Geophysical Journal International

Geophys. J. Int. (2026) **247**, 1–17 https://doi.org/10.1093/gji/ggag317
Advance Access publication 2026 August 08
Research Paper

Deep learning-based seismic reflectivity estimation by pre-training on labelled synthetic data and physics-guided fine-tuning in field data

Yutting Wang [1,2,3] Jintao Li [4] Xiaoming Sun[5] and Xinming Wu [1,2,3]

[1] *School of Earth and Space Sciences, University of Science and Technology of China, Hefei 230026, China. E-mail: xinmwu@ustc.edu.cn*
[2] *State Key Laboratory of Precision Geodesy, University of Science and Technology of China, Hefei 230026, China*
[3] *Mengcheng National Geophysical Observatory, University of Science and Technology of China, Hefei 230026, China*
[4] *State Key Laboratory of Ocean Sensing & Ocean College, Zhejiang University, Zhoushan, Zhejiang 316021, China. E-mail: lijintao@zju.edu.cn*
[5] *Institute of Advanced Technology, University of Science and Technology of China, Hefei 230088, China*

Accepted 2026 July 31. Received 2026 June 28; in original form 2026 May 12

SUMMARY

Reflectivity estimation aims to enhance the resolution of seismic data, providing crucial support for the detailed inversion of reservoir parameters. We propose a method that combines supervised pre-training with synthetic data and physics-guided fine-tuning in field data to estimate reasonable reflectivity from seismic data. Initially, a U-shaped network is pre-trained by supervised learning on a large amount of synthetic seismic data. Subsequently, multiple geophysically meaningful constraints including structure-oriented smoothness, reflectivity sparsity and data reconstruction, are introduced to formulate a self-supervised or unsupervised learning mechanism to fine-tune the pre-trained model so that it is better adapted to field data for obtaining more reasonable reflectivity. The pre-trained model provides an initial reflectivity that aligns with fundamental structural features. Based on the initial estimate, the model is further optimized through physics-guided fine-tuning to obtain a more reasonable reflectivity estimation that better reflects the characteristics of the field data and geophysical priors. Furthermore, the well-log-based correlation evaluation metric is applied to automatically and adaptively determine the early-stopping point of the fine-tuning process where favourable results are achieved. Experiments on synthetic and field seismic data confirm that the proposed method yields reasonable and high-resolution reflectivity estimation.

Key words: Image processing; Inverse theory; Machine learning; Neural networks, fuzzy logic; Waveform inversion.

1 INTRODUCTION

Enhancing the resolution of seismic data for detailed reservoir characterization has long been a central goal in seismic exploration (E.A. Robinson 1967). Seismic reflectivity estimation is commonly employed to improve the vertical resolution of seismic data, offering valuable insights into the stratigraphic framework and subsurface physical properties. Although reflectivity inversion techniques have the potential to enhance resolution under appropriate assumptions, such improvement is not always guaranteed (L. Liang & J.P. Castagna 2017). Nevertheless, reflectivity inversion remains a powerful tool for the characterization and delineation of reservoir structures (G.F. Margrave *et al.* 2011).

Seismic reflectivity represents the amplitude contrast generated by changes in impedance across subsurface interfaces and therefore links seismic observations to subsurface rock properties. In a general setting (including elastic and viscoelastic media), impedance contrasts depend on variations in density and wave velocities and may also be influenced by attenuation and dispersion. These impedance variations are directly related to geological properties such as lithology and fluid content, making reflectivity a valuable indicator for quantitative interpretation (R. O'doherty & N.A. Anstey 1971; R.E. Sheriff & L.P. Geldart 1995; Ö. Yilmaz 2001; K. Aki & P.G. Richards 2002). In this study, we focus on post-stack reflectivity estimation, where the normal-incidence approximation is commonly adopted and the reflectivity is typically modelled using an acoustic formulation.

Over the past few decades, numerous methods for reflectivity estimation have been proposed, such as homomorphic deconvolution (T.J. Ulrych 1971; T.G. Stockham *et al.* 1975), maximum entropy spectral deconvolution (T.J. Ulrych & T.N. Bishop 1975), least-squares deconvolution (A. Berkhout 1977), statistical wavelet deconvolution (S. Levy & P.K. Fullagar 1981), inverse Q filtering (Y. Wang 2002, 2006) and sparse spike deconvolution (M.D. Sacchi 1997; O. Portniaguine & J. Castagna 2005; D. Velis 2008), all aimed at enhancing the resolution and interpretability of seismic data. The most fundamental of these techniques is least-squares deconvolution, which attempts to recover the reflection coefficients by convolving the inverse seismic wavelet with the seismic data (Ö. Yilmaz 2001). However, this method assumes that noise is white and follows a Gaussian distribution, whereas field seismic data often contain non-white noise, potentially correlated, thus impacting deconvolution results. Furthermore, least-squares deconvolution tends to amplify noise in the presence of missing or attenuated frequency components, making it unsuitable for bandwidth extension. As a result, it often fails to estimate sharp or sparse seismic reflectivity series. Sparse spike deconvolution has gained notable attention for its ability to produce high-resolution reflectivity profiles by assuming the reflectivity series is sparse. R. Zhang & J. Castagna (2011) propose a basis pursuit inversion method for seismic reflectivity, and S. Yuan & S. Wang (2013) combine spectral reflectivity estimation with sparse Bayesian learning. Both approaches are applied on a trace-by-trace basis. However, these trace-by-trace methods may fail to fully exploit the information from adjacent traces and the inherent spatial continuity of seismic data. Therefore, multitrace sparse spike deconvolution methods (K.F. Kaaresen & T. Taxt 1998; F. Sroubek & P. Milanfar 2011; N. Kazemi *et al.* 2016; M. Ma *et al.* 2017; W. Sun *et al.* 2023) have been extensively studied. By simultaneously integrating multitrace information, these methods enhance reflectivity continuity and reduce noise impact. However, their limitations stem from the need to solve a nonlinear and non-convex objective function and adjust various hyperparameters based on different seismic data. Inaccurate prior models can lead to error accumulation in multitrace data processing, affecting deconvolution results (X. Chai *et al.* 2021; Y. Wang *et al.* 2023).

In recent years, deep learning has been widely applied in seismic exploration, achieving remarkable results in various processes such as seismic data processing (F. Yang & J. Ma 2019; S. Yu *et al.* 2019; J. Li *et al.* 2021; J. Zhang *et al.* 2021b; C. Liu *et al.* 2022; Y. Wang *et al.* 2023; F. Brandolin *et al.* 2024; Y. Xiao *et al.* 2024), interpretation (Y. Shi *et al.* 2019; H. Wu *et al.* 2019a; X. Wu *et al.* 2019b; H. Gao *et al.* 2021a; H. Zhang *et al.* 2021a) and inversion (C. Song & T.A. Alkhalifah 2021; L. Zhao *et al.* 2021; H. Huang *et al.* 2022; H. Chen *et al.* 2023; O.M. Saad *et al.* 2024; J. Li *et al.* 2026) including reflectivity inversion. X. Chai *et al.* (2021) propose a deep learning-based multitrace sparse spike deconvolution method that can invert high-resolution results from band-limited seismic data, outperforming traditional trace-by-trace sparse spike deconvolution methods. Z. Gao *et al.* (2021b) propose a generalized convolutional model that uses the fast iterative shrinkage-thresholding algorithm to solve the reflectivity inversion problem and effectively invert high-fidelity reflectivity models from complex seismic data. These methods share the same supervised-learning paradigm, which relies on paired seismic data and reflectivity labels. In practice, such labels are difficult to obtain: Manual interpretation is labour-intensive, and numerical simulations, while useful, cannot fully capture the complexity and variability of field seismic data. As a result, supervised models trained on limited or idealized data sets often struggle when applied to surveys whose noise characteristics, wavelet signatures, or geological settings differ from those seen during training. Although supervised learning can achieve good performance under controlled conditions, its dependence on high-quality labelled data inevitably limits its applicability in practical seismic workflows.

Self-supervised learning has recently demonstrated remarkable success in leveraging large-scale unlabelled data for representation learning in computer vision. Inspired by these advances in representation learning, the geophysical community has also begun to explore self-supervised strategies tailored to seismic applications. Building on these developments, self-supervised or unsupervised learning can fully leverage large amounts of unlabelled seismic data in the context of seismic reflectivity estimation. S. Phan & M.K. Sen (2021) achieved non-stationary reflectivity inversion by using autoencoders trained in an unsupervised manner. X. Chai *et al.* (2023) proposed a self-supervised learning method where a neural network maps seismic data to reflectivity, which is then convolved with a wavelet to obtain the reconstructed seismic data. This reduces reliance on labelled training data, which are often scarce or difficult to obtain in practice. However, because the supervisory signals in self-supervised learning are indirect and largely reconstruction-driven, these methods may not effectively guide the learning process for reflectivity estimation, potentially leading to unrealistic results, especially when the initial model is not sufficiently informative.

To mitigate these limitations, it is essential to incorporate more appropriate geological and geophysical constraints, as well as a reasonable initial model. Prior work has shown that embedding physical assumptions into neural-network-based inversion can enhance stability and promote more realistic solutions (R. Biswas *et al.* 2019) in impedance inversion. This need for both a reasonable initial estimate and appropriate physical regularization mirrors long-established experience in full-waveform inversion, where inadequate initialization or insufficient constraints often lead to non-physical results (A. Guitton *et al.* 2012; D. Datta & M.K. Sen 2016). These observations motivate a framework that combines a stable initial estimate with physics-guided optimization when performing reflectivity inversion on field seismic data. Although this bears some conceptual resemblance to unsupervised domain adaptation (Y. Ganin & V. Lempitsky 2015; M. Long *et al.* 2015; E. Tzeng *et al.* 2017) in transfer-learning strategies, the intention here is not to achieve broad generalization but to obtain a reliable initialization for subsequent survey-specific adaptation.

In this work, we focus on post-stack acoustic reflectivity estimation and propose a two-stage physics-guided deep inversion framework, named Seismic Reflectivity estimation by Pre-training on labelled synthetic data and Physics-guided Fine-tuning in field data (SRPPF). The framework is designed to address the limitations of current deep learning reflectivity estimation methods, specifically the insufficient convergence stability of pure self-supervised models and the limited adaptability of supervised models to field data. Our method relies on the close integration of two complementary stages, both aimed at ensuring process stability and geophysical rigor. The pre-training stage provides a structurally reasonable and stable initialization, mitigating the instability that arises when self-supervised inversion starts from random weights. It is not intended to deliver a general reflectivity model and cannot be expected to succeed under all seismic conditions, but it offers a practical and reliable starting point for the subsequent survey-specific fine-tuning. Subsequently, building on this initial estimate, the model is subjected to physics-guided self-supervised fine-tuning to adapt to the specific characteristics of the target field data. During fine-tuning, we introduce multiple geophysical priors constraints, including sparsity, structure-oriented smoothness and data reconstruction. These constraints guide the optimization process and enable the model to adapt in a survey-specific manner, ultimately producing reflectivity estimates that are consistent with both seismic data features and physical priors. We evaluate the proposed method on both synthetic (Marmousi) and multiple field seismic data sets, and compare it with representative traditional inversion approaches. The results show that our method provides geologically consistent and competitive reflectivity estimates across different survey conditions.

The paper is organized as follows: First, we introduce the architecture of our deep neural network and the synthetic seismic data with reflectivity labels for pre-training, along with the pre-trained model and its initial applications to synthetic and field test data. Next, we provide a detailed explanation of the physics-guided fine-tuning process. Finally, we present the results of different training strategies and loss functions on both synthetic and field data and validate the accuracy of our method using actual well-log data.

2 METHOD

To estimate reasonable reflectivity, we use a strategy of seismic reflectivity estimation by pre-training on labelled synthetic data and physics-guided fine-tuning in field data. As shown in Fig. 1, we first pre-train a model using synthetic data to obtain an initial reflectivity estimate. Then, we fine-tune this pre-trained model using a self-supervised approach to enhance its adaptability to field data, resulting in reasonable reflectivity. The pre-trained model, trained on synthetic data via supervised learning, might not fully adapt to the specific noise and characteristics of field data, as synthetic data cannot entirely replicate field data. Therefore, fine-tuning the model is necessary to further optimize it, making it better suited to the properties of field data. Additionally, we introduce sparsity and structure-oriented smoothness loss functions as geophysical constraints to guide the initial reflectivity towards more physically meaningful results. By minimizing these loss functions, we iteratively update the network parameters to achieve reasonable reflectivity.

2.1 Network and pre-training with labelled synthetic data

2.1.1 Network architecture

The network is a variant of U-Net, consisting mainly of two parts: a basic U-Net network (O. Ronneberger *et al.* 2015; E. Shelhamer *et al.* 2016) and several residual blocks. The U-Net architecture is symmetrical, comprising an encoder on one side and a decoder on the other, with skip connections introduced between the encoder and decoder. Our network contains four downsampling modules. Initially, the network uses a convolutional layer to increase the number of feature channels of the input seismic images to 64. Then, it performs four downsamplings, with each downsampling module containing a $2 \times 2 \times 2$ kernel and a stride-2 max-pooling layer, as well as two convolutional layers with $3 \times 3 \times 3$ kernels. Each convolutional layer is followed by a batch normalization layer and a rectified linear unit (ReLU). After downsampling, the feature channels are 128, 256, 512 and 1024, respectively. Our network also has four upsampling modules, designed oppositely to the downsampling modules, where the upsampling method used is transposed convolution. The second part of our network comprises residual blocks, with each residual block consisting of two convolutional layers, each followed by a batch normalization layer and a ReLU. A skip connection bypasses these two convolutional layers. This design helps mitigate the vanishing gradient and degradation problems that occur in very deep networks. After three residual blocks, a $1 \times 1 \times 1$ convolutional layer is used to reduce the number of feature channels to match those of the field seismic image labels.

2.1.2 Synthetic data generation

First, based on the automatic modelling method proposed by X. Wu *et al.* (2020), we construct a set of horizontally layered heterogeneous impedance models through stochastic simulation, as shown in Fig. 2(a). Considering that folds and faults are two important types of geological structures, we subsequently add fold and fault information to the initial impedance model to make it more realistic.

We simulate folds in the 3-D structure by vertically shearing the horizontally layered initial impedance model, which is statistically modelled based on well-log data, obtaining an initial model with fold structures, as shown in Fig. 2(b). Next, we simulate fault structures in the model using a volume vector field and randomly selected fault parameters to simulate various fault structures. By integrating the volume displacement field into the previously discussed fold model, we obtain an impedance model with fold and fault structures that approximate field data, as shown in Fig. 2(c).

Following these steps, we construct an impedance model containing both fold and fault structures. According to the theory by B. Russell *et al.* (2006), the reflectivity can be derived from impedance, as shown in Fig. 2(d). Then, using the convolution model formula (E.A. Robinson 1957), the reflectivity is convolved with Ricker wavelets of different dominant frequencies, randomly selected within the range of 25–75 Hz, to obtain the corresponding 3-D noise-free seismic data, as shown in Fig. 2(e). The 25–75 Hz range is commonly

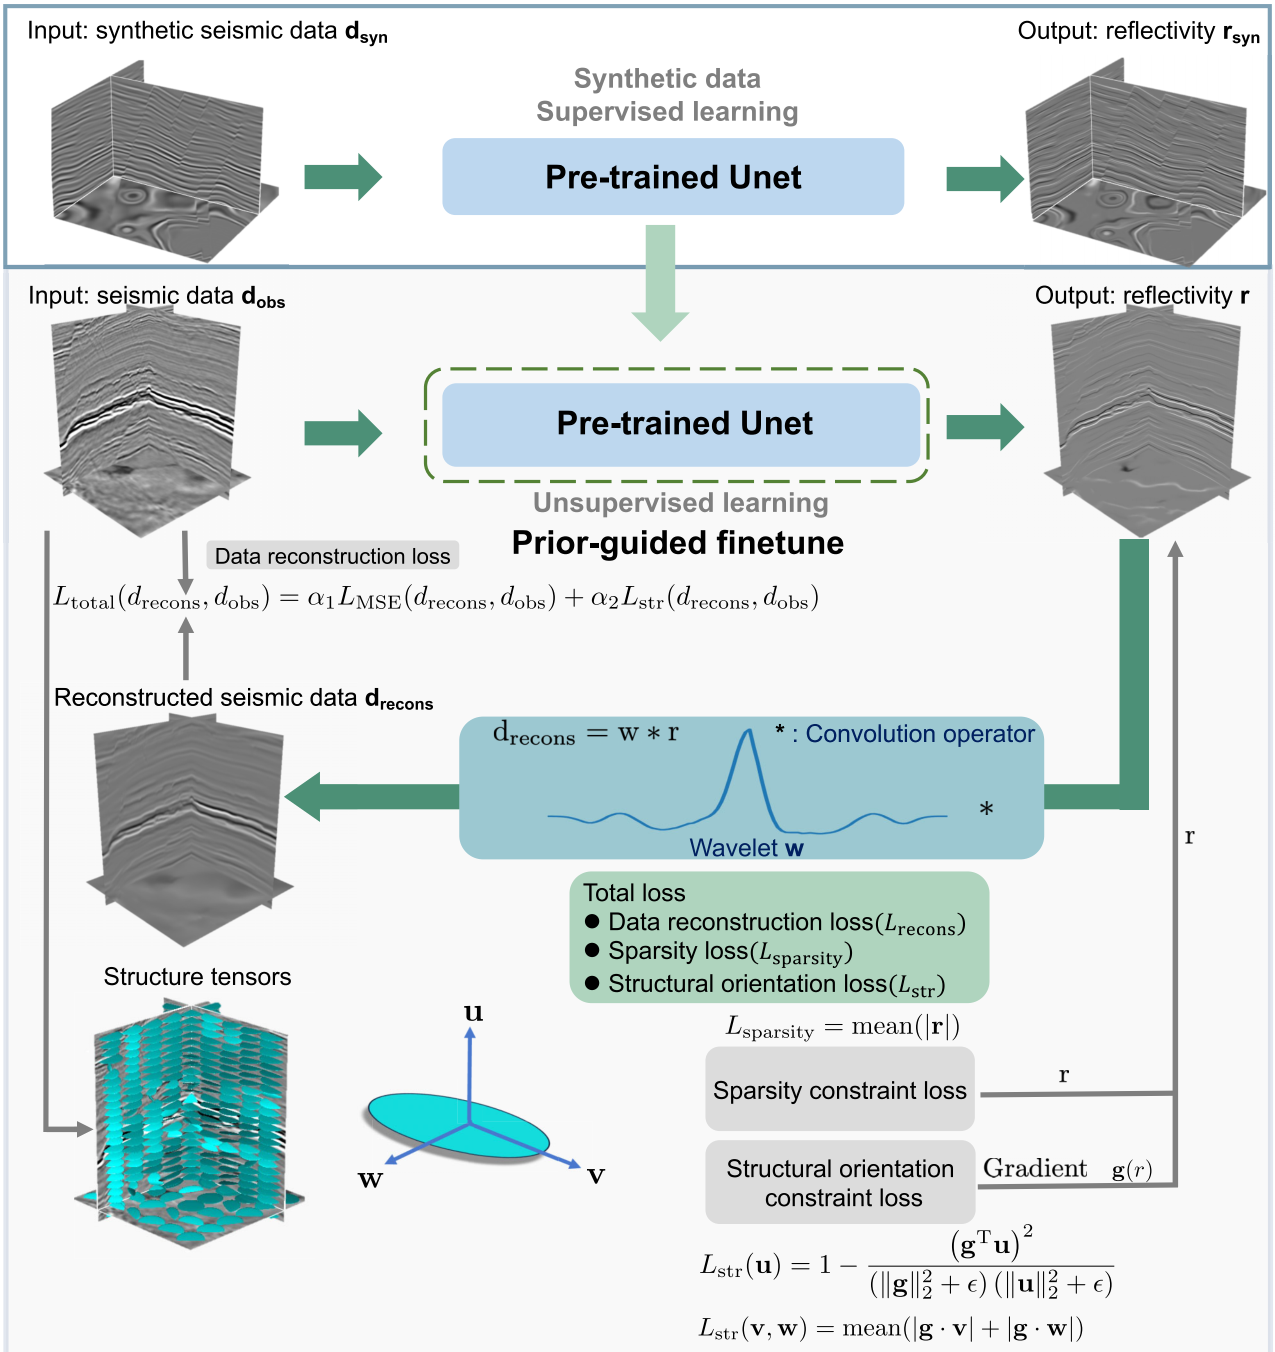

Figure 1. Overview of **S**eismic **R**eflectivity Estimation by **P**re-training on labelled synthetic data and **P**hysics-guided **F**ine-tuning in field data (SRPPF): The U-shaped network is pre-trained with a large amount of synthetic seismic data (d_{syn}, r_{syn}) to obtain an initial reflectivity model that aligns with the structural characteristics of seismic data. Subsequently, this model is fine-tuned using field data d_{obs} in a physics-guided manner, incorporating multiple geophysically meaningful constraints such as structure-oriented smoothness, sparsity and data reconstruction. The reconstructed data d_{recons} is obtained by convolving the predicted reflectivity r with a wavelet $\mathbf{w}$, that is, $d_{recons} = \mathbf{w} * r$.

used in seismic exploration and reflects the spectral characteristics of most real seismic signals, ensuring that the generated synthetic data closely aligns with the frequency characteristics of field seismic data.

After obtaining the 3-D noise-free synthetic seismic data, we add field noise to simulate field seismic data as closely as possible. First, we smooth the field 3-D seismic data volume along the structural direction (D. Hale 2009). By subtracting the smoothed data from the original 3D field seismic data, we obtain the field noise. Within a predetermined range, we randomly select a value as the weight to amplify the noise level accordingly. This noise is then added to the 3-D noise-free synthetic seismic data to obtain synthetic seismic data with field noise. The field noise is derived from multiple 3-D field seismic data sets. We use multiple 3-D field seismic data sets, adding noise with different weights randomly, as shown in Fig. 2(f). This approach helps increase the generalization ability and

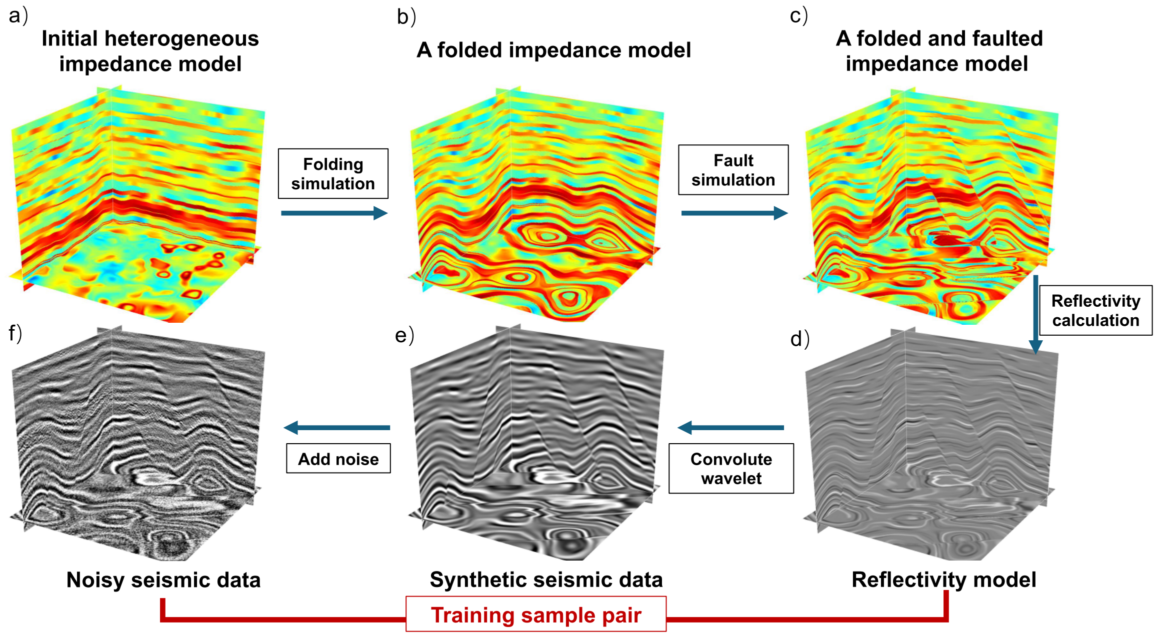

Figure 2. Data generation for the pre-trained model. (a) Initial heterogeneous flat-layer impedance model, (b) impedance model with fold structures, (c) impedance model with fold and fault structures, (d) reflectivity model, (e) synthetic seismic data, (f) noisy synthetic seismic data.

diversity of the data set. The extracted noise is directly derived from field seismic data and therefore inherently retains the band-limited spectral characteristics of the original seismic data.

In the testing of the 3-D pre-trained model, we use the constructed 3-D noisy synthetic seismic data and their corresponding 3-D reflectivity labels, as shown in Fig. 1(d_{syn} and r_{syn}), as training sample pairs.

2.1.3 Supervised pre-training with synthetic data

For the pre-training model, we synthesize 400 pairs of training samples with dimensions of 256 × 256 × 256. We apply data augmentation by randomly rotating these samples. During training, 80 per cent of the noisy synthetic seismic data is used to train the network, 10 per cent serves as the validation set and the remaining 10 per cent is used as the test set. Before training, we normalize the input seismic data and their corresponding labels to have a mean of 0 and a variance of 1. This normalization aims to standardize data of different categories to the same scale, facilitating subsequent data processing, accelerating the convergence rate of gradient descent and potentially improving the model's accuracy. The mathematical expression for mean–variance normalization is as follows

$$\mathrm{X}_{\mathrm{scale}} = \frac{x - \mu}{s}, \tag{1}$$

where $\mathrm{X}_{\mathrm{scale}}$ is the normalized data, μ is the mean and s is the standard deviation.

During model training, we use the adaptive moment estimation (Adam) optimizer with an initial learning rate of 1×10^{-4}. We also implement a dynamic learning rate adjustment strategy: If the validation loss did not decrease for 5 epochs, the learning rate was reduced to 80 per cent of its current value. Our goal is to input noisy synthetic seismic data and, through the constructed network, predict a 3-D reflectivity data set that fits our labels.

During our network training, we select the following loss functions: Mean Squared Error (MSE), Multi-scale structural similarity loss (MS-SSIM) and sparsity loss functions.

MSE is a common loss function in regression problems, and its formula can be defined as

$$L_{\mathrm{MSE}} = \frac{\sum_{i=1}^{n} (\hat{y}_i - y_i)^2}{N}, \tag{2}$$

where N is the number of points, $\hat{y}_i$ is the predicted reflectivity and y_i is the real reflectivity.

MS-SSIM (Z. Wang *et al.* 2003) can better preserve high frequency detailed geometric features by adaptively assigning higher weights at structural boundaries with large local contrast differences, and it helps to optimize our inversion model to better fit the real reflectivity.

Reflectivity is commonly used to delineate the boundaries between different subsurface layers and is typically sparsely distributed. Therefore, we introduced a sparsity constraint by calculating the L1 loss between the predicted reflectivity and zero, essentially the mean of the reflectivity, to achieve this constraint. This approach enhances the stability and reliability of the inversion results.

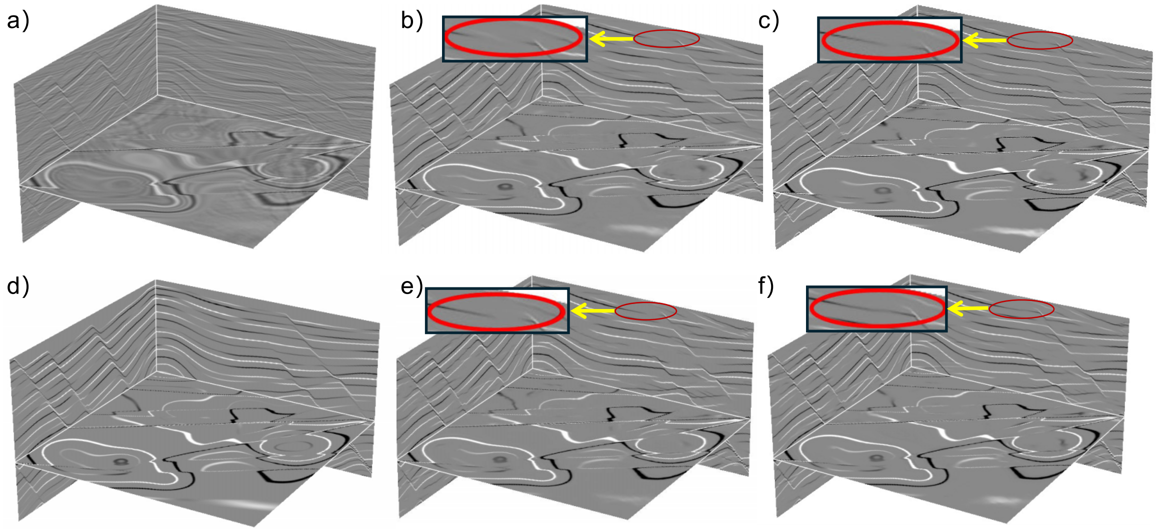

Figure 3. 3-D synthetic data testing for the pre-trained model. (a) Input synthetic seismic data. (b) Result from the model pre-trained with MSE loss. (c) Result from the model pre-trained with MS-SSIM loss. (d) True reflectivity labels. (e) Result from the model pre-trained with a combination of MSE and MS-SSIM losses. (f) Result from the model pre-trained with a combination of MSE, MS-SSIM and sparsity losses.

The sparsity loss function can be defined as

$$L_{\text{sparsity}} = \text{mean}(|r|), \tag{3}$$

where r is the predicted reflectivity.

To more comprehensively evaluate the similarity between the reflectivity and the label data, we combined three loss functions. The expression for the total loss function is as follows

$$L_{\text{pre-train}} = \lambda_1 L_{\text{MSE}} + \lambda_2 L_{\text{MS-SSIM}} + \lambda_3 L_{\text{sparsity}}, \tag{4}$$

where λ_1, λ_2 and λ_3 are the weights of the three loss functions, respectively.

By combining these loss functions, we can fully leverage their respective advantages, simultaneously considering pixel-level differences, structural similarity and sparsity during the optimization process. This comprehensive approach to using loss functions can yield more reasonable and high-resolution inversion results. However, the weights and combinations of these three loss functions need to be adjusted according to specific problems and requirements to achieve more reasonable and reliable inversion results. In this study, the three parameters $\lambda_1 = 1$, $\lambda_2 = 5$ and $\lambda_3 = 1$ were determined empirically. The relatively larger weight assigned to the MS-SSIM term (λ_2) is intended to place greater emphasis on structural similarity during optimization. Compared with pixel-wise loss functions, MS-SSIM is more sensitive to structural continuity and multiscale features, which are important for maintaining structural coherence in the predicted reflectivity.

Through this pre-trained model, initial estimates of reflectivity can be obtained to aid in fine-tuning the model for more reasonable reflectivity estimation. Next, we will introduce the fine-tuning process and provide a detailed comparison of reflectivity results obtained using different methods.

2.1.4 Synthetic reflectivity estimation with pre-trained model

First, we test and validate the feasibility of our pre-trained model on synthetic data and two sets of field seismic data, A and B. The experimental results for the synthetic data are shown in Fig. 3. Fig. 3(a) displays the network input (synthetic seismic data with added noise). Fig. 3(b) shows the reflectivity obtained using L_{MSE}, while Fig. 3(c) presents the results using only the $L_{\text{MS-SSIM}}$. Fig. 3(d) shows the reflectivity labels, and Fig. 3(e) depicts the results using a combination of L_{MSE} and $L_{\text{MS-SSIM}}$. Finally, Fig. 3(f) displays the results using a combination of L_{MSE}, $L_{\text{MS-SSIM}}$ and L_{sparsity}. Our goal is for the predicted reflectivity to match the reflectivity labels in Fig. 3(d). It can be observed that the results are generally consistent with the labels, but the effects vary under different loss functions. In the black box (enlarged in the red circle) in Fig. 3, we find that the reflectivity obtained using the combined loss function $L_{\text{pre-train}}$ shows clear stratigraphic structures and more continuity. In contrast, the stratigraphy in the results from other loss functions exhibits gaps and discontinuity. Next, we apply the model to the two sets of field seismic data, as shown in Fig. 4. The results indicate that the overall resolution improves when the sparsity loss function is included, demonstrating that the added sparsity loss function positively impacts the reflectivity estimation results.

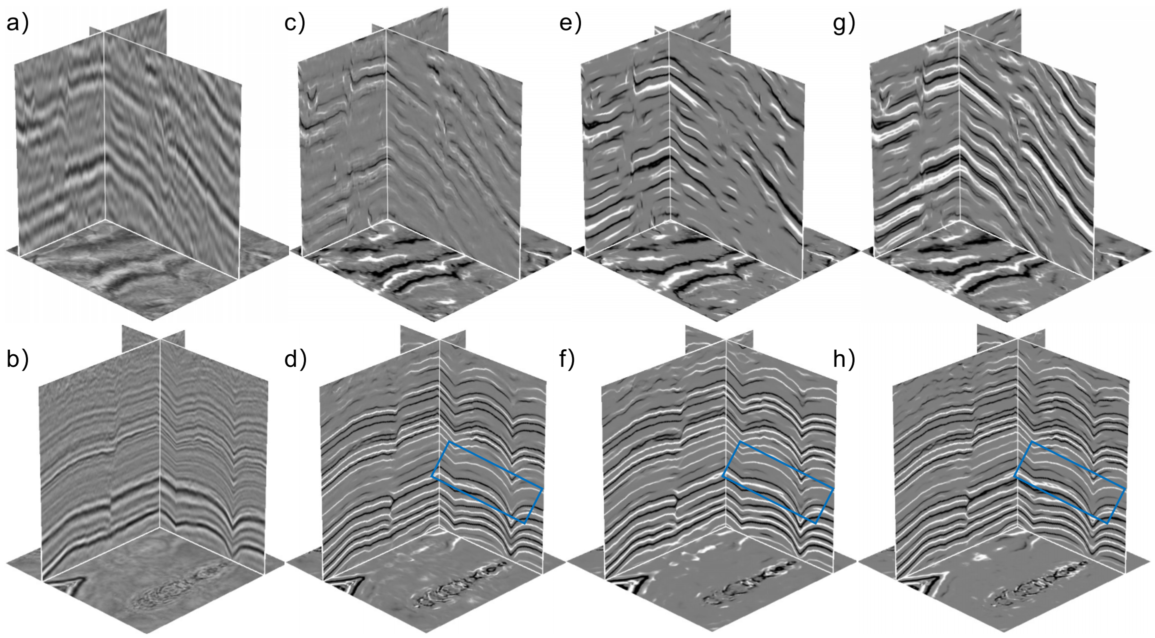

Figure 4. 3-D field data testing for the pre-trained model on data sets A (first row) and B (second row). Panels (a) and (b) are the input field seismic data. The results obtained using different loss functions are shown as follows: (c) and (d) with MSE, (e) and (f) with a combination of MSE and MS-SSIM and (g) and (h) with a combination of MSE, MS-SSIM and sparsity.

2.2 Physics-guided fine-tuning

Considering that the pre-trained model is trained on synthetic data, it cannot fully simulate the distribution and characteristics of field seismic data and lacks generalization ability when applied to field seismic data. To address this issue, we propose a physics-guided fine-tuning approach and introduce appropriate physical constraints to further enhance the network's performance and generalization ability on the field seismic data. Specifically, we incorporate sparsity and structure-oriented smoothness as geophysical constraints to guide the self-supervised learning process in producing reasonable reflectivity models that align with geophysical principles.

The fine-tuning of the model employs a self-supervised learning approach to further optimize the pre-trained model (as depicted in the lower part of Fig. 1). The network architecture used in fine-tuning is identical to the one mentioned earlier for the pre-trained model. The input to the network is the observed field seismic data d_{obs}. After passing through the network similar to the pre-trained model, it produces reflectivity r. This reflectivity is convolved with the wavelet $\mathbf{w}$ to reconstruct seismic data d_{recons}. Subsequently, the reconstruction loss function L_{recons} is computed between d_{recons} and d_{obs} to minimize the difference between them. This process helps the model learn the correspondence between seismic reflectivity and seismic data. Continuously refining the model parameters to ensure that the convolved reconstructed seismic data closely fits the observed seismic data, thereby enhancing the quality of the inversion results. Our self-supervised process is similar to the method proposed by X. Chai *et al.* (2023), but we provide an initial reflectivity model through pre-training, giving the self-supervised process a better starting point. Additionally, we introduce more physical constraints as supervisory information to guide the self-supervised process in continually fine-tuning the network. By inputting field seismic data into the fine-tuned network, we obtain a reasonable reflectivity model that aligns with geophysical understanding.

2.2.1 Physics-guided loss for fine-tuning

Compared to the three loss functions (L_{MSE}, $L_{\mathrm{MS\text{-}SSIM}}$, L_{sparsity}) used in the pre-trained model, we have additionally incorporated the structure-oriented smoothness loss L_{str} as part of the loss function to enhance the lateral continuity of the reflectivity results. This ensures that the structure-oriented smoothness of reflectivity remains consistent with the structure-oriented smoothness of the original seismic data. Integrating structure-oriented smoothness constraints with other conditions helps ensure that the inversion model's structure aligns with the directional features of the subsurface geology, enhancing the geological consistency and interpretability of the inversion results. From d_{obs}, we estimate the 3-D structure tensor (X. Wu 2017) of the seismic data volume, which captures the local geometric and structural information of the seismic data. The structure tensor is derived by calculating the local gradient of the seismic data and forming a 3×3 gradient outer product matrix. In our implementation, the structure tensor was computed using a Gaussian-smoothed local orientation analysis with parameters (6, 2, 2). Through eigen-decomposition, three orthogonal eigenvectors, denoted as the normal component $\mathbf{u}$ and the structural direction components $\mathbf{v}$ and $\mathbf{w}$, are obtained. The normal component $\mathbf{u}$ is perpendicular to linear or planar features such as seismic reflectors, representing the normal direction of the local structure. In contrast, the components $\mathbf{v}$ and $\mathbf{w}$ lie within the planes of planar features, with $\mathbf{v}$ being orthogonal to stratigraphic features and $\mathbf{w}$ parallel to them. This decomposition enables us to effectively characterize and analyse subsurface structural features. Utilizing these components, we construct two distinct loss functions: $L_{\text{str}}(\mathbf{u})$, which enforces smoothness in the normal direction, and $L_{\text{str}}(\mathbf{v}, \mathbf{w})$, which regularizes smoothness along the structural directions within the planes. The combination of these loss functions forms the structure-oriented smoothness loss function L_{str}, which enhances the model's capability to handle noisy environments while preserving critical structural details in the seismic data. To further ensure that the imposed smoothness does not smear structural discontinuities, the structure-oriented smoothness constraint is designed to be spatially adaptive. Specifically, the strength of the smoothing is modulated based on seismic discontinuity attributes—regions with strong discontinuities, such as faults or unconformities, are assigned weaker smoothing, while more continuous areas receive stronger regularization. This spatially variant design enables the model to better preserve sharp geological features while still promoting lateral continuity in stratified zones. Consequently, the proposed constraint enhances the robustness of the inversion to noise while reducing the risk of oversmoothing critical structural details that are important for geological interpretation.

$$L_{\text{str}}(\mathbf{u}) = 1 - \frac{\left(\mathbf{g}^{\mathrm{T}}\mathbf{u}\right)^2}{\left(\|\mathbf{g}\|_2^2 + \epsilon\right)\left(\|\mathbf{u}\|_2^2 + \epsilon\right)}, \tag{5}$$

$$L_{\text{str}}(\mathbf{v}, \mathbf{w}) = \text{mean}(|\mathbf{g} \cdot \mathbf{v}| + |\mathbf{g} \cdot \mathbf{w}|). \tag{6}$$

In the above equations, $\mathbf{g}$ represents the gradient of the predicted reflectivity r. The normal component loss function $L_{\text{str}}(\mathbf{u})$ ensures that the model's predictions are consistent with the actual geological structural features by making the normal components and gradients as similar as possible. This enables the model to better capture the continuity and consistency of geological structures. The structural component loss function $L_{\text{str}}(\mathbf{v}, \mathbf{w})$ ensures the balance of predicted reflectivity in different structural directions by evaluating their perpendicularity. These two loss functions incorporate stratigraphic information, guiding the gradient direction in reflectivity estimation, enabling the model to more reasonably reflect subsurface geological structures. Although the calculation of seismic data directional information is required in advance, its computational cost is relatively low, as the structure tensor only involves local gradient smoothing and does not require extensive iterative computations.

After adding structure-oriented smoothness loss, the formula is as follows

$$L_{\text{fine-tune}} = \alpha_1 L_{\text{MSE}} + \alpha_2 L_{\text{MS-SSIM}} + \alpha_3 L_{\text{sparsity}} + \alpha_4 L_{\text{str}}, \tag{7}$$

where α_1, α_2, α_3 and α_4 are the weights of the four loss functions, respectively.

It is important to note that the inputs for L_{MSE} and $L_{\text{MS-SSIM}}$ here are the reconstructed seismic data and the field seismic data, whereas in the previous section, the inputs for $L_{\text{pre-train}}$ in L_{MSE} and $L_{\text{MS-SSIM}}$ are the predicted reflectivity and the reflectivity labels.

The structure-oriented smoothness loss function proves particularly effective on seismic data with high noise levels. Noise can notably accentuate the directional differences between the inversion results and the seismic data, introducing additional uncertainty and increasing the instability of the inversion results. By incorporating the structure-oriented smoothness loss function, we can effectively constrain the directionality of the inversion model, ensuring it better aligns with the seismic data's direction. This reduces the impact of noise on the inversion results and can partially denoise the data. When dealing with noisy data, this method effectively suppresses noise, enhancing the stability and accuracy of the inversion results, which is particularly evident in field seismic data. We provide a detailed introduction in the subsequent field seismic data test.

Since more iterations are not always better, we do not rely on the conventional method of plotting the loss curve and stopping training once it stabilizes and converges. Instead, we choose to assess the training process by evaluating the match between the predicted reflectivity and the true labels, with well-log data often serving as the labels in field data. As shown in eq. (8), we compute the overall average of the cross-correlation coefficients between the predicted reflectivity and the well-log data at all well-log locations. The formula is as follows

$$\text{corr} = \frac{\sum_j x_j y_j}{\sqrt{\sum_j x_j^2 \sum_j y_j^2}}, \tag{8}$$

where $\mathbf{x}$ is the log label and $\mathbf{y}$ is the forecast result.

When this average reaches a relative maximum, we consider the reflectivity to be the most reasonable. The main reason for this approach is that real-world seismic data, which serves as network input, contains substantial noise. Similar to the Deep Image Prior (D. Ulyanov *et al.* 2018) training process, the model first fits the reflectivity and then fits the noise, which may lead to overfitting the noise—an undesirable outcome. Therefore, using the cross-correlation coefficient as the evaluation metric for model selection offers a practical and effective stopping criterion for unsupervised training. If this metric is considered unreliable, it is entirely possible to run more epochs during training and manually select the best result from the intermediate models based on the user's own criteria, such as the convergence behaviour of the loss curve. Although not optimal, this provides a fallback solution that ensures flexibility in practical use.

Next, we conduct comparative tests on synthetic data and two additional sets of field seismic data to analyse the performance of our proposed SRPPF model. This analysis includes the impact of different loss functions, training strategies and noise levels.

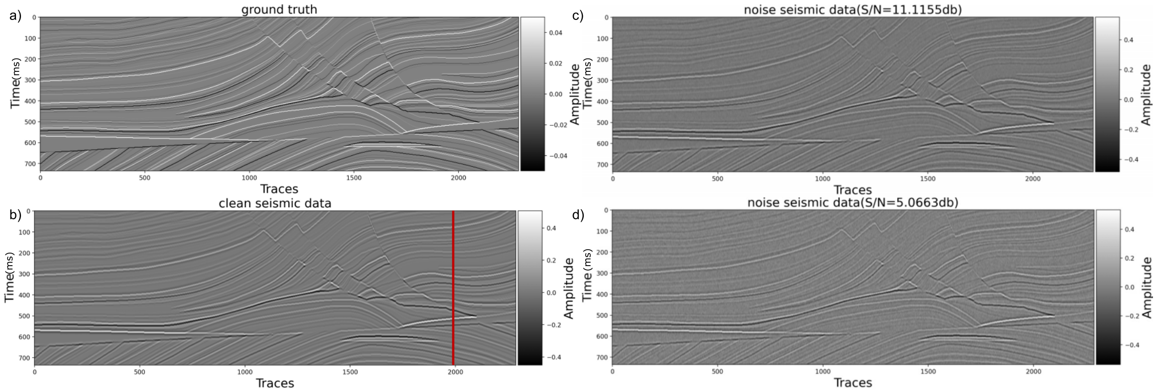

Figure 5. Marmousi test data (a) real reflectivity model, (b) noiseless synthetic seismic data, (c) synthetic seismic data with a signal-to-noise ratio of approximately 11 dB and (d) synthetic seismic data with a signal-to-noise ratio of approximately 5 dB.

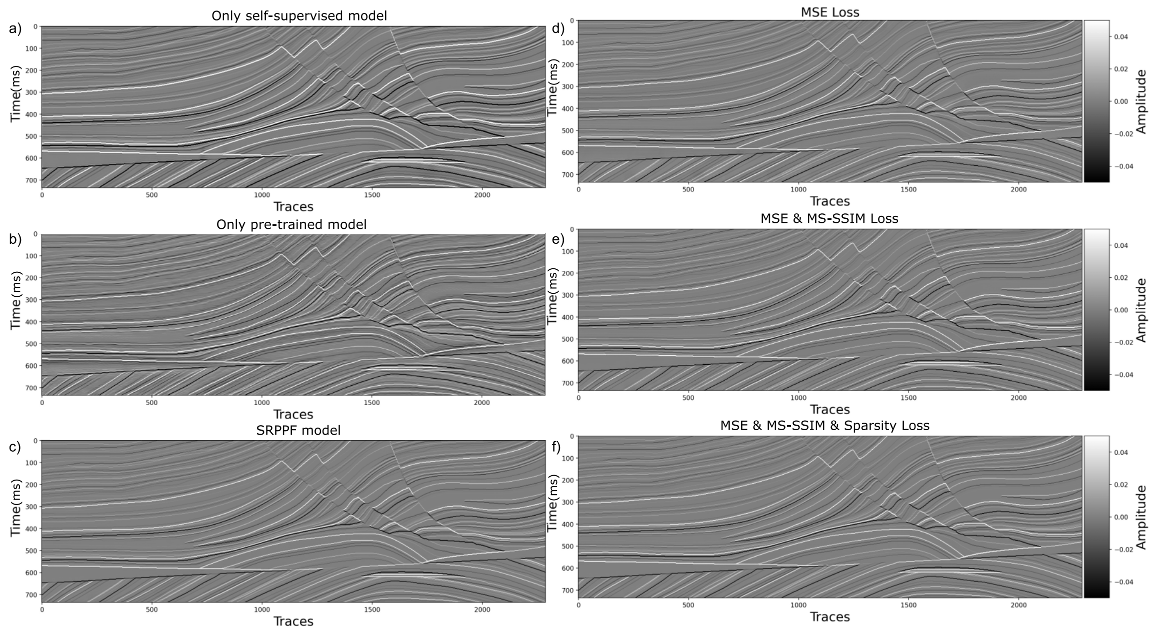

Figure 6. Test results of different training strategies for 2-D synthetic data. Reflectivity predicted by: (a) the self-supervised model, (b) the pre-trained model with synthetic data and (c) our proposed SRPPF method. Using our SRPPF method, reflectivity was predicted with different loss functions: (d) MSE, (e) a combination of MSE and MS-SSIM and (f) a combination of MSE, MS-SSIM and sparsity.

2.2.2 Enhanced reflectivity estimation with fine-tuning

The Marmousi model is a classic geological model. Therefore, we used it as the test data in this study. During testing, we used a 25 Hz Ricker wavelet as the seismic wavelet. By convolving the 25 Hz Ricker wavelet with a 2-D reflectivity profile (shown in Fig. 5a), we obtained a clean, noise-free seismic profile with 2-D dimensions, as depicted in Fig. 5(b). This data set contains 736 time samples and 2288 traces. To assess the model's ability to handle noise and evaluate its robustness and stability against noise, we added different levels of Gaussian noise to the clean seismic data. To ensure the seismic data remains band-limited, we applied filtering to the noise before adding it to the clean seismic data. This resulted in noisy seismic data with signal-to-noise ratios (SNR) of approximately 11 dB (Fig. 5c) and 5 dB (Fig. 5d).

On the synthetic data, we compare the effects of three different training strategies: Using only supervised training with synthetic data (only pre-trained model), using self-supervised training (only self-supervised model), and SRPPF. This comparison illustrates the effectiveness of our SRPPF approach. The difference between using only self-supervised training and the SRPPF approach lies in the presence of pre-training; the self-supervised process remains consistent across both strategies. Figs 6(a)–(c) present the 2-D reflectivity profiles obtained from the three training strategies. It is evident that the reflectivity profile in Fig. 6(c) is the most similar to the ground truth shown in Fig. 5(a). For a clearer comparison, we extract the trace along the red vertical line in Fig. 5(b) and plot it on the left side of Fig. 7. This comparison demonstrates that our method better fits the true reflectivity waveform than the other two methods. Additionally, as shown in Table 1, we calculate the cross-correlation coefficients between the predicted and true reflectivity traces.

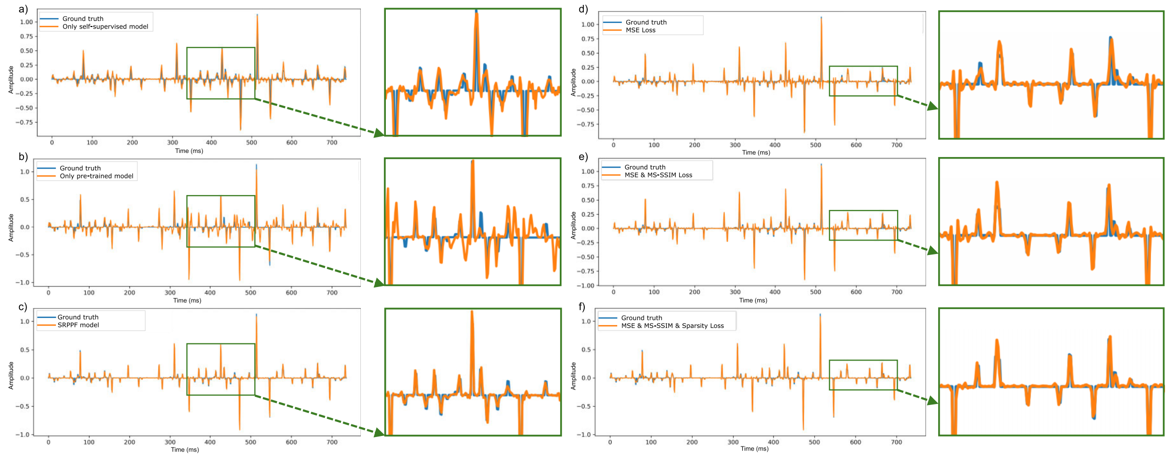

Figure 7. Comparison of predicted 1-D waveforms with the ground truth reflectivity (label), extracted along the red line in Fig. 5: (a) Prediction from the self-supervised model versus ground truth; (b) prediction from the pre-trained model versus ground truth; (c) prediction from SRPPF versus ground truth; (d) SRPPF trained with MSE loss versus ground truth; (e) SRPPF trained with MSE and MS-SSIM losses versus ground truth; (f) SRPPF trained with MSE, MS-SSIM and sparsity losses versus ground truth.

Table 1. Comparison of different training strategies on clean seismic data: cross-correlation and MSE with ground truth.

Method	Only pre-trained model ($L_{\text{MSE}} + L_{\text{MS-SSIM}} + L_{\text{sparsity}}$)	Only self-supervised model ($L_{\text{MSE}} + L_{\text{MS-SSIM}} + L_{\text{sparsity}}$)	SRPPF model ($L_{\text{MSE}} + L_{\text{MS-SSIM}} + L_{\text{sparsity}}$)	SRPPF model ($L_{\text{MSE}} + L_{\text{MS-SSIM}}$)	SRPPF model (L_{MSE})
Correlation coefficient	0.8453	0.8012	**0.9832**	0.9436	0.9385
MSE	0.0052	0.0098	**0.0006**	0.0016	0.0017

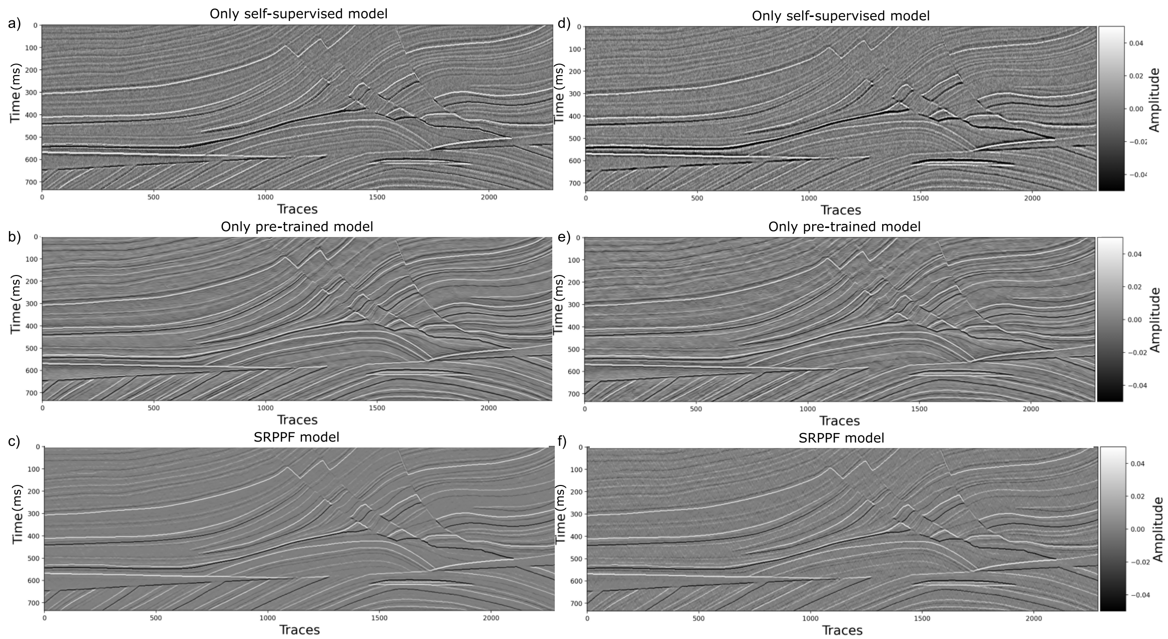

Figure 8. Panels (a)–(c) and (d)–(f) show the test results of different training strategies for 2-D synthetic data with signal-to-noise ratios of approximately 11 and 5 dB, respectively. Reflectivity was predicted using: (a) only the self-supervised model, (b) only the pre-trained model with synthetic data and (c) our proposed SRPPF method. The same applies to panels (d)–(f).

This quantitative comparison further confirms that our SRPPF method achieves the best results, demonstrating accuracy and fidelity in fitting the true reflectivity.

To assess the robustness of our model to noise, we utilize noisy seismic data with SNR of approximately 11 dB (as shown in Fig. 5c) and 5 dB (as shown in Fig. 5d) as inputs. As depicted on the left side of Fig. 8, when a small amount of noise (approximately 11 dB SNR) is introduced, using either the physics-guided fine-tuning model alone (Fig. 8a) or the pre-trained model alone (Fig. 8b) results in lower overall resolution. Conversely, our SRPPF method (Fig. 8c) yields higher resolution outcomes. This advantage is further reflected in the MSE between predicted reflectivity and ground truth: As shown in Table 2, our method consistently achieves the lowest MSE

Table 2. Comparison of different training strategies on noisy seismic data: cross-correlation and MSE with ground truth.

Method	Only pre-trained model	Only self-supervised model	SRPPF model
Correlation coefficient (SNR ≈ 11 dB)	0.8303	0.5292	**0.9302**
Correlation coefficient (SNR ≈ 5 dB)	0.7959	0.4504	**0.9010**
MSE (SNR ≈ 11 dB)	0.0060	0.0167	**0.0015**
MSE (SNR ≈ 5 dB)	0.0061	0.0249	**0.0020**

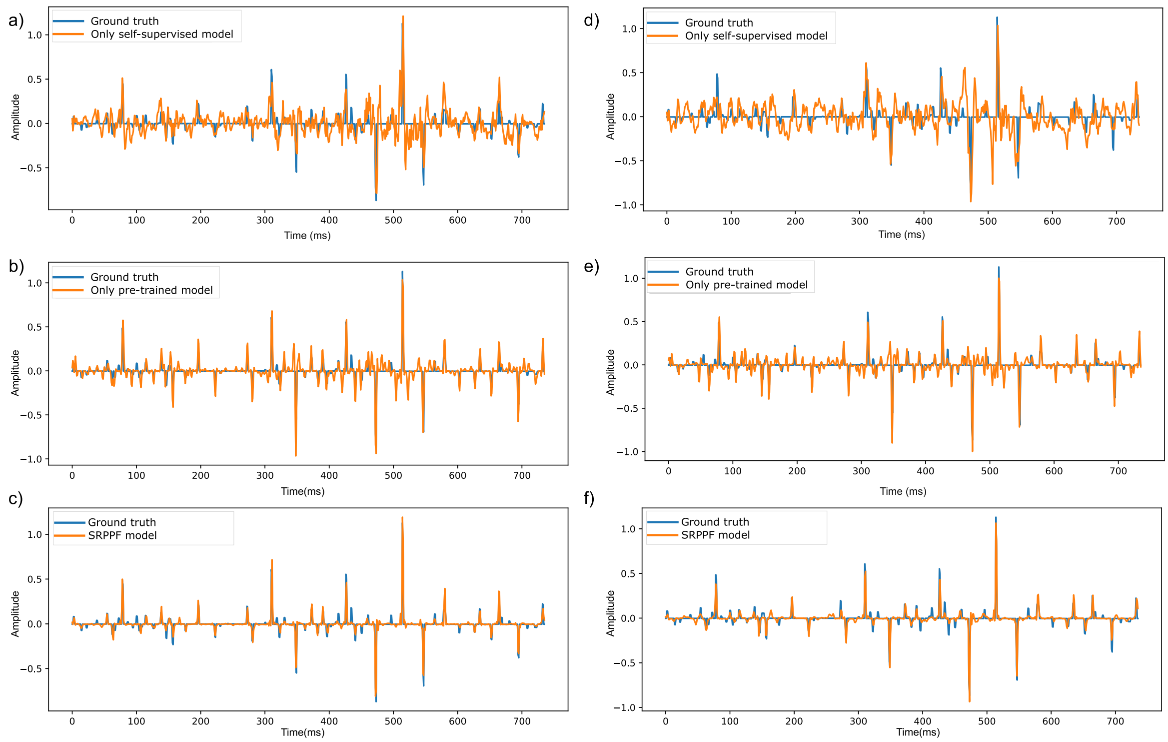

Figure 9. Comparison of predicted 1-D waveforms with the ground truth reflectivity (label), extracted along the red line in Fig. 5 from the results shown in Fig. 8: (a) Self-supervised model for data with SNR ≈ 11 dB versus ground truth; (b) pre-trained model for data with SNR ≈ 11 dB versus ground truth; (c) SRPPF for data with SNR ≈ 11 dB versus ground truth; (d) self-supervised model for data with SNR ≈ 5 dB versus ground truth; (e) pre-trained model for data with SNR ≈ 5 dB versus ground truth; (f) SRPPF for data with SNR ≈ 5 dB versus ground truth.

under both noise levels. Furthermore, as illustrated in Fig. 9, we compared the single-trace seismic waveforms at the position of the red line in Fig. 5(b), demonstrating the strong fit of our SRPPF method to the labels. Even under increased noise levels (approximately 5 dB SNR), our proposed SRPPF method still yields reasonable reflectivity estimates, maintaining high overall resolution and lateral continuity. In contrast, the other two training strategies exhibit noise and detail loss in their results. As shown in Table 2, our SRPPF method achieves correlation coefficients above 0.9 even in high-noise conditions, further highlighting its robustness and ability to provide stable and reasonable reflectivity estimates. These findings underscore the advantage of our model in enhancing seismic data resolution and result reliability when dealing with low SNR data.

To test the accuracy and generalization ability of our method, we select two field seismic data sets that include well-log data. The seismic wavelets for the two field data sets were extracted from well logs using the Hampson–Russell software (D.P. Hampson *et al.* 2005), assuming a constant-phase wavelet over the target interval. The extracted wavelets were directly used in the convolutional reconstruction step during fine-tuning without additional filtering or processing.

Next, we also conduct tests on field data, and the results of the 3-D data we use are shown in Fig. 10. In Fig. 10, we can see that the reflectivity results produced by our proposed SRPPF method (Figs 10g and h), exhibit notably better overall spatial consistency compared to the results obtained using only the pre-training strategy (Figs 10c and d). The key reason is that after pre-training the model, we employ physics-guided fine-tuning, enabling the model to learn the characteristics and distribution of the subsurface structure from the input field seismic data. This allows it to capture the intrinsic relationships and spatial continuity between the data without the need for labelled seismic reflectivity information, thereby enhancing the model's adaptability to various geological conditions and improving its generalization capability.

Furthermore, comparing the reflectivity results of our proposed SRPPF method (Figs 10g and h) with those obtained using self-supervised learning alone (Figs 10e and f), we observe that the SRPPF method achieves higher overall resolution and more clearly depicts fine geological structures (thin layers, small-scale faults as indicated by the yellow arrow and blue box). This is because the

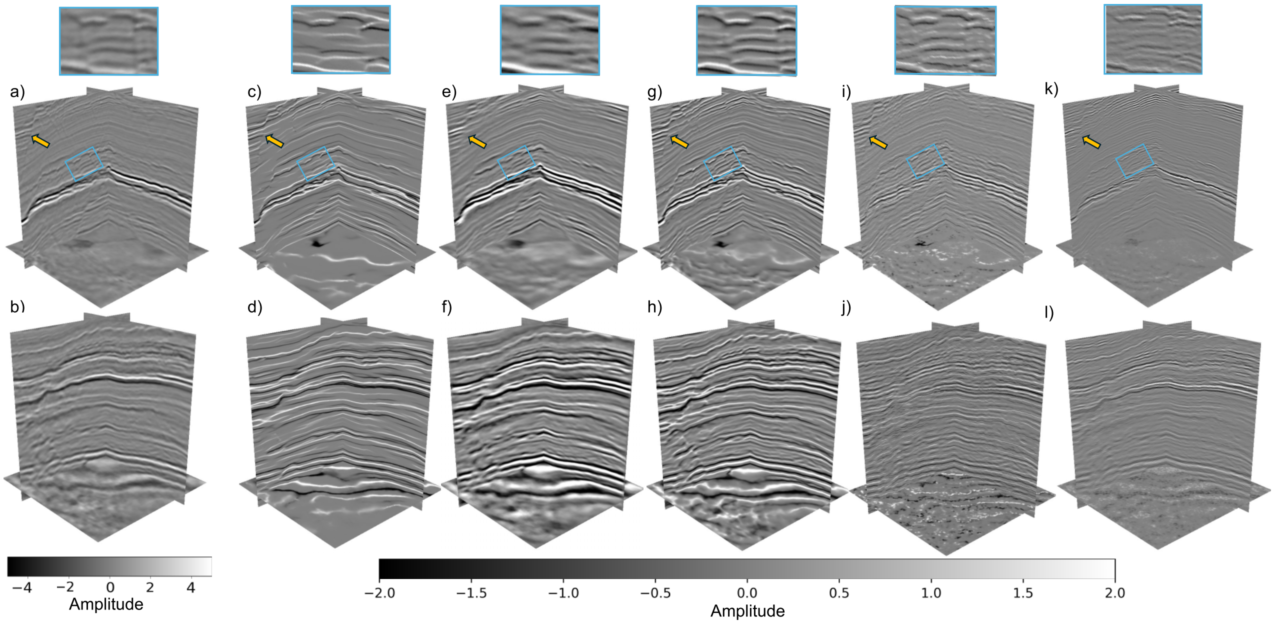

Figure 10. Field data comparison test results. Panels (a) and (b) show the input 3-D seismic data sets C and D, with dimensions of 128 × 128 × 512 along the crossline, inline and time directions. Reflectivity results are displayed as follows: (c) and (d) from the pre-training model alone, (e) and (f) from the self-supervised method alone, (g) and (h) from our proposed SRPPF method, (i) and (j) from the results obtained using the sparse-spike deconvolution method, (k) and (l) from the results obtained using the BPI method.

initial reflectivity model essentially matches the geological structures, which helps to fine-tune the model for more reasonable results, avoiding local minima.

In contrast, the results obtained using a traditional sparse spike deconvolution method (Figs 10i and j), which we implemented based on the general principles of sparse reflectivity inversion, suffer from several drawbacks. While sparse spike deconvolution theoretically aims to enhance resolution by enforcing sparsity in the reflectivity model, its results exhibit poor spatial consistency, as evident from the disrupted continuity of reflectors and the loss of fine structural details. Additionally, it amplifies random noise, introducing artefacts that obscure weak reflectors, particularly in deeper layers. These issues can be attributed to the method's sensitivity to noise, reliance on wavelet estimation and the strong sparsity constraints that may suppress weak but geologically notable signals. For a complete comparison, we also include results from basis pursuit inversion (BPI), as shown in Figs 10(k) and (l). It can be observed that BPI is able to reasonably enhance the main reflection events and provide an overall plausible reflectivity distribution for the field data. At the same time, we note that in deeper regions with relatively low signal-to-noise ratio, the continuity of some reflection in-phase axes is slightly weakened, and the distinguishability of weak-amplitude features is reduced at certain locations. In contrast, three deep learning methods, particularly our SRPPF method, achieve superior spatial continuity, better preservation of weak reflectors and higher robustness to noise.

As shown in Fig. 11, we also compare the reconstructed seismic data obtained from different methods. The reconstructed seismic data are generated by convolving the estimated reflectivity with the wavelet. Our SRPPF method produces reconstructed seismic data that most closely matches the input seismic data, demonstrating its effectiveness in capturing and reconstructing field seismic data.

To further validate the efficacy of the structure-oriented smoothness loss function, we conduct ablation experiments on field data. Synthetic data, being overly simplistic, do not adequately demonstrate the benefits of the structure-oriented smoothness loss function. Thus, we utilize field seismic data to illustrate its impact. For these ablation experiments, we keep training details, strategies and methods constant, with the only variable being the inclusion or exclusion of the structure-oriented smoothness loss function. The overall results improved notably when the structure-oriented smoothness loss function was included (Figs 10g and h), compared to the results without it (Figs 13a and c). The incorporation of structure-oriented smoothness loss functions effectively reduces noise interference and enhances the resolution of the data. This improvement is from the function's ability to guide the model in better preserving structure-oriented smoothness within the data. This is particularly evident in field data sets with strong noise, where the structure-oriented smoothness loss function proves to be more beneficial. Our approach, which combines pre-training with physics-guided fine-tuning, achieves its peak performance with only 10 iterations of evaluation metrics. In contrast, the self-supervised method alone requires 100 iterations to reach its peak. This highlights the importance of pre-training models.

To verify the accuracy of the predicted reflectivity, we extract a trace from the predicted reflectivity of the 3-D seismic data sets C, D, and compare it with the well-log data. Given that the resolution of the well-log impedance data is notably higher than that of seismic data, we first resample the impedance data to match the sampling interval of the seismic data. This is achieved by applying median filtering to the well-log impedance data to reduce its resolution to the seismic scale. Next, we calculate the corresponding well-log

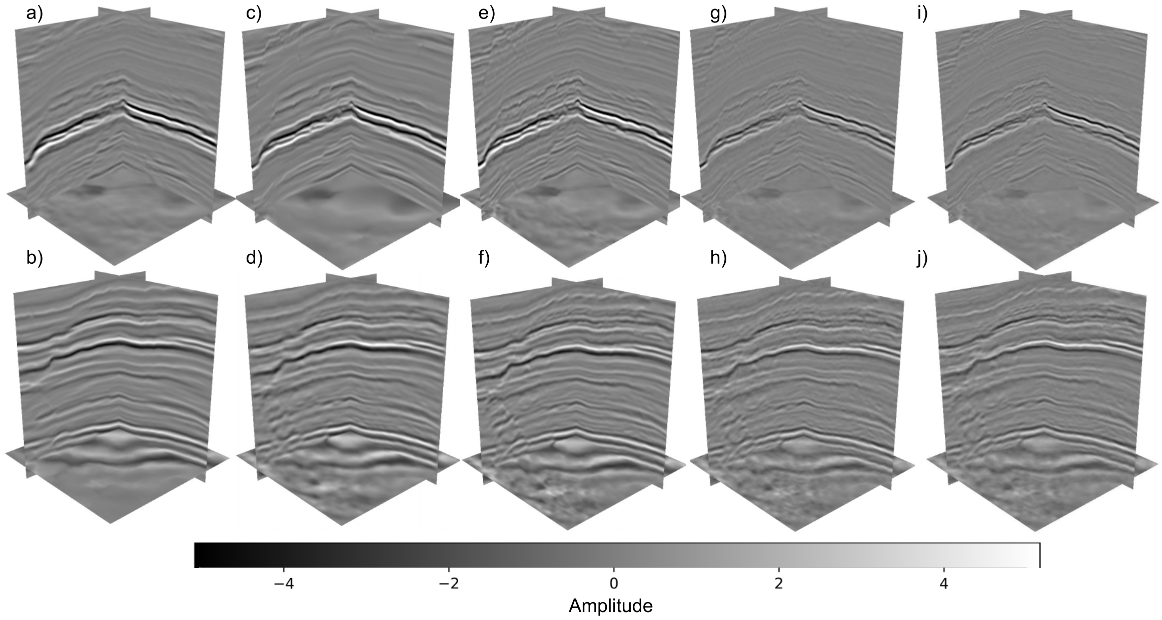

Figure 11. Reconstructed seismic data corresponding to the reflectivity results in Figs 10(c)–(l). Panels (a) and (b) are from the pre-training model alone, (c) and (d) from the self-supervised method alone, (e) and (f) from our proposed SRPPF method and (g) and (h) are obtained using the sparse-spike deconvolution method, (i) and (j) are obtained using the BPI method.

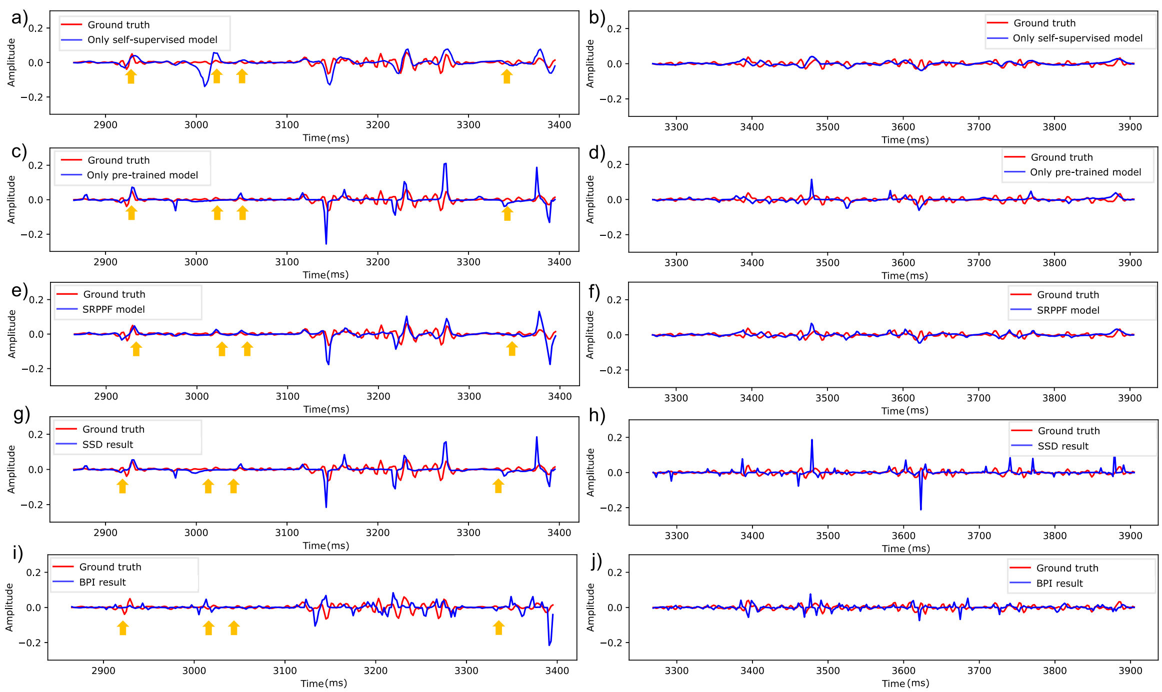

Figure 12. Comparison of real well-log data (red) with reflectivity predicted by different methods (blue) for 1-D trace analysis of 3-D field seismic data C and D. Panels (a) and (b) show the reflectivity from the self-supervised method. Panels (c) and (d) display the reflectivity from the pre-training model. Panels (e) and (f) present the reflectivity from our proposed SRPPF method. Panels (g) and (h) show the reflectivity from the sparse spike deconvolution method. Panels (i) and (j) show the reflectivity from the basis pursuit inversion method, all compared with well-log data.

reflectivity from the resampled impedance data. To further align the well-log reflectivity with the seismic data bandwidth, we apply a 200 Hz low-pass filter to the reflectivity, ensuring that its frequency range is higher than the estimated reflectivity, serving as a reference for comparison. As shown in Fig. 12, the blue curve represents the trace extracted from the predicted 3-D reflectivity at the well-log position. We observe that the blue curve and the red curve exhibit similar trends, and at the locations indicated by the yellow arrows, our proposed SRPPF method shows better alignment with the well-log data. In summary, we propose that the SRPPF method achieves more reasonable reflectivity estimation. This improved accuracy and resolution will facilitate more detailed reservoir characterization.

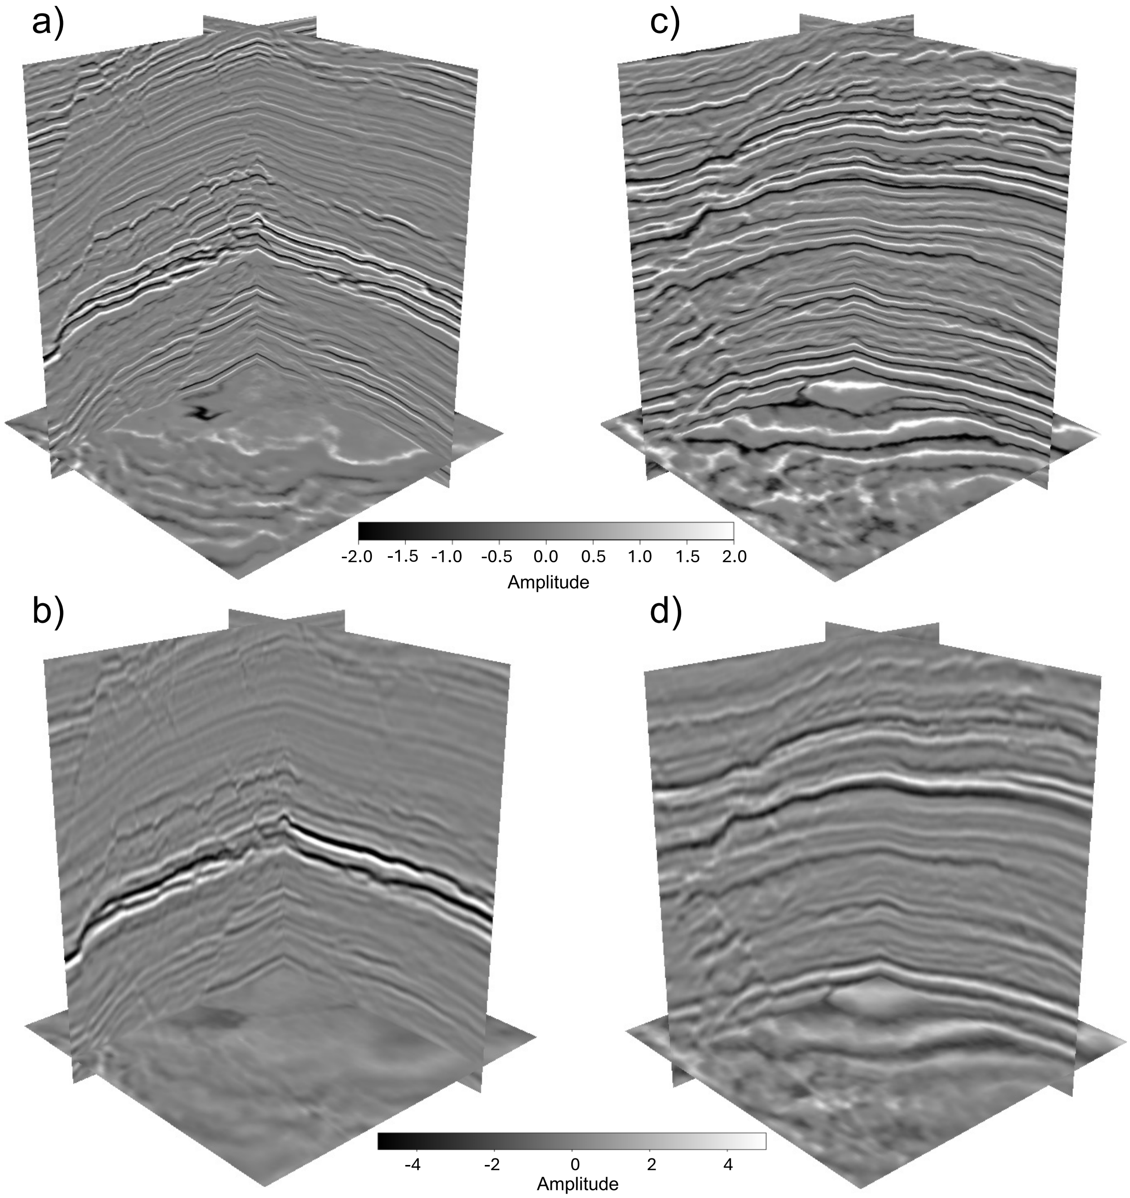

Figure 13. Panels (a) and (b) present the reflectivity and corresponding reconstructed seismic data obtained from input seismic volume C, respectively. Panels (c) and (d) show the reflectivity and reconstructed seismic data derived from input seismic volume D, respectively.

3 DISCUSSION

The proposed framework couples supervised pre-training on labelled synthetic data with physics-guided self-supervised fine-tuning on field seismic data. The two stages play distinct but complementary roles. Pre-training on a large set of synthetic data with known reflectivity labels enables the network to learn the basic mapping from seismic amplitudes to sparse, high-resolution reflectivity and to capture characteristic stratigraphic patterns such as folds, faults and layered structures. The resulting pre-trained model provides a structurally reasonable initial reflectivity model and mitigates the convergence instability that commonly arises when self-supervised inversion is initialized from random weights. Fine-tuning then adapts the network to the specific wavelet signatures, noise statistics and structural styles of each field survey, which inevitably differ from those represented in the synthetic training set.

During fine-tuning, three geophysical constraints jointly regularize the inversion. The sparsity constraint enforces the impulsive nature of the reflectivity series and highlights primary reflection interfaces; the structure-oriented smoothness constraint, guided by the local orientation of seismic reflectors, preserves lateral continuity and dip coherence while spatially down-weighting smoothing across faults and unconformities; and the data reconstruction loss ensures that the predicted reflectivity, when convolved with the wavelet, reproduces the observed seismic traces. These constraints translate well-established geophysical priors into differentiable loss terms, allowing them to be optimized jointly through a single end-to-end network. Compared with purely data-driven or purely self-supervised approaches, this design improves both the geological plausibility and the data fidelity of the estimated reflectivity.

A further advantage of the proposed framework lies in its flexibility. Because all constraints are imposed through loss functions and optimized via automatic differentiation, additional physical priors can be readily incorporated without redesigning the inversion procedure. For example, non-stationary wavelets, Q-compensated forward modelling, or more sophisticated structural and stratigraphic constraints could be added in future work to further enhance the framework's adaptability to complex geological settings.

3.1 Limitations

Despite these advantages, two limitations of the current framework should be acknowledged. First, the synthetic 3-D data sets used in pre-training remain simplified relative to real seismic data, with restricted structural diversity, Ricker wavelets, simplified noise models and an acoustic-based procedure for generating reflectivity. These simplifications inevitably limit the standalone generalizability of the pre-trained model, which is precisely the motivation for introducing the subsequent fine-tuning stage. Future synthetic data sets with higher geological realism and more complete forward modelling are expected to further expand the standalone applicability of the pre-trained model.

Secondly, the fine-tuned model is inherently survey-specific, since differences in noise statistics, wavelet properties and structural styles across surveys generally require tailored choices of fine-tuning hyperparameters and constraint weights. Future work could introduce adaptive or data-driven mechanisms for selecting these weights to improve robustness across surveys, and incorporate more sophisticated wavelet modelling—such as non-stationary or Q-compensated wavelets—into the fine-tuning stage to further enhance inversion performance.

4 CONCLUSION

This paper presents SRPPF, a two-stage reflectivity inversion framework that couples supervised pre-training on labelled synthetic data with physics-guided self-supervised fine-tuning on field seismic data. The pre-training stage provides a stable and structurally reasonable initial reflectivity model, while the fine-tuning stage incorporates geophysical constraints—namely data reconstruction, reflectivity sparsity and structure-oriented smoothness—to adapt the network to survey-specific wavelet signatures, noise statistics and structural styles. Experiments on synthetic and field data sets demonstrate that the proposed framework produces reliable, high-resolution reflectivity estimates that are consistent with both observed seismic data and geophysical priors, providing a practical tool for seismic interpretation, thin-layer characterization and subsurface reservoir analysis.

ACKNOWLEDGMENTS

This research was financially supported by the Fundamental and Interdisciplinary Disciplines Breakthrough Plan of the Ministry of Education of China (JYB2025XDXM301) , the China Postdoctoral Science Foundation (grant no. 2026M790497) and the Postdoctoral Fellowship Program of CPSF (grant no. GZB20260111).

AUTHOR CONTRIBUTIONS

Y.W. designed the study, implemented the network and conducted the experiments; J.L. and X.S. contributed to data preparation and analysis; X.W. supervised the project and revised the manuscript. All authors discussed the results and contributed to the final manuscript.

CONFLICT OF INTEREST

The authors declare no conflicts of interest.

DATA AVAILABILITY

The synthetic 3-D seismic data sets and the corresponding reflectivity labels used in the pre-training stage were generated following the procedure described in Section 2. The implementation of the proposed SRPPF framework is publicly available on GitHub.[1]

The field seismic data sets C and D are proprietary and cannot be publicly shared due to confidentiality agreements with the data providers. Access to these data sets should be requested directly from the respective data owners.

[1] https://github.com/Wytclean/SRPPF-reflectivity-estimation

REFERENCES

Aki, K. & Richards, P.G., 2002. *Quantitative Seismology,* 2nd edn. University Science Books, Sausalito, CA.

Berkhout, A., 1977. Least-squares inverse filtering and wavelet deconvolution, *Geophysics,* **42**(7), 1369–1383.

Biswas, R., Sen, M.K., Das, V. & Mukerji, T., 2019. Prestack and poststack inversion using a physics-guided convolutional neural network, *Interpretation,* **7**(3), SE161–SE174.

Brandolin, F., Ravasi, M. & Alkhalifah, T., 2024. Pinnslope: seismic data interpolation and local slope estimation with physics informed neural networks, *Geophysics,* **89**(4), V331–V345.

Chai, X., Tang, G., Lin, K., Yan, Z., Gu, H., Peng, R., Sun, X. & Cao, W., 2021. Deep learning for multitrace sparse-spike deconvolution, *Geophysics,* **86**(3), V207–V218.

Chai, X., Yang, T., Gu, H., Tang, G., Cao, W. & Wang, Y., 2023. Geophysics-steered self-supervised learning for deconvolution, *Geophys. J. Int.,* **234**(1), 40–55.

Chen, H., Sacchi, M.D., Haghshenas Lari, H., Gao, J. & Jiang, X., 2023. Nonstationary seismic reflectivity inversion based on prior-engaged semisupervised deep learning method, *Geophysics,* **88**(1), WA115–WA128.

Datta, D. & Sen, M.K., 2016. Estimating a starting model for full-waveform inversion using a global optimization method, *Geophysics,* **81**(4), R211–R223.

Ganin, Y. & Lempitsky, V., 2015. Unsupervised domain adaptation by backpropagation, *International Conference on Machine Learning,* pp. 1180–1189, PMLR.

Gao, H., Wu, X. & Liu, G., 2021a. Channelseg3d: channel simulation and deep learning for channel interpretation in 3D seismic images, *Geophysics,* **86**(4), IM73–IM83.

Gao, Z., Hu, S., Li, C., Chen, H., Jiang, X., Pan, Z., Gao, J. & Xu, Z., 2021b. A deep-learning-based generalized convolutional model for seismic data and its application in seismic deconvolution, *IEEE Trans. Geosci. Remote Sens.,* **60,** 1–17.

Guitton, A., Ayeni, G. & Díaz, E., 2012. Constrained full-waveform inversion by model reparameterization, *Geophysics,* **77**(2), R117–R127.

Hale, D., 2009. Structure-oriented smoothing and semblance. CWP Report 635, *Colorado School of Mines,* 261–270.

Hampson, D.P., Russell, B.H. & Bankhead, B., 2005. Simultaneous inversion of pre-stack seismic data, *SEG Technical Program Expanded Abstracts 2005,* pp. 1633–1637, Society of Exploration Geophysicists.

Huang, H., Wang, T., Cheng, J., Xiong, Y., Wang, C. & Geng, J., 2022. Self-supervised deep learning to reconstruct seismic data with consecutively missing traces, *IEEE Trans. Geosci. Remote Sens.,* **60,** 1–14.

Kaaresen, K.F. & Taxt, T., 1998. Multichannel blind deconvolution of seismic signals, *Geophysics,* **63**(6), 2093–2107.

Kazemi, N., Bongajum, E. & Sacchi, M.D., 2016. Surface-consistent sparse multichannel blind deconvolution of seismic signals, *IEEE Trans. Geosci. Remote Sens.,* **54**(6), 3200–3207.

Levy, S. & Fullagar, P.K., 1981. Reconstruction of a sparse spike train from a portion of its spectrum and application to high-resolution deconvolution, *Geophysics,* **46**(9), 1235–1243.

Li, J., Wu, X. & Hu, Z., 2021. Deep learning for simultaneous seismic image super-resolution and denoising, *IEEE Trans. Geosci. Remote Sens.,* **60,** 1–11.

Li, J., Wu, X., Zhang, X., Du, X., Sun, X., Deng, B. & Wang, G., 2026. High-fidelity seismic super-resolution using prior-informed deep learning with 3D awareness, *IEEE Trans. Image Process,* **35,** 2152–2166.

Liang, L. & Castagna, J.P., 2017. Resolution limits of seismic inversion: An uncertainty analysis, *Geophysics,* **82**(6), R301–R315..

Liu, C., Sun, M., Dai, N., Wu, W., Wei, Y., Guo, M. & Fu, H., 2022. Deep learning-based point-spread function deconvolution for migration image deblurring, *Geophysics,* **87**(4), S249–S265.

Long, M., Cao, Y., Wang, J. & Jordan, M.I., 2015. Learning transferable features with deep adaptation networks, *International Conference on Machine Learning,* pp. 97–105, PMLR.

Ma, M., Wang, S., Yuan, S., Wang, J. & Wen, J., 2017. Multichannel spatially correlated reflectivity inversion using block sparse bayesian learning, *Geophysics,* **82**(4), V191–V199.

Margrave, G.F., Lamoureux, M.P. & Henley, D.C., 2011. Gabor deconvolution: estimating reflectivity by nonstationary deconvolution of seismic data, *Geophysics,* **76**(3), W15–W30.

O'doherty, R. & Anstey, N.A., 1971. Reflections on amplitudes, *Geophys. Prospect.,* **19**(3), 430–458.

Phan, S. & Sen, M.K., 2021. Seismic nonstationary deconvolution with physics-guided autoencoder, *SEG International Exposition and Annual Meeting,* pp. D011S067R002, SEG.

Portniaguine, O. & Castagna, J., 2005. Spectral inversion: lessons from modeling and Boonesville case study, *SEG Technical Program Expanded Abstracts 2005,* pp. 1638–1641, Society of Exploration Geophysicists.

Robinson, E.A., 1957. Predictive decomposition of seismic traces, *Geophysics,* **22**(4), 767–778.

Robinson, E.A., 1967. Predictive decomposition of time series with application to seismic exploration, *Geophysics,* **32**(3), 418–484.

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

Russell, B., Hampson, D. & Bankhead, B., 2006. An inversion primer, *CSEG Recorder,* **31**(2), 96–103.

Saad, O.M., Harsuko, R. & Alkhalifah, T., 2024. SiameseFWI: a deep learning network for enhanced full waveform inversion, *J. Geophys. Res.: Machine Learning and Computation,* **1**(3), e2024JH000227.

Sacchi, M.D., 1997. Reweighting strategies in seismic deconvolution, *Geophys. J. Int.,* **129**(3), 651–656.

Shelhamer, E., Long, J. & Darrell, T., 2016. Fully convolutional networks for semantic segmentation, *IEEE Trans. Pattern Anal. Mach. Intell.,* **39**(4), 640–651.

Sheriff, R.E. & Geldart, L.P., 1995. *Exploration Seismology,* Cambridge University Press.

Shi, Y., Wu, X. & Fomel, S., 2019. Saltseg: automatic 3D salt segmentation using a deep convolutional neural network, *Interpretation,* **7**(3), SE113–SE122.

Song, C. & Alkhalifah, T.A., 2021. Wavefield reconstruction inversion via physics-informed neural networks, *IEEE Trans. Geosci. Remote Sens.,* **60,** 1–12.

Sroubek, F. & Milanfar, P., 2011. Robust multichannel blind deconvolution via fast alternating minimization, *IEEE Trans. Image Process.,* **21**(4), 1687–1700.

Stockham, T.G., Cannon, T.M. & Ingebretsen, R.B., 1975. Blind deconvolution through digital signal processing, *Proc. IEEE,* **63**(4), 678–692.

Sun, W., Sacchi, M. & Gu, Y., 2023. Multichannel sparse deconvolution of teleseismic receiver functions with f-x preconditioning, *J. Geophys. Res.: Solid Earth,* **128**(4), e2022JB025625.

Tzeng, E., Hoffman, J., Saenko, K. & Darrell, T., 2017. Adversarial discriminative domain adaptation, in 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), IEEE Computer Society, 2962–2971. doi:10.1109/CVPR.2017.316.

Ulrych, T.J., 1971. Application of homomorphic deconvolution to seismology, *Geophysics,* **36**(4), 650–660.

Ulrych, T.J. & Bishop, T.N., 1975. Maximum entropy spectral analysis and autoregressive decomposition, *Rev. Geophys.,* **13**(1), 183–200.

Ulyanov, D., Vedaldi, A. & Lempitsky, V., 2018. Deep image prior, in D., Forsyth, I., Laptev, A., Oliva & D., Ramanan (Eds.), Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, IEEE Computer Society, pp. 9446–9454.

Velis, D., 2008. Stochastic sparse-spike deconvolution, *Geophysics,* **73,** R1–R9..

Wang, Y., 2002. A stable and efficient approach of inverse q filtering, *Geophysics,* **67**(2), 657–663.

Wang, Y., 2006. Inverse q-filter for seismic resolution enhancement, *Geophysics,* **71**(3), V51–V60.

Wang, Y., Xu, J., Zhao, Z., Gao, Y. & Zhang, H., 2024. Structurally-constrained unsupervised deep learning for seismic high-resolution reconstruction, *IEEE Trans. Geosci. Remote Sens,* **62,** 1–15.

Wang, Z., Simoncelli, E.P. & Bovik, A.C., 2003. Multiscale structural similarity for image quality assessment, *The Thrity-Seventh Asilomar Conference on Signals, Systems and Computers, 2003,* Vol. **2,** pp. 1398–1402, IEEE.

Wu, H., Zhang, B., Lin, T., Cao, D. & Lou, Y., 2019a. Semiautomated seismic horizon interpretation using the encoder-decoder convolutional neural network, *Geophysics,* **84**(6), B403–B417.

Wu, X., 2017. Structure-, stratigraphy- and fault-guided regularization in geophysical inversion, *Geophys. J. Int.,* **210**(1), 184–195.

Wu, X., Liang, L., Shi, Y. & Fomel, S., 2019b. Faultseg3d: using synthetic data sets to train an end-to-end convolutional neural network for 3D seismic fault segmentation, *Geophysics,* **84**(3), IM35–IM45.

Wu, X., Geng, Z., Shi, Y., Pham, N., Fomel, S. & Caumon, G., 2020. Building realistic structure models to train convolutional neural networks for seismic structural interpretation, *Geophysics,* **85**(4), WA27–WA39.

Xiao, Y., Li, K., Dou, Y., Li, W., Yang, Z. & Zhu, X., 2024. Diffusion models for multidimensional seismic noise attenuation and super-resolution, *Geophysics,* **89**(5), 1–66.

Yang, F. & Ma, J., 2019. Deep-learning inversion: a next-generation seismic velocity model building method, *Geophysics,* **84**(4), R583–R599.

Yilmaz, Ö., 2001. *Seismic Data Analysis: Processing, Inversion, and Interpretation of Seismic Data,* Society of Exploration Geophysicists.

Yu, S., Ma, J. & Wang, W., 2019. Deep learning for denoising, *Geophysics,* **84**(6), V333–V350.

Yuan, S. & Wang, S., 2013. Spectral sparse bayesian learning reflectivity inversion, *Geophys. Prospect.,* **61**(4), 735–746.

Zhang, H., Chen, T., Liu, Y., Zhang, Y. & Liu, J., 2021a. Automatic seismic facies interpretation using supervised deep learning, *Geophysics,* **86**(1), IM15–IM33.

Zhang, J., Gao, Y., Meng, Q., Wang, D., Zhang, Z., Li, G. & Jiang, L., 2021b. Well-controlled seismic resolution improvement: a machine learning approach, *82nd EAGE Annual Conference and Exhibition,* Vol. **2021,** pp. 1–5, European Association of Geoscientists and Engineers.

Zhang, R. & Castagna, J., 2011. Seismic sparse-layer reflectivity inversion using basis pursuit decomposition, *Geophysics,* **76**(6), R147–R158.

Zhao, L., Zou, C., Chen, Y., Shen, W., Wang, Y., Chen, H. & Geng, J., 2021. Fluid and lithofacies prediction based on integration of well-log data and seismic inversion: a machine-learning approach, *Geophysics,* **86**(4), M151–M165.