, the method recovers broader bandwidth and produces sharper reflection events, especially in low-amplitude regions. Compared with recent learning-based baselines, the gain is more moderate but consistent, indicating that the improvement mainly comes from a better balance between spectral extension and structural preservation rather than from aggressive sharpening alone.

5.1. Interpretation of the Performance Gains

The results clarify the role of the methodological design. The bidirectional domain formulation reduces the risk of unsupported high-frequency hallucination because every enhanced trace must remain recoverable in the LR domain. This mechanism is especially important for seismic traces, where visually plausible oscillations can still be geologically misleading if they are not constrained by the original band-limited observation. The cycle path therefore acts as a physical plausibility filter: it allows the model to sharpen reflectors only when the sharpened output remains compatible with the information content of the LR input.

The structural term plays a complementary role. In the revised formulation adopted in this manuscript, SSIM is applied to cycle-reconstructed traces rather than to unmatched cross-domain samples, so the structural penalty is always computed on sample-consistent signal pairs. This design is better aligned with the unpaired training assumption and explains why the proposed method improves ATC in the field experiment instead of merely increasing dominant frequency. In other words, the structural term does not simply reward local sharpness; it rewards the preservation of waveform morphology that survives a forward-and-back translation process.

The effectiveness of the field-guided wavelet strategy is also supported by the transfer results. A purely analytic synthetic training set can easily bias a model toward over-idealized waveforms and deterministic degradations. By incorporating wavelets extracted from field data, the constructed LR and HR domains better reflect realistic source signatures and bandwidth variations. The resulting model is therefore trained on a more representative distribution of oscillatory patterns, which helps explain the improved performance on the SEG Open Data benchmark.

5.2. Practical Applicability and Deployment Considerations

From a practical viewpoint, the proposed framework is attractive because inference requires only the forward generator. Once training is completed, each input trace is enhanced in a single pass without iterative inversion, explicit Q estimation, or repeated optimization. The architecture is therefore suitable for offline section processing and for integration into broader seismic interpretation workflows in which computational cost, reproducibility, and pipeline simplicity all matter.

The experiments further suggest several conditions under which the method is most useful. First, the framework is well suited to scenarios where the input data suffer from moderate bandwidth loss but still retain stable large-scale reflector geometry. In such cases, the cycle constraint can guide the model toward conservative yet meaningful spectral recovery. Second, the method is particularly valuable when classical inverse- Q filtering is unstable or when accurate attenuation estimation is unavailable. Third, the ablation and robustness results indicate that the approach tolerates moderate mismatch in noise level and bandwidth, which is important because real acquisition conditions are rarely uniform across an entire line or survey.

Classical sparse reconstruction methods such as orthogonal matching pursuit address bandwidth extension by selecting atoms from an overcomplete dictionary through orthogonal least-squares updates. Their computational stability depends on dictionary coherence, stopping criteria, regularization, and noise level, and high-frequency energy is commonly controlled through sparsity constraints or residual thresholds. The proposed method does not explicitly use OMP. Instead, spectral stability is imposed implicitly by cycle consistency, adversarial domain matching, SSIM regularization, and validation of the recovered bandwidth. Therefore, high-frequency components are encouraged only when they are supported by both the learned HR domain and the recoverability of the LR input.

At the same time, the proposed approach should not be interpreted as a universal replacement for physics-based processing. If the input data are severely contaminated by coherent noise, footprint artifacts, or grossly incorrect amplitudes, the network may enhance the trace in a way that is visually sharper but still limited by the quality of the incoming signal. For this reason, the method is better viewed as a complementary post-processing tool that should be used after basic denoising, amplitude balancing, and quality control have been completed.

5.3. Reproducibility and Validity Considerations

Another important point is how the reported gains should be interpreted from the standpoint of reproducibility. In seismic-resolution-enhancement studies, performance can be overstated when different methods are trained on different synthetic generators, different frequency bands, or different data splits. To reduce this risk, the present manuscript uses unified LR/HR bandwidth settings, common train-validation-test partitions, and the same evaluation windows across the learning-based baselines whenever reproduction is possible from the published descriptions. This does not eliminate all uncertainty, but it makes the comparison substantially fairer than a setting in which each method is allowed to use a different synthetic protocol.

At the same time, several sources of evaluation bias remain unavoidable. Synthetic and public benchmark experiments rely on pseudo-HR references rather than on directly observed field HR signals. Consequently, similarity metrics such as SSIM, PSNR, SNR, and PCC should be interpreted as indicators of consistency with the adopted degradation model, not as absolute proof that every recovered oscillation corresponds to a uniquely correct geological reflector. This is precisely why the field experiment remains necessary: it complements reference-based metrics with continuity-oriented and spectrum-oriented indicators that are closer to real deployment conditions.

When only limited data are available, the model should be interpreted through validation curves, ablation results, spectra, and continuity metrics rather than through a single accuracy value. In practice, overfitting is mitigated by field-guided synthetic augmentation, amplitude perturbation, additive noise, dropout, moderate model capacity, fixed validation splits, and sensitivity analysis of alpha and beta. If validation loss decreases while bandwidth or continuity metrics deteriorate, the structural weight should be increased or the adversarial contribution reduced; if enhancement is too conservative, the cycle or structural weights can be moderately relaxed.

The field evaluation also has its own limitations. The selected 100-trace subset is useful for controlled visual comparison, but it cannot capture the full diversity of noise patterns, structural styles, and acquisition artifacts that may appear in larger field surveys. Likewise, dominant frequency and absolute bandwidth are informative but not sufficient on their own, because aggressive spectral amplification can sometimes increase both quantities without improving interpretability. For this reason, the manuscript reports ATC together with spectral indicators and emphasizes balanced improvement across multiple metrics rather than the maximization of one spectral statistic in isolation.

5.4. Limitations and Future Work

Several limitations should nevertheless be noted. First, although the SEG Open Data experiment improves reproducibility, the HR target on the public benchmark is still generated under a controlled bandwidth-degradation protocol rather than obtained from direct field measurement. The benchmark is therefore useful for standardized comparison, but it does not fully eliminate the uncertainty inherent in defining a “true” HR seismic target.

Second, the current backbone is trace-wise and only indirectly captures cross-trace dependencies through the training distribution and structural loss. This choice keeps the network lightweight and computationally efficient, but it may limit the recovery of larger-scale lateral structures, stratigraphic continuity, and multi-trace interference patterns. A two-dimensional or hybrid trace-section architecture may offer additional gains when sufficient training data and computational resources are available.

Third, the present training corpus mainly reflects bandwidth variation and random noise, while other practical degradations such as statics errors, footprint contamination, anisotropy-related distortions, and strong nonstationary attenuation are not modeled explicitly. Extending the framework toward physics-guided or acquisition-aware training is therefore a natural next step. Future work will focus on introducing stronger cross-trace structural coupling, explicit attenuation priors, and more diverse open benchmarks so that the method can be assessed under a broader range of geological and acquisition conditions.

6. Conclusions

This paper presents a structure-preserving bidirectional translation framework for high-resolution seismic reconstruction. Instead of describing the method merely as a GAN variant, the proposed formulation interprets seismic resolution enhancement as unpaired translation between low-bandwidth and high-bandwidth seismic domains, where adversarial learning is combined with cycle consistency and SSIM-based structural preservation.

The main contributions of this work can be summarized as follows. First, a bidirectional domain formulation is introduced for seismic resolution enhancement, enabling stable LR-to-HR translation while using the reverse mapping as a regularizer that suppresses physically inconsistent spectral enhancement. Second, the combination of field-extracted wavelets, synthetic reflectivity construction, and a lightweight one-dimensional residual backbone provides a practical training strategy for trace-wise seismic bandwidth extension. Third, by redefining the SSIM term on cycle-reconstructed traces, the final objective remains fully consistent with unpaired training while explicitly preserving waveform morphology and reflector continuity.

Comprehensive experiments on synthetic data, a field profile, and the SEAM Phase I SEG Open Data benchmark show that the proposed method achieves the best overall balance between bandwidth gain, waveform fidelity, and structural continuity among the compared methods. The ablation, sensitivity, and robustness studies further show that the performance gain does not come from one isolated design choice, but from the coordinated effect of bidirectional translation, structural regularization, and field-guided domain construction.

These results indicate that the proposed framework is a competitive and extensible solution for seismic resolution enhancement. Future work will focus on introducing stronger cross-trace structural coupling, broader public benchmarks, and additional physics-guided constraints for more challenging geological scenarios.

Author Contributions: Conceptualization, S.-Y.C. and M.Y.; Methodology, S.-Y.C. and M.Y.; Software, S.-Y.C.; Formal analysis, S.-Y.C.; Investigation, S.-Y.C.; Data curation, S.-Y.C.; Writing—original draft, S.-Y.C.; Writing—review & editing, M.Y.; Supervision, M.Y.; Project administration, M.Y.; Funding acquisition, M.Y. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the Key Laboratory of Flight Techniques and Flight Safety under Grant Nos. FZ2025ZX32 and FZ2025ZX17.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The public benchmark data used in this study were obtained from the SEAM Phase I Elastic 2DEW Classic subset released through SEG Open Data (<https://seg.org/SEAM/open-data/>, accessed on 30 March 2026). The proprietary field seismic example used for qualitative evaluation is not publicly available. The synthetic domain construction rules, random split strategy, and main implementation settings are described in Sections 2 and 3 to facilitate reproduction of the public benchmark experiments.

Acknowledgments: The authors would like to thank Xi-Jun Liu for his helpful suggestions and constructive comments, which improved the clarity of this manuscript.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Robinson, E.A. Predictive decomposition of seismic traces. *Geophysics* **1957**, *22*, 767–778. [\[CrossRef\]](#)
2. Clarke, G.K.C. Time-varying deconvolution filters. *Geophysics* **1968**, *33*, 936–944. [\[CrossRef\]](#)
3. Griffiths, L.J.; Smolka, F.R.; Tremblay, L.D. Adaptive deconvolution: A new technique for processing time-varying seismic data. *Geophysics* **1977**, *42*, 742–759. [\[CrossRef\]](#)
4. Margrave, G.F.; Lamoureux, M.P.; Henley, D.C. Gabor deconvolution: Estimating reflectivity by nonstationary deconvolution of seismic data. *Geophysics* **2011**, *76*, W15–W30. [\[CrossRef\]](#)
5. Futterman, W.I. Dispersive body waves. *J. Geophys. Res.* **1962**, *67*, 5279–5291. [\[CrossRef\]](#)
6. Kjartansson, E. Constant Q -wave propagation and attenuation. *J. Geophys. Res.* **1979**, *84*, 4737–4748. [\[CrossRef\]](#)
7. Hargreaves, N.D.; Calvert, A.J. Inverse Q filtering by Fourier transform. *Geophysics* **1991**, *56*, 519–527. [\[CrossRef\]](#)
8. Wang, Y. A stable and efficient approach of inverse Q filtering. *Geophysics* **2002**, *67*, 657–663. [\[CrossRef\]](#)
9. Wang, Y. Inverse Q -filter for seismic resolution enhancement. *Geophysics* **2006**, *71*, V51–V60. [\[CrossRef\]](#)
10. Zhang, C.; Ulrych, T.J. Seismic absorption compensation: A least-squares inverse scheme. *Geophysics* **2007**, *72*, R109–R114. [\[CrossRef\]](#)
11. Qian, F.; Yin, M.; Liu, X.; Wang, Y.; Lu, C.; Hu, G. Unsupervised seismic facies analysis via deep convolutional autoencoders. *Geophysics* **2018**, *83*, A39–A43. [\[CrossRef\]](#)
12. Kim, Y.; Nakata, N. Geophysical inversion versus machine learning in inverse problems. *Lead. Edge* **2018**, *37*, 894–901. [\[CrossRef\]](#)
13. Chen, H.; Gao, J.; Jiang, X.; Gao, Z.; Zhang, W. Optimization-inspired deep learning high-resolution inversion for seismic data. *Geophysics* **2021**, *86*, R265–R276. [\[CrossRef\]](#)
14. Gao, Y.; Zhang, J.; Li, H.; Li, G. Incorporating structural constraint into the machine learning high-resolution seismic reconstruction. *IEEE Trans. Geosci. Remote Sens.* **2022**, *60*, 5912712. [\[CrossRef\]](#)
15. Jo, Y.; Choi, Y.; Seol, S.J.; Byun, J. Machine learning-based vertical resolution enhancement considering the seismic attenuation. *J. Petrol. Sci. Eng.* **2022**, *208*, 109657. [\[CrossRef\]](#)
16. Gao, Y.; Zhao, D.; Li, T.; Li, G.; Guo, S. Deep learning vertical resolution enhancement considering features of seismic data. *IEEE Trans. Geosci. Remote Sens.* **2023**, *61*, 5900913. [\[CrossRef\]](#)
17. Zhao, X.; Gao, Y.; Guo, S.; Gu, W.; Li, G. Seismic resolution enhancement using a cycle generative adversarial neural network with pseudo-well data. *Appl. Sci.* **2023**, *13*, 12980. [\[CrossRef\]](#)
18. Yin, X.; Xu, W.; Yang, Z.; Wu, B. Seismic blind deconvolution based on self-supervised machine learning. *Appl. Sci.* **2024**, *14*, 5214. [\[CrossRef\]](#)
19. Cheng, S.; Zhang, H.; Alkhalifah, T. Self-supervised seismic resolution enhancement. *IEEE Trans. Geosci. Remote Sens.* **2025**, *63*, 5904115. [\[CrossRef\]](#)
20. Chen, S.; Cao, S.; Sun, Y. Enhancing the resolution of seismic data based on the non-local similarity. *Geophys. Prospect.* **2022**, *70*, 1116–1128. [\[CrossRef\]](#)
21. Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; Bengio, Y. Generative adversarial nets. In *Advances in Neural Information Processing Systems 27*; Curran Associates, Inc.: Red Hook, NY, USA, 2014; pp. 2672–2680.
22. Zhu, J.-Y.; Park, T.; Isola, P.; Efros, A.A. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Venice, Italy, 22–29 October 2017; pp. 2242–2251.
23. Wang, Z.; Simoncelli, E.P.; Bovik, A.C. Multiscale structural similarity for image quality assessment. In *Proceedings of the 37th Asilomar Conference on Signals, Systems & Computers*, Pacific Grove, CA, USA, 9–12 November 2003; Volume 2, pp. 1398–1402.
24. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Las Vegas, NV, USA, 27–30 June 2016; pp. 770–778.
25. Ba, J.L.; Kiros, J.R.; Hinton, G.E. Layer normalization. *arXiv* **2016**, arXiv:1607.06450. [\[CrossRef\]](#)

26. Srivastava, N.; Hinton, G.E.; Krizhevsky, A.; Sutskever, I.; Salakhutdinov, R. Dropout: A simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* **2014**, *15*, 1929–1958.
27. SEG. SEAM Open Data. Available online: <https://seg.org/SEAM/open-data/> (accessed on 30 March 2026).
28. Fehler, M.; Lerner, K. SEG Advanced Modeling (SEAM): Phase I first year update. *Lead. Edge* **2008**, *27*, 1006–1007. [[CrossRef](#)]
29. Kingma, D.P.; Ba, J. Adam: A method for stochastic optimization. *arXiv* **2014**, arXiv:1412.6980.
30. Zhou, H.; Wang, C.; Marfurt, K.J.; Jiang, Y.; Bi, J. Enhancing the resolution of non-stationary seismic data using improved time-frequency spectral modelling. *Geophys. J. Int.* **2016**, *205*, 203–219. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.