Preprint Coversheet

Title

ADMM-Guided Physics-Informed Deep Learning for Two-Dimensional Acoustic Impedance Inversion with Reweighted ℓ_1 Sparse Regularization

Authors and affiliations

Deepak Kumar[1,*], Jayant Nath Tripathi[2]

[1] Institute of Geophysics, Polish Academy of Sciences, Warsaw, Poland

[2] University of Allahabad, Prayagraj, India

[*] Corresponding author. E-mail: deepak.kumar@igf.edu.pl

ORCID (Deepak Kumar): https://orcid.org/0000-0001-7677-860X

Preprint statement

This manuscript is a non-peer-reviewed preprint submitted to *EarthArXiv*. It has been submitted to *IEEE Transactions on Geoscience and Remote Sensing* for peer review. Subsequent versions of this manuscript may have different content. If a peer-reviewed version is published, the final version will be accessible through the "Peer-reviewed Publication DOI" link on the EarthArXiv page for this preprint.

ADMM-Guided Physics-Informed Deep Learning for Two-Dimensional Acoustic Impedance Inversion with Reweighted ℓ_1 Sparse Regularization

Deepak Kumar[1,*] and Jayant Nath Tripathi[2]

[1]*Institute of Geophysics, Polish Academy of Sciences, Warsaw, Poland*

[2]*University of Allahabad, Prayagraj, India*

[*]Corresponding author. E-mail: deepak.kumar@igf.edu.pl

Abstract

Acoustic impedance inversion is ill-posed because post-stack seismic data are band-limited and noisy. Reweighted ℓ_1 sparse inversion sharpens impedance boundaries but, applied trace by trace, ignores lateral geological continuity. We present a two-dimensional ADMM-guided physics-informed neural framework in which a reweighted ℓ_1 ADMM estimate (after He et al., 2022) serves as a physics prior for a convolutional refinement network trained with a differentiable wavelet-convolution data term, reweighted ℓ_1 sparsity, model proximity, and lateral smoothness. Three architectures (a 2D U-Net, a residual CNN, and an Attention ResUNet) are compared against two classical baselines: trace-wise reweighted ℓ_1 and a stronger spatially-coupled 2D inversion with lateral total variation. On a controlled synthetic section and two geologically distinct Marmousi-2 crops, the networks markedly improve noisy-data reconstruction: the U-Net reduces RMSE by about 45% over trace-wise ADMM and, after transfer, by 44% on Marmousi-2 (a 5 dB SNR gain). The advantage is stable across five training seeds, persists against the spatially-coupled baseline, and holds at every input-noise level under matched training. A wavelet-mismatch study identifies a deployment caveat: without retraining, the networks are only as reliable as the estimated wavelet, whereas the spatially-coupled classical inversion degrades most gracefully.

Keywords: acoustic impedance inversion; ADMM; reweighted ℓ_1 regularization; physics-informed learning; 2D U-Net; ResCNN; Marmousi-2; transfer learning

Highlights:

- ADMM-guided physics-informed neural framework for 2D impedance inversion
- U-Net reduces noisy-data RMSE by 44.89% over trace-wise ADMM baseline
- Transfer learning validated on Marmousi-2 fault-zone benchmark crop
- Three architectures compared: 2D U-Net, ResCNN, Attention ResUNet
- Full source code and numerical results publicly available on GitHub

1. Introduction

Post-stack acoustic impedance inversion estimates the subsurface impedance model from seismic reflection data. In the convolutional model, the seismic trace is represented as the wavelet-convolved reflectivity series, and reflectivity is related to impedance contrasts. This inverse problem is non-unique because seismic data lack low frequencies, contain finite-bandwidth wavelets, and are contaminated by noise. Classical seismic impedance inversion therefore depends on regularization and prior information, including low-frequency initial models, sparse reflectivity assumptions, and structural constraints (Lindseth, 1979; Cooke and Schneider, 1983; Wang, 2010).

Sparse-spike methods use the fact that major impedance changes occur at a relatively small number of interfaces. The ℓ_1 norm became a standard sparsity-promoting penalty for deconvolution and impedance inversion because it favors compact reflection-coefficient sequences more strongly than quadratic smoothing penalties (Taylor et al., 1979; Wang, 2010). Reweighted ℓ_1 minimization further strengthens sparsity by assigning larger weights to small-amplitude coefficients and smaller weights to large-amplitude coefficients (Candès et al., 2008). He et al. (2022) adapted this idea to seismic acoustic impedance inversion and showed that a reweighted ℓ_1 sparse constraint can improve impedance-boundary resolution and reduce pseudolayer artifacts relative to a traditional ℓ_1 constraint. Their objective combines seismic data fidelity, a reweighted sparse reflectivity term, and an initial-model constraint, and is solved by ADMM (Boyd et al., 2011).

The limitation addressed here is not the one-dimensional sparse formulation itself; it is the loss of lateral context when the same trace-wise inversion is repeated independently across

a 2D section. Faults, thin beds, and channel-like or lens-like bodies are laterally organized structures. Structural noise can also leak across adjacent traces through migration artifacts, scattering, residual multiples, or acquisition footprints, so treating each trace independently can turn laterally coherent artifacts into impedance contrasts. A trace-by-trace inverse solver can recover sharp vertical impedance contrasts, but it cannot learn or enforce coherent lateral patterns across neighboring traces. Deep learning methods, including fully convolutional residual networks with transfer learning, have demonstrated effectiveness for learning complex nonlinear seismic-to-impedance mappings (Wu et al., 2020). Convolutional neural networks provide this spatial coupling. U-Net-style encoder-decoder networks preserve multi-scale context through skip connections (Ronneberger et al., 2015); residual networks improve trainability by learning corrections relative to an input representation (He et al., 2016); attention gates can emphasize informative skip features while suppressing irrelevant background responses (Oktay et al., 2018). Physics-informed neural networks introduce physical residuals directly into the loss function so that neural optimization is constrained by the governing forward model rather than acting as a purely data-driven interpolator (Raissi et al., 2019).

This work combines these two ideas. The reweighted ℓ_1 ADMM inversion supplies a physically meaningful sparse baseline, and the neural model acts as a two-dimensional refiner that uses seismic data and the ADMM-guided impedance prior as input channels. The goal is not to replace the sparse inverse problem with an unconstrained black-box network. Instead, the network is trained with the same physical ingredients that make the classical inversion interpretable: seismic wavelet convolution, impedance-derived reflectivity, sparse reflectivity, and proximity to a stable initial estimate.

2. Forward Model and Classical Reweighted ℓ_1 Inversion

Let $L = \log Z$ denote logarithmic acoustic impedance and let S denote a post-stack seismic trace. Following the standard linearized impedance inversion model, the reflectivity is approximated by a half-difference operator applied to L,

$$r \approx D_{1/2} L, \qquad D_{1/2} = 0.5 D, \tag{1}$$

where D is the first-difference matrix. If W is the Toeplitz matrix corresponding to the source wavelet, the forward model is

$$S = WD_{1/2}L + n, \tag{2}$$

where n represents noise.

The classical baseline used in this paper follows the reweighted ℓ_1 sparse formulation of He et al. (2022). For one trace, the optimization problem is

$$\min_L \ \frac{1}{2}\|S - WD_{1/2}L\|_2^2 + \mu\|MD_{1/2}L\|_1 + \frac{\alpha}{2}\|L - L_0\|_2^2, \tag{3}$$

where L_0 is the log-impedance initial model, μ is the sparse penalty weight, α controls proximity to the initial model, and M is a diagonal reweighting matrix. The weights are updated from the current reflectivity estimate as

$$M_i = \frac{1}{|r_i| + \epsilon}, \tag{4}$$

with a small ϵ for numerical stability. This weighting makes small reflectivity coefficients more expensive and preserves strong impedance boundaries.

ADMM introduces an auxiliary variable R and scaled dual variable C so that the sparse term can be updated through soft thresholding. In this implementation, the L update is a linear solve, the R update is soft thresholding of $MD_{1/2}L + C$, and the dual variable is updated by residual accumulation. The operator and parameter values used are: $\mu = 5 \times 10^{-6}$, $\alpha = 5 \times 10^{-5}$, $\lambda = 10^{-4}$, maximum ADMM iterations of 200, and convergence tolerance 10^{-6}.

2.1. 2D Spatially-Coupled Classical Baseline

The reweighted ℓ_1 formulation in Equation 3 is applied independently to each trace and therefore imposes no lateral coupling between neighboring impedance profiles. To ensure that any advantage of the neural refinement is attributable to learned geometric priors rather than to spatial coupling alone, we additionally construct a stronger, spatially-coupled classical baseline. Let $L \in \mathbb{R}^{N_z \times N_x}$ denote the full log-impedance section, with $L_{:,j}$ the j-th trace. The 2D objective augments the per-trace term with an anisotropic lateral total-variation (TV) penalty (Rudin et al.,

1992),

$$\min_{L} \sum_{j} \left[\tfrac{1}{2} \| S_{:,j} - W D_{1/2} L_{:,j} \|_2^2 + \mu_t \| M_j D_{1/2} L_{:,j} \|_1 + \tfrac{\alpha}{2} \| L_{:,j} - L_{0,:,j} \|_2^2 \right] + \mu_x \sum_{i} \| D_x L_{i,:} \|_1, \tag{5}$$

where D_x is the lateral first-difference operator and μ_x weights the lateral TV term. We solve Equation 5 by consensus (variable-splitting) ADMM (Boyd et al., 2011), which decouples the objective into two subproblems linked by a consensus variable V and scaled dual U. The vertical subproblem reduces to the per-trace reweighted ℓ_1 inversion of Equation 3, augmented by a proximal term $\frac{\rho}{2}\|L - (V - U)\|_2^2$ and solved exactly and simultaneously across all traces through a batched linear solve. The lateral subproblem is a 1-D TV denoising of each impedance row, $V = \arg\min_V \frac{\rho}{2}\|V - (L + U)\|_2^2 + \mu_x \|D_x V\|_1$, solved by an inner ADMM that shares a single Cholesky factorization across rows. The dual update is $U \leftarrow U + L - V$. We use μ_t, α, λ identical to the trace-wise baseline, $\mu_x = 100\,\mu_t$, $\rho = 40\,\alpha$, 40 outer iterations, and 10 inner vertical sweeps per outer iteration; setting $\mu_x = 0$ reproduces the trace-wise result to within 0.01 dB, confirming the implementation.

3. ADMM-Guided Physics-Informed Neural Refinement

3.1. Input Representation

The neural models operate on 2D patches. Each patch has two input channels: the seismic section patch and an impedance-prior patch. The prior is the reweighted ℓ_1 ADMM result when available; otherwise it can be the smoothed initial impedance model. In the controlled numerical benchmark reported here, the neural input prior is the ADMM result, so the network learns a correction to a sparse physics-based inversion rather than learning impedance directly from seismic amplitudes alone.

3.2. Architectures

Three neural architectures are compared in the same training framework. The first is a 2D U-Net with encoder-decoder paths and skip connections. The second is a residual CNN that learns a spatial correction through stacked residual blocks. The third is an Attention ResUNet, which combines residual encoding/decoding with attention-gated skip connections. Figure 1

summarizes the full ADMM-guided workflow, and Figure 2 summarizes the three neural architectures used in the compact 2D benchmark code.

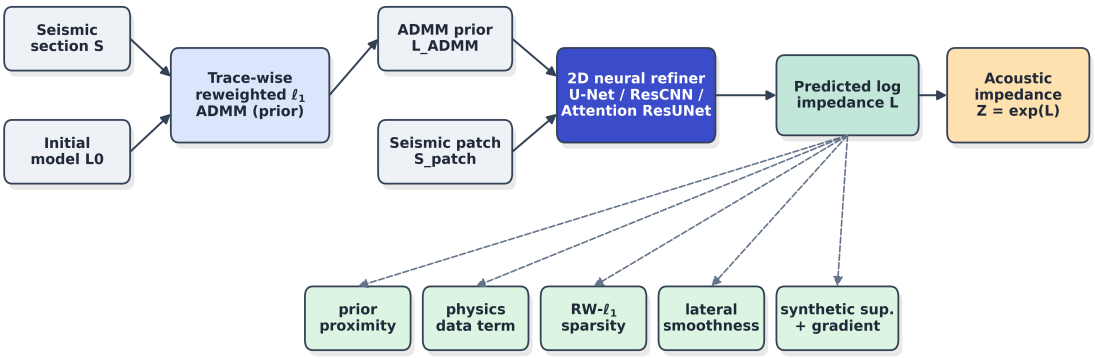

Figure 1: ADMM-guided physics-informed neural inversion workflow. The neural models receive seismic data and an ADMM-guided impedance prior, then optimize a loss composed of seismic physics, reweighted sparse reflectivity, prior proximity, lateral regularity, and synthetic benchmark supervision with gradient sharpening.

3.3. Loss Function

For a predicted log-impedance patch $\hat{L}$, the neural model computes the exact nonlinear reflection coefficient

$$\hat{r}_j = \frac{\hat{Z}_{j+1} - \hat{Z}_j}{\hat{Z}_{j+1} + \hat{Z}_j + \eta}, \qquad \hat{Z} = \exp(\hat{L}), \tag{6}$$

where η prevents division by zero. This exact nonlinear step is part of the PyTorch (Paszke et al., 2019) computation graph in the differentiable physics layer. Backpropagation therefore uses automatic differentiation through the exponential, division, and wavelet-convolution operations; no linearized reflection-coefficient approximation is substituted during the neural backward pass. The predicted seismic patch is obtained by wavelet convolution of $\hat{r}$. The total training loss is

$$\begin{aligned}\mathcal{L} = {} & w_p \|\hat{S} - S\|_2^2 + w_s \|M\hat{r}\|_1 + w_i \|\hat{L} - L_{\text{prior}}\|_2^2 \\ & + w_{lat} \|\nabla_x \hat{L}\|_1 + w_g \|\nabla \hat{L} - \nabla L_{\text{true}}\|_1 + w_{sup} \|\hat{L} - L_{\text{true}}\|_2^2.\end{aligned} \tag{7}$$

The reported synthetic benchmark uses $w_p = 3.0$, $w_s = 0.01$, $w_{lat} = 0.005$, $w_g = 0.6$, and $w_{sup} = 2.0$. The initial-model proximity weight is linearly annealed from $w_i^{(0)} = 0.5$ to

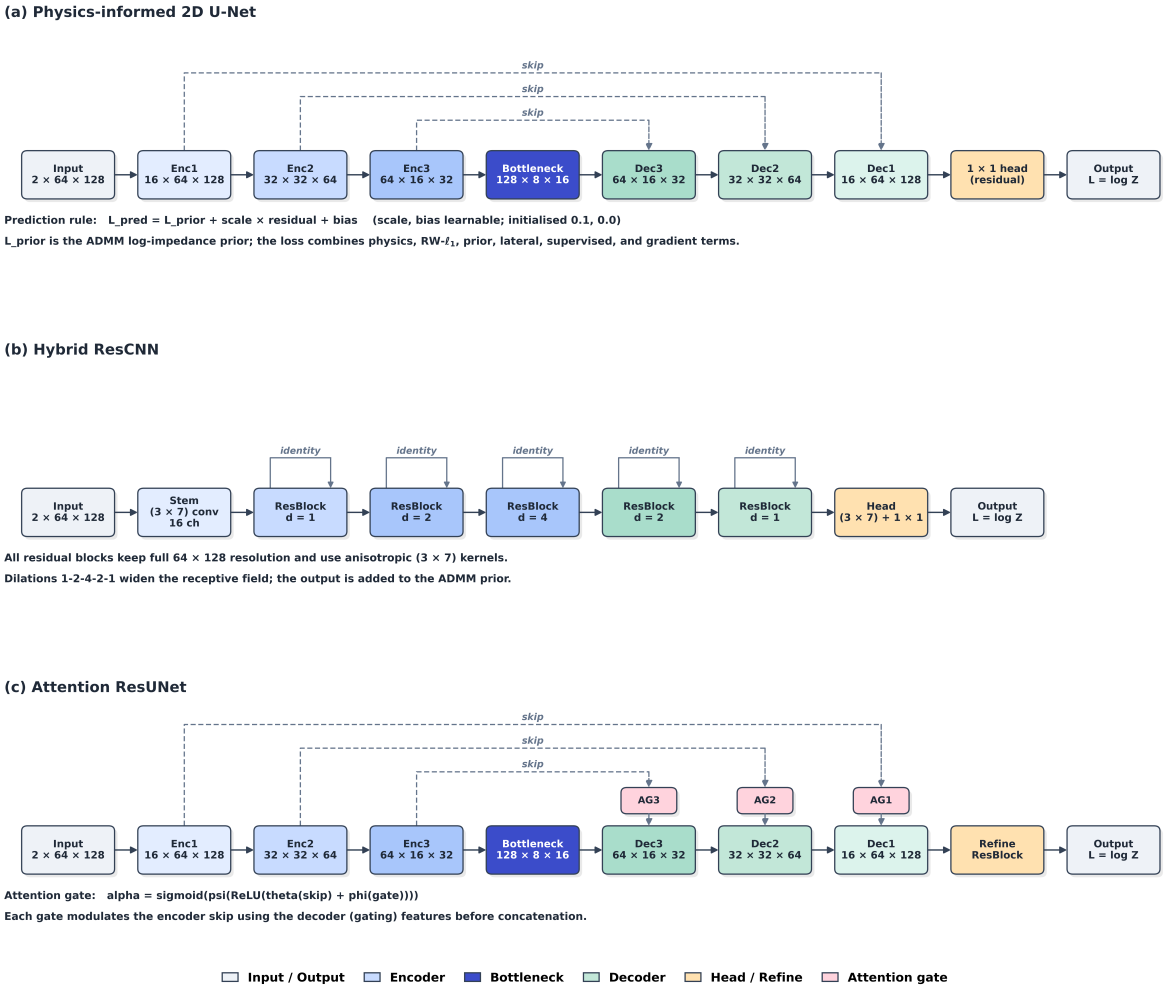

Figure 2: Deep learning architectures used in the controlled 2D benchmark: physics-informed 2D U-Net, hybrid ResCNN, and Attention ResUNet.

$w_i^{(\mathrm{end})} = 0.01$ during the first 25 epochs:

$$w_i(e) = \begin{cases} 0.5 - (0.5 - 0.01)\, e/25, & e < 25, \\ 0.01, & e \geq 25, \end{cases} \tag{8}$$

where e is the zero-based epoch index. The terms involving L_{true} are included because the synthetic model is known and this experiment is a controlled benchmark. They should be omitted for blind field-data inversion, leaving the field-compatible physics, sparsity, prior-proximity, and lateral-regularity terms.

4. Synthetic 2D Controlled Benchmark

4.1. Model and Data

The experiment uses a 2D synthetic acoustic impedance section with 400 lateral traces and 300 time samples at a sampling interval of 1 ms. The model is intentionally more complex than a simple horizontally layered section: it includes dipping layers, a normal fault, thin beds, and an oval impedance body. The Ricker wavelet has a dominant frequency of 30 Hz. A clean seismic section and a noisy section are generated from the same impedance model. The noisy case uses the code setting `NOISE_SNR=8.0`, giving a deliberately difficult test of noise robustness.

The low-frequency starting model is produced by two-dimensional Gaussian smoothing of the true impedance. This initial model is not a solution by itself; it contains the large-scale impedance trend but lacks sharp interfaces, thin beds, and detailed fault structure. Figure 3 shows the true impedance, the smoothed initial model, the clean and noisy seismic sections, the true reflectivity, and the Ricker wavelet used in the forward model.

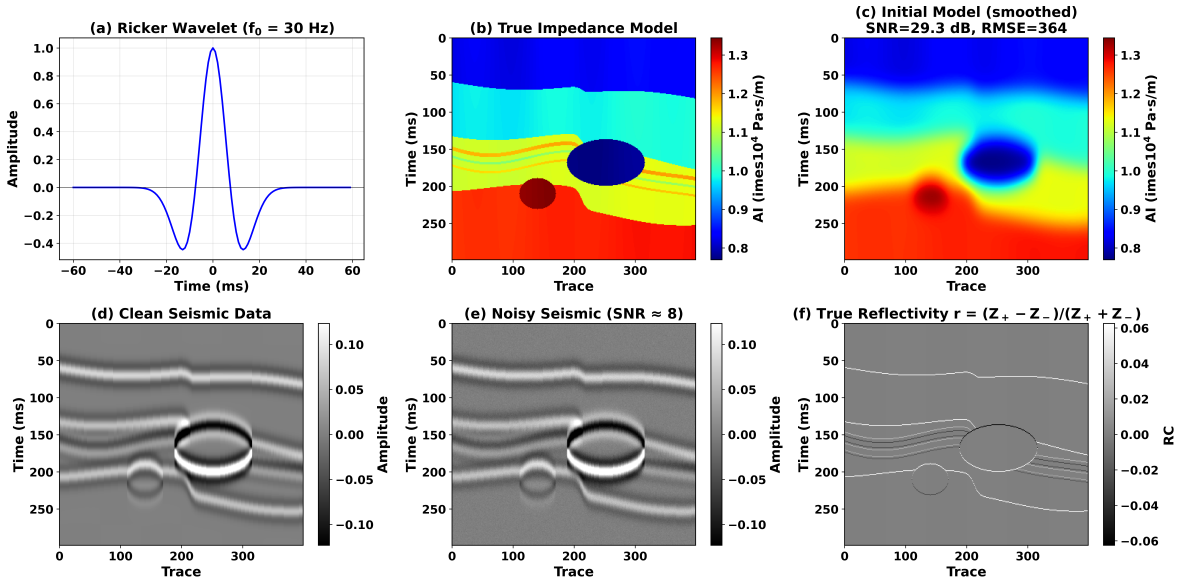

Figure 3: Synthetic 2D controlled benchmark: true impedance, smoothed initial model, clean seismic data, noisy seismic data, true reflectivity, and Ricker wavelet.

4.2. Training and Numerical Settings

The compact benchmark used the following settings: patch size 64×128, stride 32×64, batch size 8, base channel count 16, learning rate 5×10^{-4}, and 80 epochs for each neural model/data condition. The training patches cover the full 400×300 section with anchored corner patches to

avoid uncovered boundary regions. The optimizer is Adam with cosine annealing to a minimum learning rate of 10^{-6}.

5. Results

5.1. Clean Data

On clean data, reweighted ℓ_1 ADMM already performs strongly because the sparse reflectivity assumption matches the synthetic layered model. The neural refinements therefore produce modest but measurable improvements rather than dramatic changes. The ResCNN achieves the best clean-data score with SNR 40.85 dB and RMSE 95.89, compared with ADMM SNR 40.51 dB and RMSE 99.76. The U-Net gives SNR 40.79 dB and RMSE 96.56. The Attention ResUNet remains comparable but does not outperform the simpler ResCNN or U-Net in this clean-data case. Figure 4 compares the initial model, the reweighted ℓ_1 ADMM result, and the U-Net refinement on clean data, and Figure 5 shows the corresponding ResCNN comparison against the ADMM baseline under both clean and noisy data.

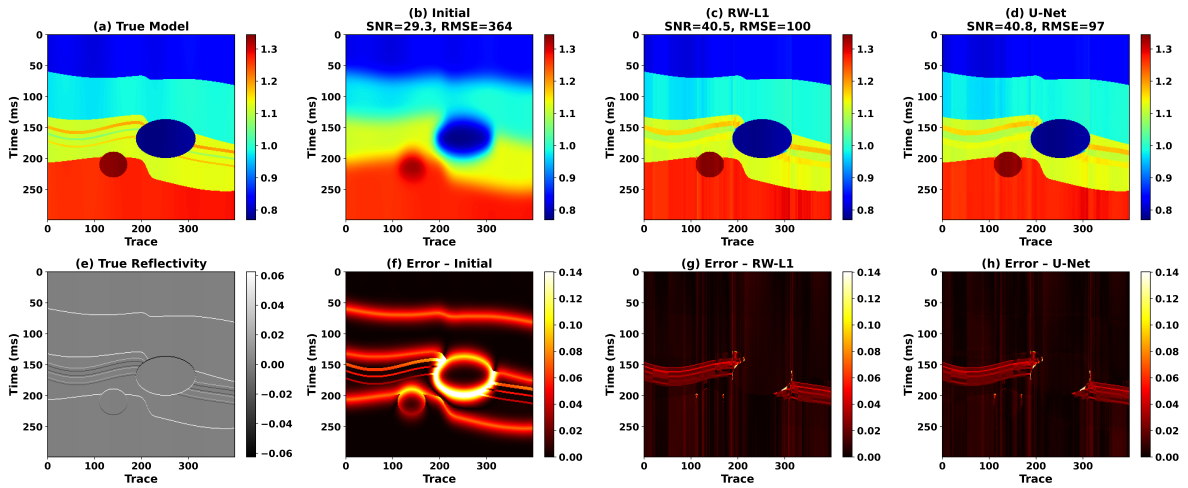

Figure 4: Clean-data inversion comparison for the initial model, reweighted ℓ_1 ADMM, and physics-informed 2D U-Net.

5.2. Noisy Data

The noisy case is where the two-dimensional neural refinement is most valuable. The trace-wise ADMM baseline remains sparse but is more sensitive to noise because each trace is inverted independently. The 2D neural models exploit neighboring traces and suppress incoherent artifacts while preserving major impedance boundaries. The U-Net gives the best noisy result, increasing

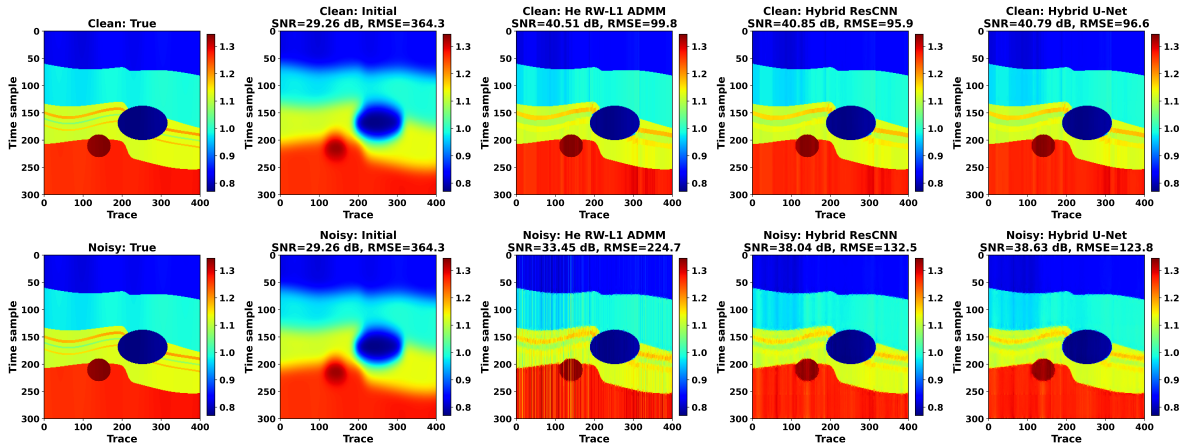

Figure 5: ResCNN comparison with the reweighted ℓ_1 ADMM baseline under clean and noisy seismic data.

SNR from 33.45 dB to 38.63 dB and reducing RMSE from 224.67 to 123.81. The ResCNN also improves substantially, reaching SNR 38.04 dB and RMSE 132.49. The Attention ResUNet improves over ADMM but is not the best architecture in this specific controlled benchmark. Figure 6 compares the initial model, trace-wise ADMM, and U-Net refinement on the noisy section; Figure 7 maps where the ResCNN improves or degrades relative to ADMM; and Figure 8 places all four methods, including the Attention ResUNet, side by side.

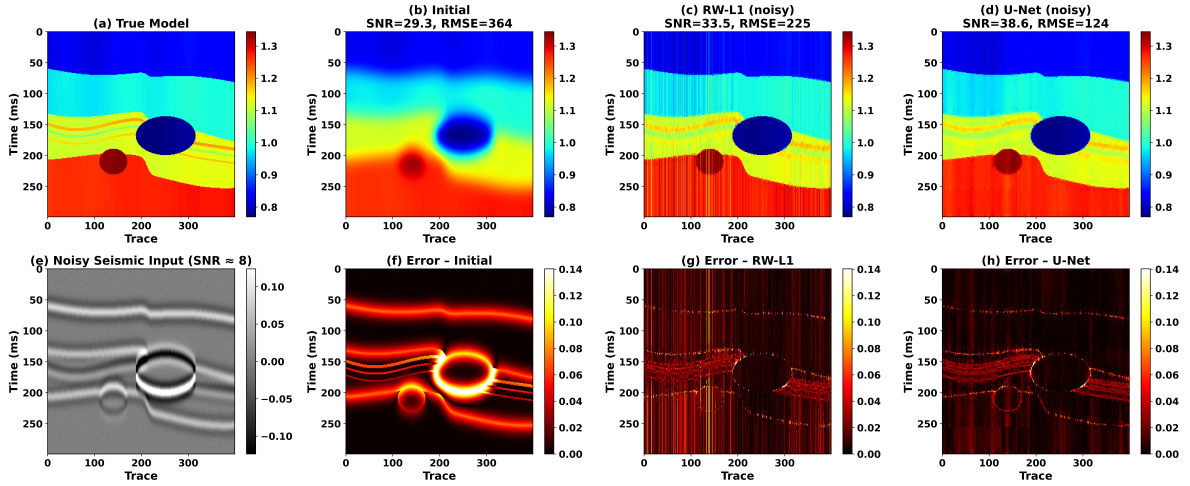

Figure 6: Noisy-data inversion comparison for the initial model, reweighted ℓ_1 ADMM, and physics-informed 2D U-Net.

5.3. Structural Diagnostics

Global metrics are useful but can hide localized structural failure. For this reason, Figure 9 isolates two geologically important windows: the faulted interval around traces 165–255 and samples 125–235, and a thin-bed interval evaluated by trace extraction. In the noisy fault window,

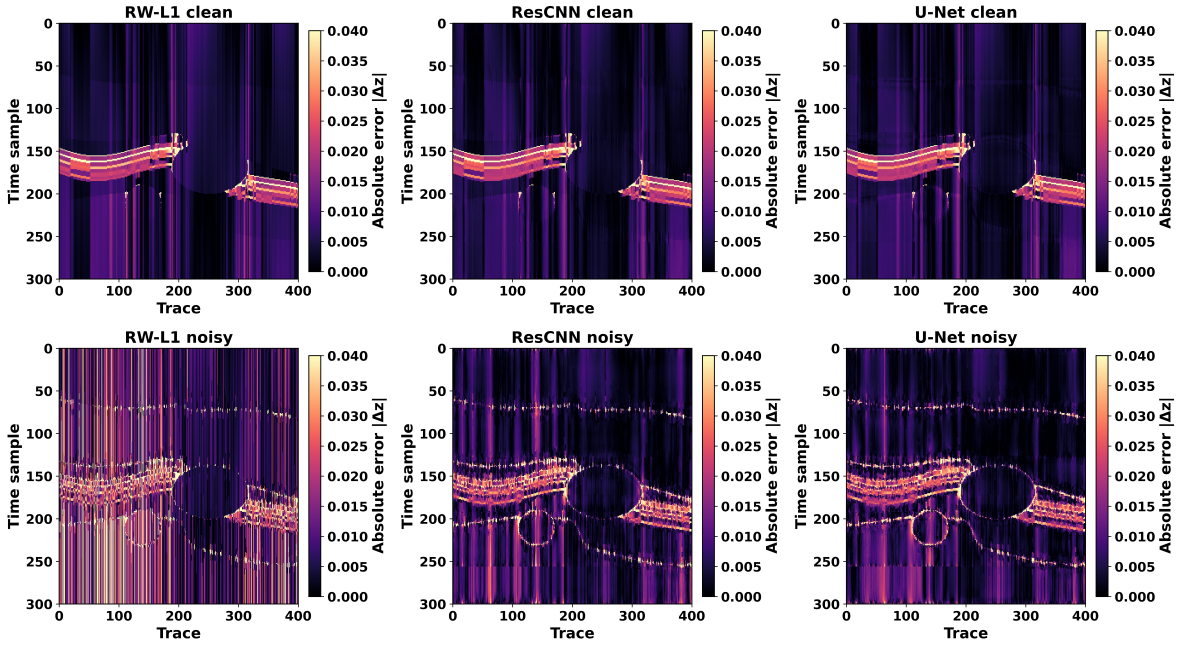

Figure 7: Error-map comparison showing where the ResCNN improves or degrades relative to the reweighted ℓ_1 ADMM baseline.

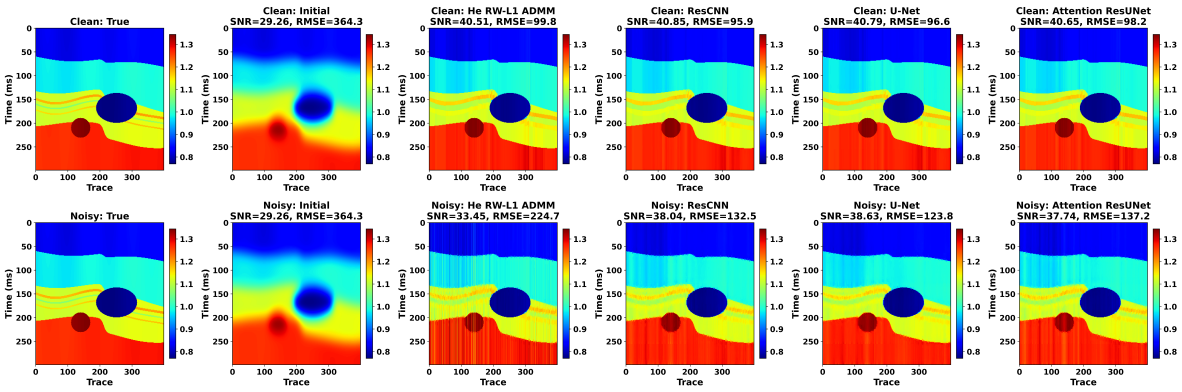

Figure 8: Advanced architecture comparison including reweighted ℓ_1 ADMM, ResCNN, 2D U-Net, and Attention ResUNet.

the U-Net reduces RMSE from 219.88 for trace-wise ADMM to 195.56. In the thin-bed window, the reduction is larger, from 305.54 to 208.99. These local diagnostics support the global metrics while also showing that the neural refinement does not simply improve a section-average score; it improves the structurally challenging regions where trace-wise noise and tuning effects are most visible.

5.4. Quantitative Metrics

Table 1 reports SNR, RMSE, and MAE computed directly from the numerical results. The key conclusion is conditional rather than overstated. On clean data, the neural models are close to the ADMM result, with ResCNN giving the best metric. On noisy data, all three neural models

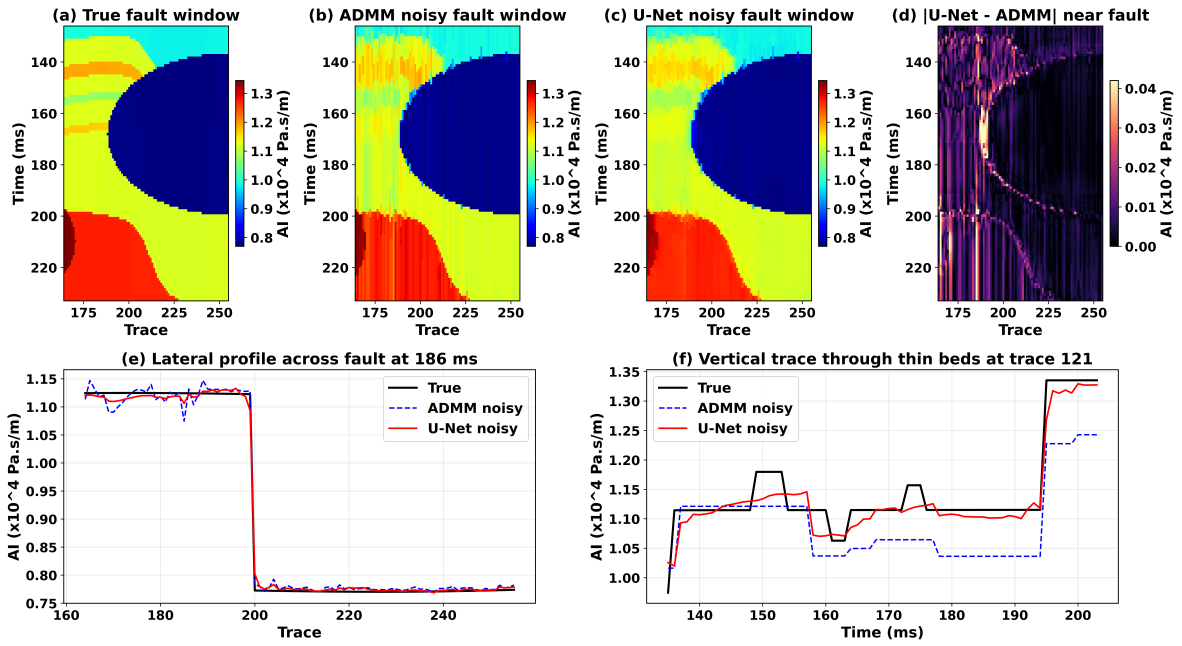

Figure 9: Structural diagnostic windows for the noisy experiment. The top row compares the true model, trace-wise ADMM, U-Net refinement, and absolute U-Net-ADMM difference around the fault. The bottom row extracts a lateral profile across the fault and a vertical trace through the thin beds.

outperform ADMM by a large margin, and the 2D U-Net gives the best noisy-data performance. Figure 10 summarizes these SNR, RMSE, and MAE values graphically for all four methods on clean and noisy data.

Table 1: Quantitative impedance-inversion metrics from the controlled 2D benchmark.

Method	Data	SNR (dB)	RMSE	MAE
Initial model	–	29.26	364.30	208.75
Reweighted ℓ_1 ADMM	clean	40.51	99.76	62.87
Hybrid ResCNN	clean	40.85	95.89	57.83
Physics-informed 2D U-Net	clean	40.79	96.56	58.72
Attention ResUNet	clean	40.65	98.15	60.21
Reweighted ℓ_1 ADMM	noisy	33.45	224.67	142.21
Hybrid ResCNN	noisy	38.04	132.49	75.90
Physics-informed 2D U-Net	noisy	38.63	123.81	70.29
Attention ResUNet	noisy	37.74	137.17	81.13

6. Transfer Learning Validation on the Marmousi-2 Model

The controlled synthetic benchmark in Sections 4–5 demonstrates the ADMM-guided neural framework on a purpose-built impedance model. To test whether the learned convolutional features generalize to a geologically distinct and more complex structure, the pretrained U-Net

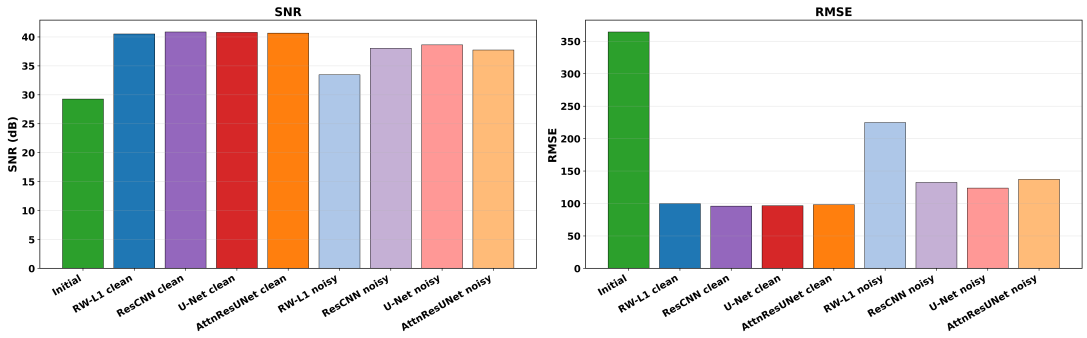

Figure 10: Metric summary for ADMM, ResCNN, U-Net, and Attention ResUNet on clean and noisy data.

and ResCNN weights are transferred to a cropped section of the Marmousi-2 acoustic impedance model (Martin et al., 2006). Transfer learning has been shown to be effective for seismic impedance inversion when the source and target domains share underlying physical relationships (Wu et al., 2020).

6.1. Marmousi-2 Crop and Data Generation

The Marmousi-2 model is a widely used elastic benchmark that spans 17 km laterally and 3.5 km in depth. Rather than inverting the full model, a 400-trace by 300-sample fault-zone crop is extracted from the region containing the most complex geology: intersecting dipping reflectors, a normal fault, and laterally varying impedance contrasts. Figure 11 shows the crop location within the full Marmousi-2 acoustic impedance section, and Figure 12 shows the cropped true impedance, smoothed initial model, clean seismic data, noisy seismic data (SNR = 8), and true reflectivity.

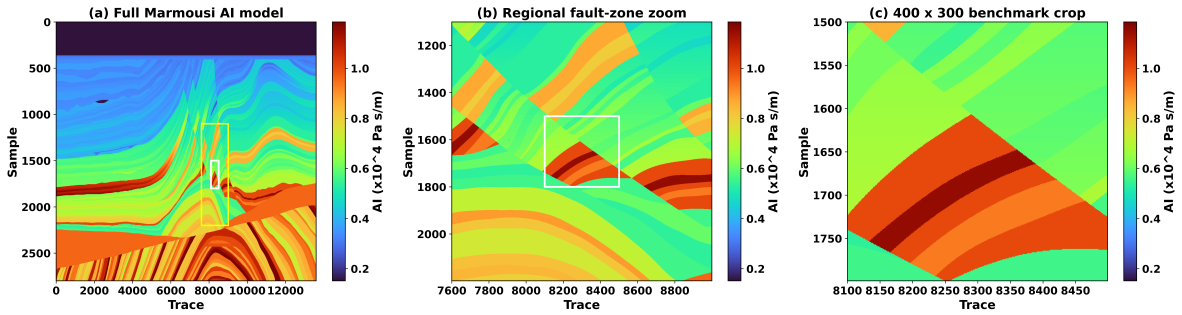

Figure 11: Marmousi-2 crop location. (a) Full Marmousi-2 acoustic impedance model. (b) Regional fault-zone zoom. (c) The 400 × 300 benchmark crop used for transfer learning validation.

The crop size matches the controlled synthetic benchmark (400 × 300), so the same patch geometry (64 × 128, stride 32 × 64) and training configuration apply. The ADMM baseline is

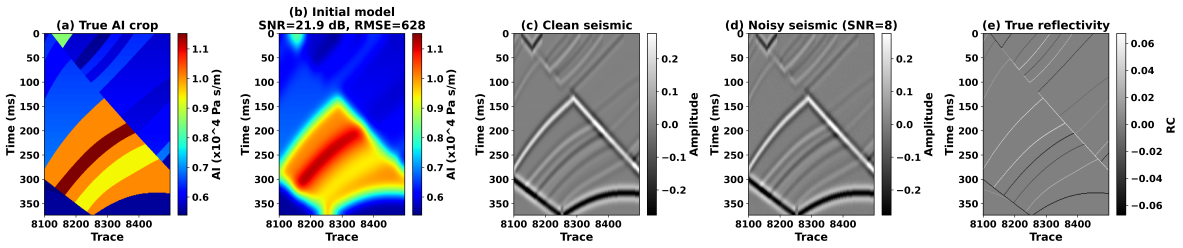

Figure 12: Marmousi-2 faulted crop: true impedance, smoothed initial model, clean seismic section, noisy seismic section (SNR = 8), and true reflectivity.

recomputed independently on the Marmousi-2 crop using the same reweighted ℓ_1 parameters. The U-Net and ResCNN are initialized from the pretrained weights obtained on the controlled synthetic model and fine-tuned for 80 epochs on the Marmousi-2 data.

6.2. Marmousi-2 Results

Figure 13 compares the true model, initial model, ADMM result, and U-Net refinement for the noisy Marmousi-2 case. The ADMM baseline recovers the major impedance boundaries but introduces noise-induced artifacts, particularly around the fault and in the laterally varying low-impedance zones. The U-Net refinement suppresses these artifacts while preserving the fault geometry and layer contrasts.

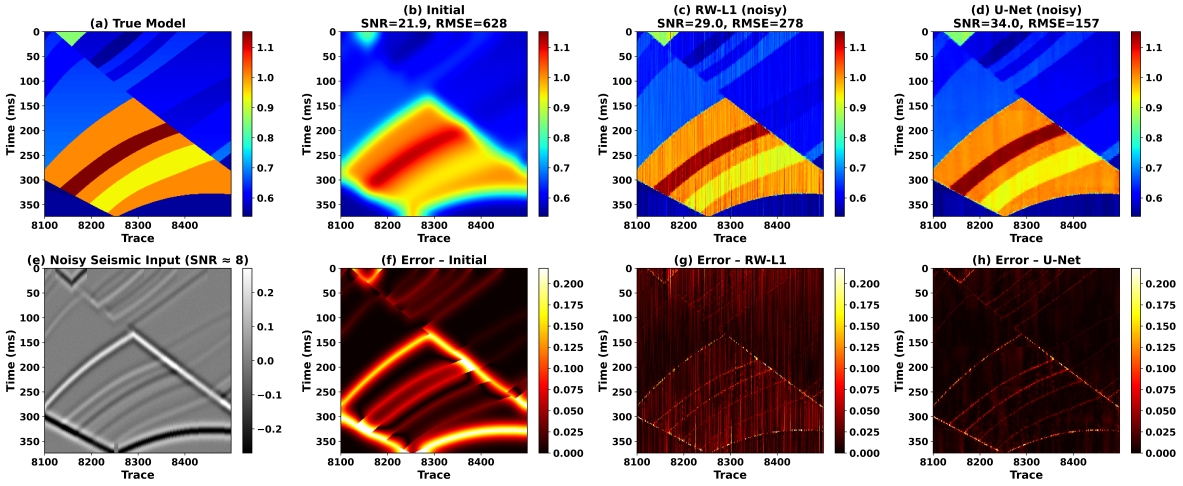

Figure 13: Noisy-data inversion on the Marmousi-2 faulted crop: true model, initial model, reweighted ℓ_1 ADMM, and physics-informed 2D U-Net, with corresponding error maps.

Figure 14 presents the architecture comparison across all tested models for both clean and noisy conditions. On the Marmousi-2 crop, the neural models achieve substantially larger improvements over ADMM than on the controlled synthetic model, reflecting the greater geological complexity and the increased difficulty of the trace-wise inverse problem.

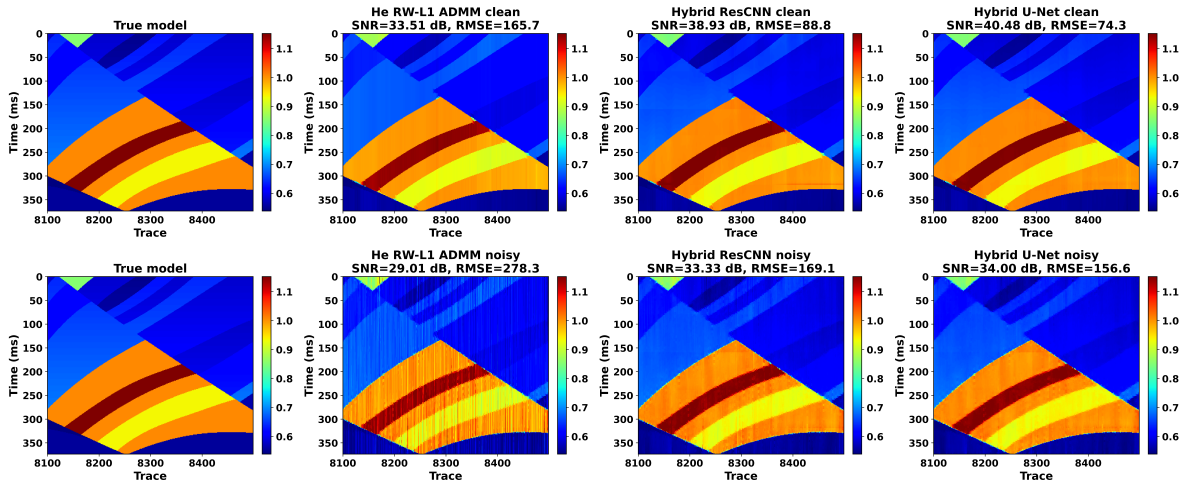

Figure 14: Architecture comparison on the Marmousi-2 faulted crop: reweighted ℓ_1 ADMM, Hybrid ResCNN, and 2D U-Net under clean and noisy conditions.

Single-trace impedance profiles at three representative locations (traces 101, 201, and 301) confirm that both the U-Net and ResCNN recover the sharp impedance contrasts more accurately than the ADMM baseline, particularly in the noisy case where trace-wise ADMM introduces spurious oscillations (Figure 15).

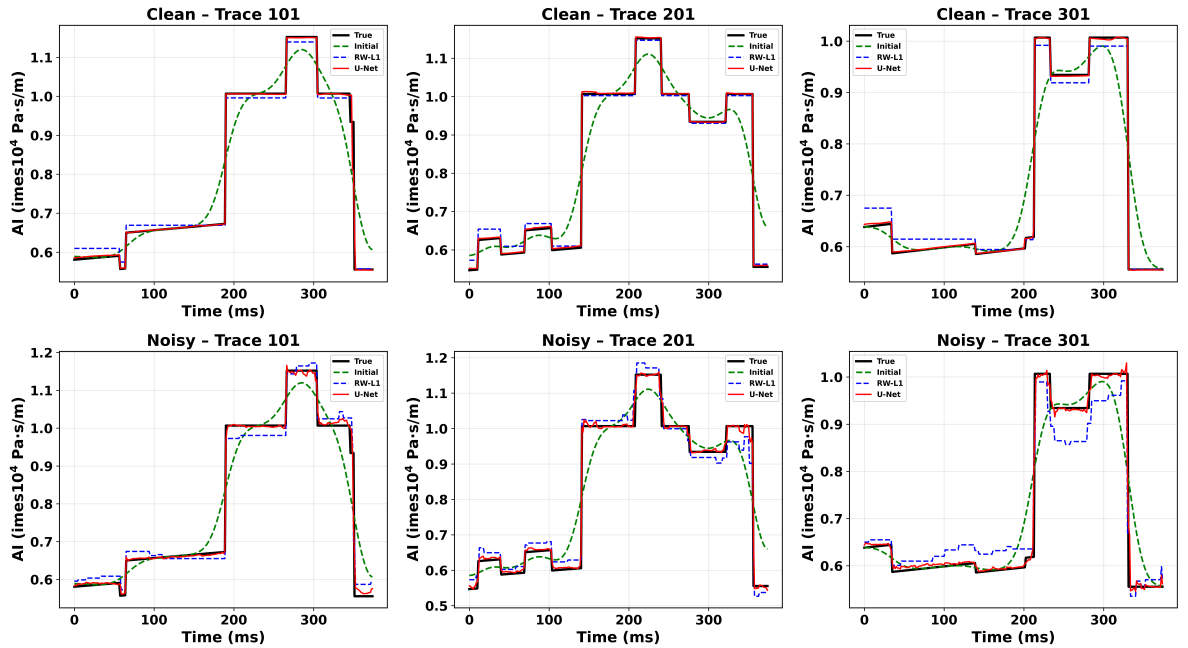

Figure 15: Single-trace impedance comparisons at three locations through the Marmousi-2 faulted crop under clean (top row) and noisy (bottom row) conditions.

6.3. Marmousi-2 Quantitative Metrics

Table 2 reports the quantitative metrics for the Marmousi-2 transfer learning experiment. The improvements are larger than those observed on the controlled synthetic model (Table 1), which

is expected because the Marmousi-2 fault-zone crop presents a harder inversion problem with more complex lateral impedance variations. On clean data, the U-Net reduces RMSE by 55.19% and the ResCNN by 46.43% relative to ADMM. On noisy data, the U-Net reduces RMSE by 43.74% with a 5.00 dB SNR gain, and the ResCNN reduces RMSE by 39.23%. The table also reports the 2D spatially-coupled classical baseline of Section 2.1, which is analyzed in Section 6.4. Figure 16 summarizes these results graphically.

Table 2: Quantitative impedance-inversion metrics from the Marmousi-2 faulted crop benchmark.

Method	Data	SNR (dB)	RMSE	MAE
Initial model	–	21.94	627.69	373.42
Reweighted ℓ_1 ADMM (trace-wise)	clean	33.51	165.74	129.35
2D-coupled ADMM (TV + RW-ℓ_1)	clean	31.53	208.22	136.78
Hybrid ResCNN	clean	38.93	88.79	50.22
Physics-informed 2D U-Net	clean	40.48	74.27	30.12
Reweighted ℓ_1 ADMM (trace-wise)	noisy	29.01	278.33	187.48
2D-coupled ADMM (TV + RW-ℓ_1)	noisy	31.48	209.46	107.54
Hybrid ResCNN	noisy	33.33	169.14	83.93
Physics-informed 2D U-Net	noisy	34.00	156.59	74.41

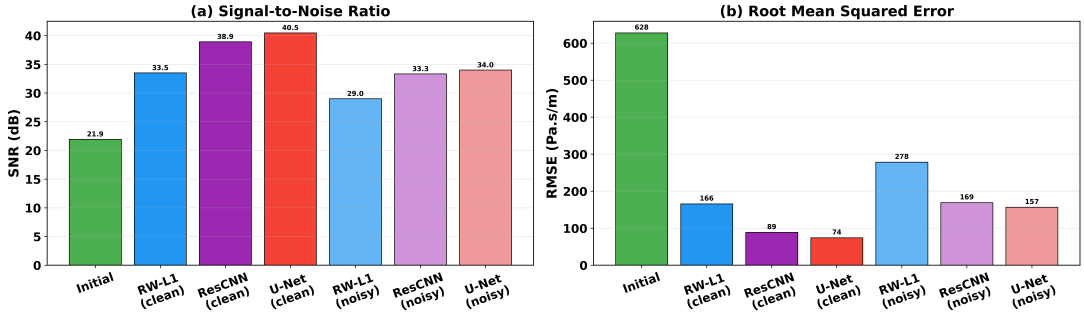

Figure 16: Metric summary for ADMM, ResCNN, and U-Net on the Marmousi-2 faulted crop under clean and noisy data.

6.4. Spatial Coupling Is Not the Source of the Neural Advantage

A natural concern is that the neural improvement over the trace-wise reweighted ℓ_1 baseline simply reflects the networks' ability to exploit lateral spatial context, which the per-trace classical inversion cannot. To isolate this factor we evaluate the 2D spatially-coupled classical baseline of Section 2.1, which adds an explicit lateral total-variation prior to the same reweighted ℓ_1 formulation. Figure 17 compares the true impedance, initial model, trace-wise reweighted ℓ_1, 2D-coupled ADMM, and the two physics-informed networks on the noisy Marmousi-2 crop, and Figure 18 summarizes the corresponding SNR and RMSE.

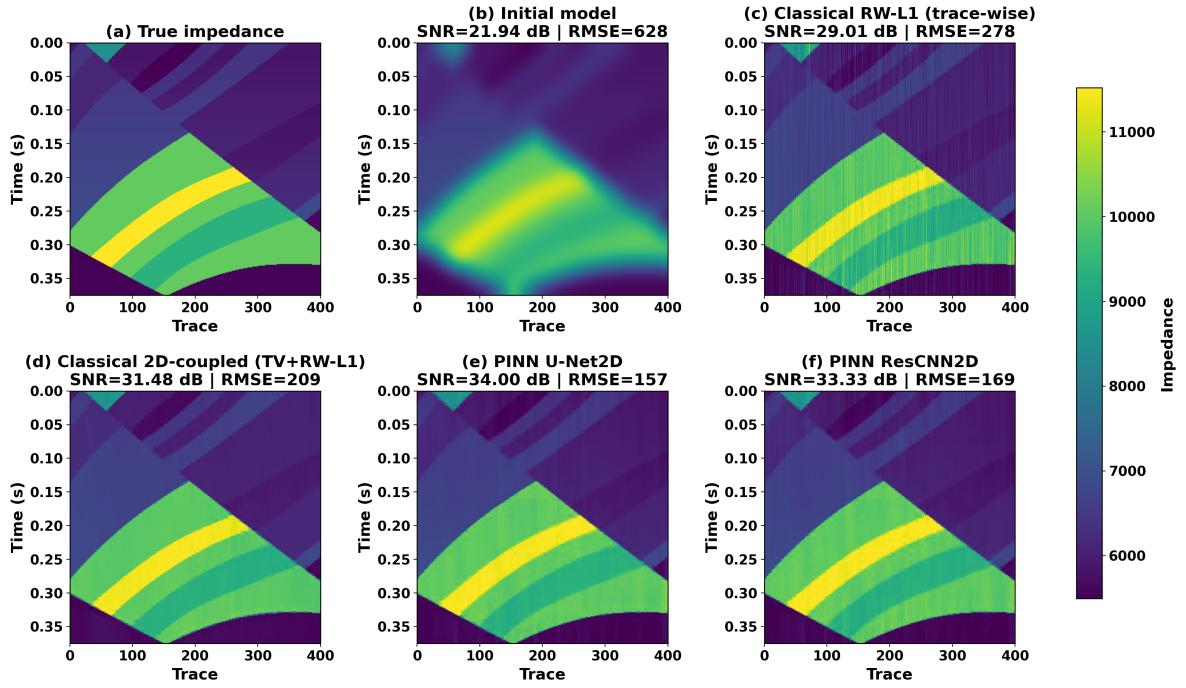

Figure 17: Noisy-data inversion on the Marmousi-2 faulted crop comparing both classical baselines against the physics-informed networks: true impedance, smoothed initial model, trace-wise reweighted ℓ_1 ADMM, 2D-coupled ADMM (lateral TV + reweighted ℓ_1, Section 2.1), physics-informed 2D U-Net, and Hybrid ResCNN. Adding lateral coupling to the classical inversion does not close the gap to the learned refinement.

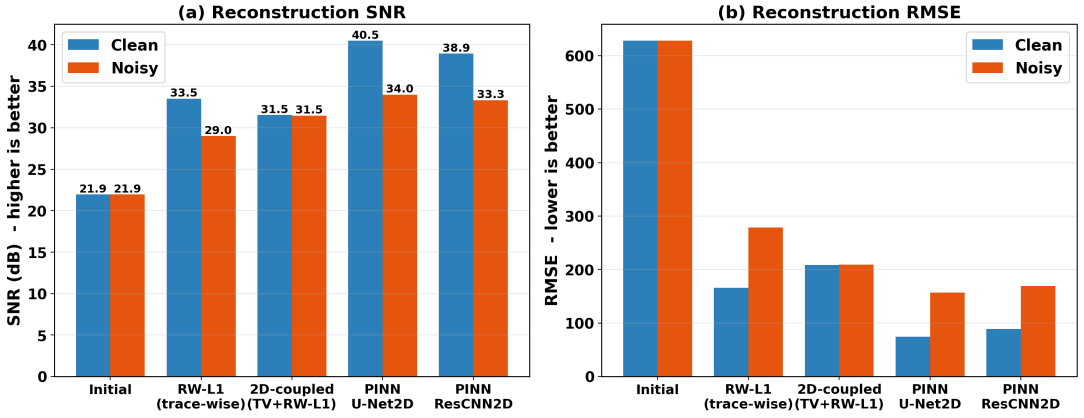

Figure 18: SNR and RMSE on the Marmousi-2 faulted crop for the trace-wise reweighted ℓ_1 ADMM, the 2D-coupled ADMM (lateral TV + reweighted ℓ_1), and the two physics-informed networks, under clean and noisy data.

The lateral TV prior is neutral-to-slightly-detrimental on clean data, where it lowers SNR from 33.51 dB to 31.53 dB (Table 2): an isotropic lateral smoothness assumption conflicts with the steeply dipping reflectors and the normal fault in this crop. Under noise the coupling is beneficial, raising SNR from 29.01 dB to 31.48 dB by suppressing trace-wise noise-induced oscillations. Crucially, in both regimes the physics-informed networks remain 6–9 dB above

the best classical result, and the 2D-coupled baseline never approaches the neural performance. Because the explicitly spatially-coupled classical inversion still falls well short of the networks, the neural advantage cannot be attributed to spatial coupling alone; it reflects the geometry-aware prior learned by the convolutional refinement.

6.5. Generalization to a Second Marmousi-2 Crop

To confirm that the results are not specific to the fault-zone crop, the same workflow is applied to a second, geologically distinct 400×300 Marmousi-2 crop extracted from a different region of the model (lateral samples 4500–4900, depth samples 1000–1300). This region is structurally simpler than the fault zone, reflected in the higher initial-model SNR (31.91 dB versus 21.94 dB). Figure 19 compares the true impedance, initial model, trace-wise reweighted ℓ_1 ADMM, 2D-coupled ADMM, and the two physics-informed networks under noise, and Table 3 reports the quantitative metrics.

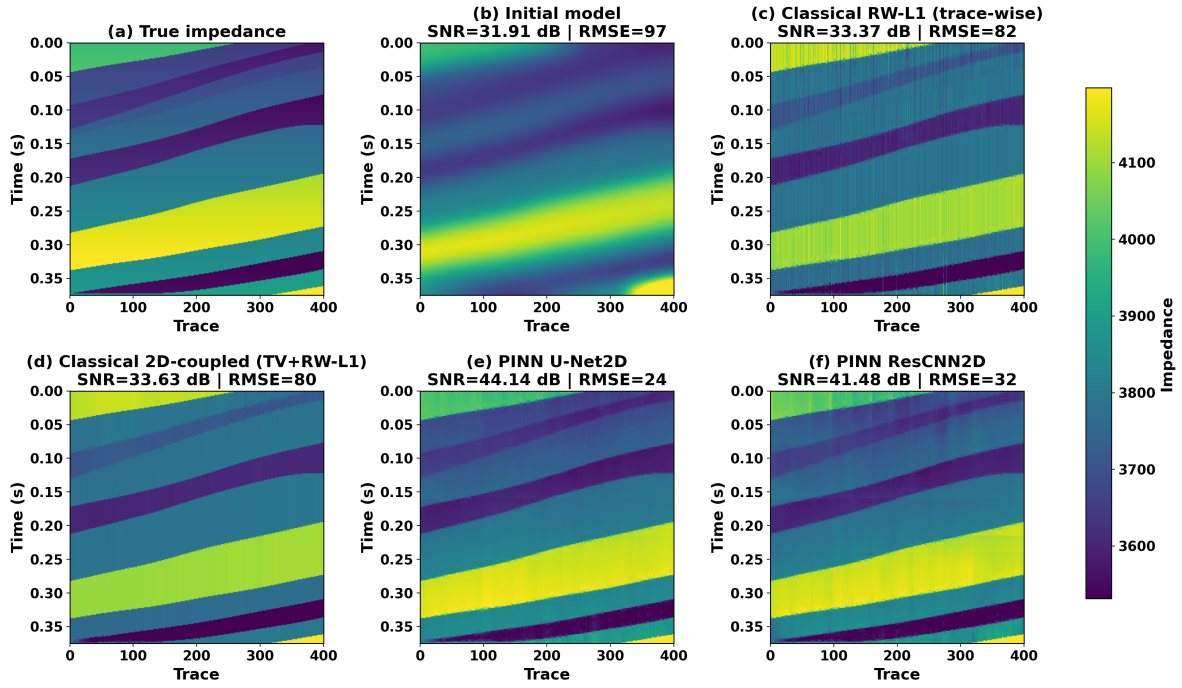

Figure 19: Noisy-data inversion on the second Marmousi-2 crop: true impedance, smoothed initial model, trace-wise reweighted ℓ_1 ADMM, 2D-coupled ADMM (lateral TV + reweighted ℓ_1), physics-informed 2D U-Net, and Hybrid ResCNN.

The pattern observed on the fault-zone crop is reproduced on this independent crop: the lateral TV coupling is essentially neutral relative to the trace-wise baseline (within 0.05 dB on clean data, +0.26 dB on noisy data), whereas both physics-informed networks improve SNR

Table 3: Quantitative impedance-inversion metrics on the second Marmousi-2 crop.

Method	Data	SNR (dB)	RMSE	MAE
Initial model	–	31.91	97.21	60.88
Reweighted ℓ_1 ADMM (trace-wise)	clean	33.46	81.27	69.05
2D-coupled ADMM (TV + RW-ℓ_1)	clean	33.42	81.62	69.17
Hybrid ResCNN	clean	45.73	19.79	15.48
Physics-informed 2D U-Net	clean	46.83	17.44	13.44
Reweighted ℓ_1 ADMM (trace-wise)	noisy	33.37	82.13	67.25
2D-coupled ADMM (TV + RW-ℓ_1)	noisy	33.63	79.74	67.40
Hybrid ResCNN	noisy	41.48	32.27	22.95
Physics-informed 2D U-Net	noisy	44.14	23.78	14.45

by 11–13 dB over the best classical result and reduce RMSE by more than 70%. That the same ordering holds across two Marmousi-2 regions of differing structural complexity further indicates that the gain is driven by the learned refinement rather than by the spatial coupling or by a single favorable crop geometry.

The consistency of the improvement pattern across two geologically different models, one purpose-built and one derived from the established Marmousi-2 benchmark, provides evidence that the ADMM-guided physics-informed neural refinement is not overfitting to a single synthetic geometry. The transferred convolutional features learned on the simpler model remain effective when applied to the more complex Marmousi-2 fault-zone structure.

6.6. Noise-Robustness Sweep

The preceding experiments fix the input noise at a single amplitude-SNR ratio of 8 (≈18.1 dB). To characterize how each method degrades as the data quality changes, the seismic section is regenerated at five amplitude-SNR ratios spanning 2 to 32 (6.0 to 30.1 dB) and all four methods are re-evaluated at each level. To keep the comparison fair, the two networks are not frozen at their nominal training noise; instead each is fine-tuned at the noise level it is evaluated on (warm-started from the deployed per-crop checkpoint, using the same supervised Marmousi loss), so that every method is given data at the input SNR at which it is scored. The classical inversions are deterministic and require no retraining. Figure 20 reports the resulting output-SNR and output-RMSE curves for both Marmousi-2 crops, and Table 4 lists the output SNR values.

Three trends are consistent across both crops. First, the physics-informed U-Net attains the highest reconstruction SNR at every noise level, with the Hybrid ResCNN second, confirming

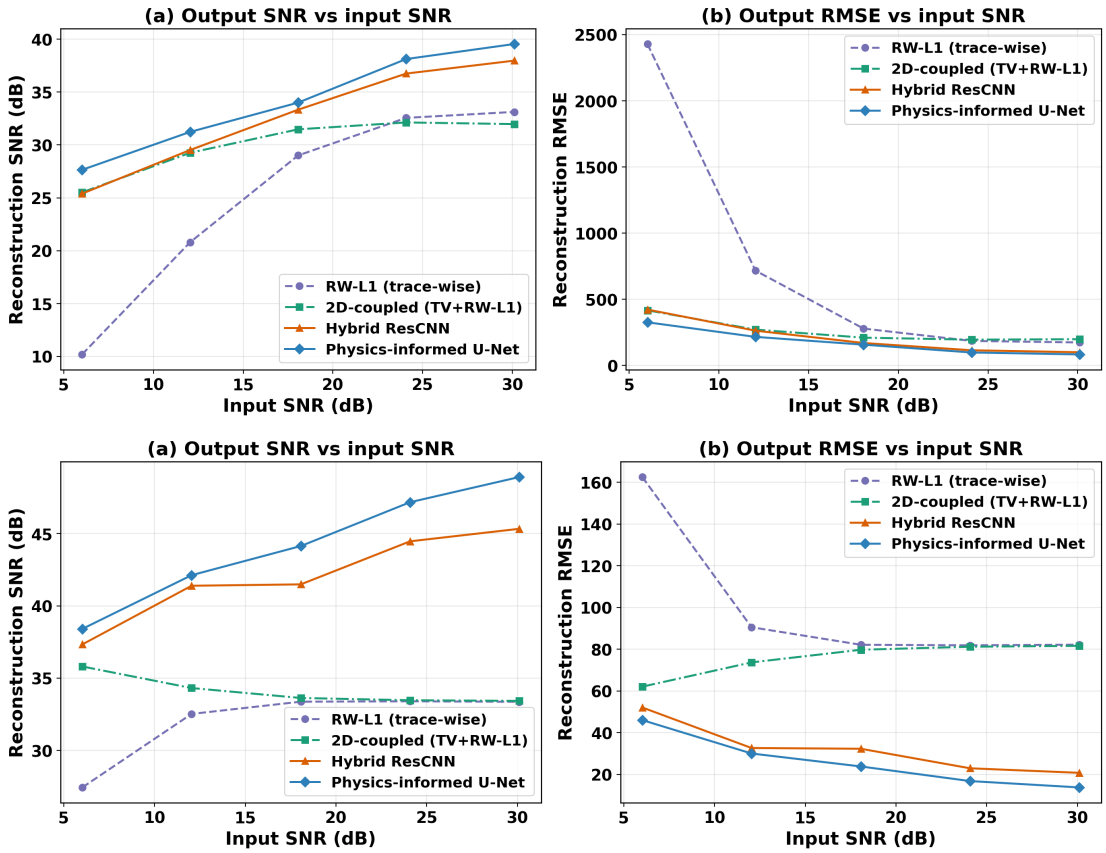

Figure 20: Noise-robustness sweep with matched retraining for the fault-zone crop (top) and the second crop (bottom). Left: reconstruction SNR versus input SNR; right: reconstruction RMSE versus input SNR. Each network is fine-tuned at the noise level it is evaluated on.

Table 4: Reconstruction SNR (dB) as a function of input SNR, with networks retrained at each noise level. Best value in each row is in bold.

Crop	Input SNR (dB)	RW-ℓ_1	2D-coupled	ResCNN	U-Net
Fault zone	6.0	10.19	25.51	25.39	**27.65**
	12.0	20.80	29.26	29.52	**31.23**
	18.1	29.01	31.48	33.33	**34.00**
	24.1	32.56	32.13	36.75	**38.12**
	30.1	33.12	31.97	37.97	**39.54**
Second	6.0	27.44	35.81	37.33	**38.40**
	12.0	32.52	34.31	41.39	**42.12**
	18.1	33.37	33.63	41.48	**44.14**
	24.1	33.40	33.47	44.46	**47.17**
	30.1	33.36	33.43	45.33	**48.91**

that the neural advantage is not an artifact of one favorable noise setting. Second, the margin widens as the data improve: on the fault-zone crop the U-Net leads the best classical result by 2.5 dB at 18.1 dB input SNR but by 6.4 dB at 30.1 dB, because the learned refinement is no longer limited by noise and can exploit fine structure that the sparsity prior smooths away. Third, the gap narrows in the heavy-noise regime: at 6.0 dB input SNR on the fault-zone crop the U-Net leads the 2D-coupled baseline by only 2.1 dB, and the trace-wise reweighted ℓ_1 inversion collapses (10.2 dB) because per-trace sparse deconvolution has no lateral information to suppress strong random noise. The spatially-coupled methods, classical and neural alike, remain far more robust there, which underscores that lateral context, however it is supplied, is the key ingredient under severe noise. This sweep reflects a matched-training protocol; the separate question of how a single deployed model generalizes to noise levels it was *not* trained on is a distinct robustness concern that we note as a limitation in Section 9.

6.7. Wavelet-Mismatch Robustness

All experiments above hand the inversion the true Ricker wavelet. Field data never afford this: the wavelet must be estimated, and residual amplitude, bandwidth, and phase errors are unavoidable. To probe sensitivity to this error, the clean seismic data are generated with the true wavelet, but every method inverts with a *wrong* wavelet, and no network is retrained. This is the realistic deployment setting, in which an already-trained model is applied with an imperfect wavelet estimate. Two error axes are swept independently: (i) a peak-frequency error $f_0' = f_0(1+\delta)$ with $\delta \in \{-0.30, -0.15, 0, 0.15, 0.30\}$, and (ii) a constant-phase rotation of the true wavelet by $\phi \in \{0°, 30°, 60°, 90°\}$ implemented through the analytic signal. Each perturbed wavelet is rescaled to the true wavelet's L_2 norm, so the test isolates shape and phase error from a trivial amplitude mismatch. Figure 21 reports the resulting reconstruction-SNR curves for both crops, and Table 5 lists representative values for the fault-zone crop.

The behavior is consistent across both crops and reveals a clear division of labour. When the wavelet is correct, the physics-informed networks retain their large advantage (U-Net 40.5 dB versus 33.5 dB for the best classical inversion on the fault-zone crop). As soon as the wavelet is wrong, however, the networks lose that advantage and track the trace-wise reweighted ℓ_1 result almost exactly, because the trace-wise inversion is the prior fed to the network: when an incorrect

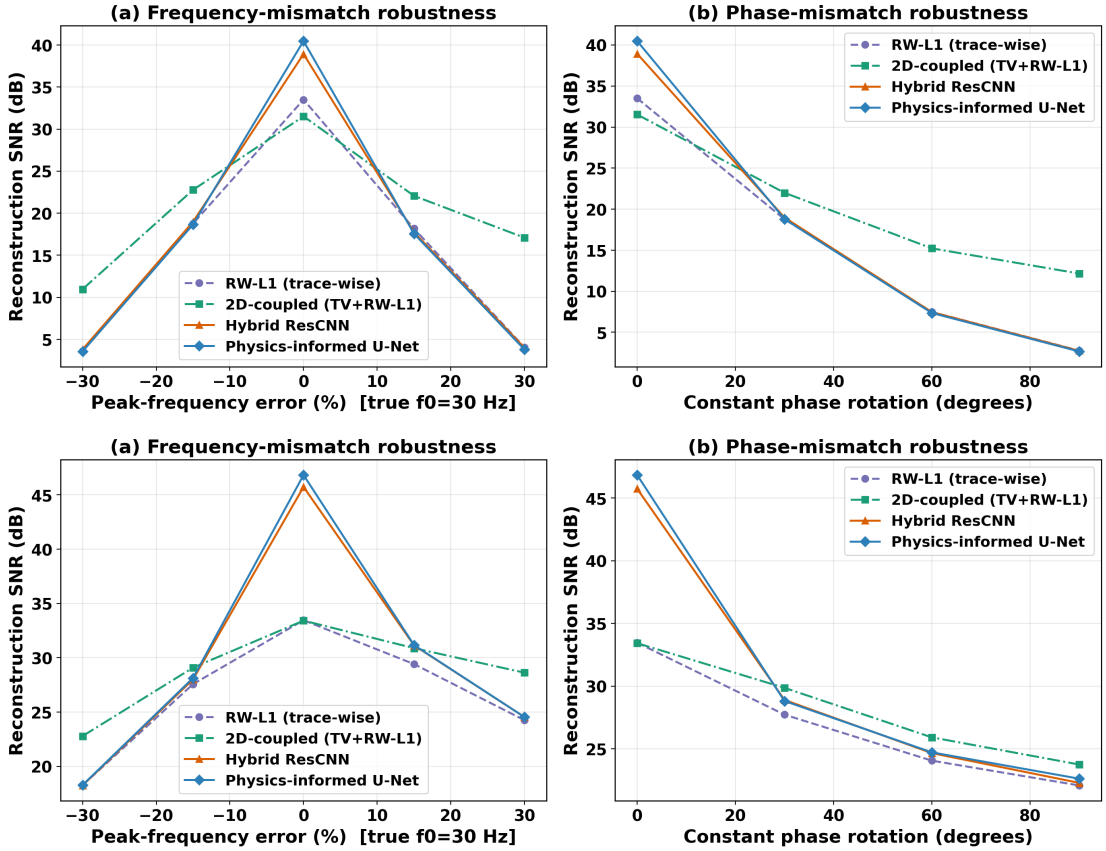

Figure 21: Wavelet-mismatch robustness on clean data (true-wavelet data, wrong-wavelet inversion, no retraining) for the fault-zone crop (top) and the second crop (bottom). Left: peak-frequency error; right: constant-phase rotation. The networks are applied as trained.

Table 5: Reconstruction SNR (dB) under wavelet mismatch on the fault-zone crop. Best value in each row is in bold.

Error axis	Value	RW-ℓ_1	2D-coupled	ResCNN	U-Net
Frequency	−30%	3.61	**10.95**	3.79	3.55
	−15%	18.77	**22.78**	18.97	18.65
	0%	33.51	31.53	38.93	**40.48**
	+15%	18.16	**22.06**	17.73	17.57
	+30%	4.03	**17.08**	3.98	3.78
Phase	0°	33.51	31.53	38.93	**40.48**
	30°	18.77	**21.99**	18.97	18.80
	60°	7.47	**15.22**	7.46	7.35
	90°	2.77	**12.16**	2.74	2.66

wavelet corrupts that prior, the network inherits the corruption and cannot recover information the prior has destroyed. The 2D-coupled baseline, in contrast, is the most robust method at every non-zero error level on both crops, retaining roughly 11–17 dB where the trace-wise and neural reconstructions collapse below 5 dB under severe frequency error or 90° phase rotation. Its lateral total-variation coupling enforces spatial continuity that partially compensates for the per-trace deconvolution error introduced by the wrong wavelet. Two practical implications follow. First, the neural refinement is only as reliable as the wavelet estimate it is given, so accurate wavelet estimation, or wavelet-robust training, is a prerequisite for deploying these networks on field data. Second, the spatially-coupled classical inversion is a valuable safeguard precisely in the regime where the learned models are least trustworthy, which argues for reporting both rather than either alone.

6.8. Seed Sensitivity and Confidence Intervals

The metrics reported above are single-seed values. To verify that the neural advantage is reproducible and not an artifact of a fortunate random seed, each network was retrained five times with independent seeds on the noisy data of both crops, and the spread of the reconstruction metrics was measured. Two initialization protocols are reported: training from random initialization (no transfer), which is the most stringent test of seed independence, and the synthetic-to-Marmousi transfer initialization used for the headline results. Table 6 lists the mean ± standard deviation over the five seeds, with the deterministic classical baselines for reference.

The conclusion is unambiguous. Even from random initialization, every seed of both networks outperforms both classical baselines by a margin that is one to two orders of magnitude larger than the seed-to-seed standard deviation: on the fault-zone crop the worst U-Net seed still exceeds the 2D-coupled baseline by more than 1 dB, and on the second crop the worst seed leads it by over 2.5 dB. The seed variability is modest under random initialization (0.09–2.22 dB in SNR) and becomes negligible under transfer initialization (≤ 0.15 dB), which additionally raises the mean. This shows both that the reported advantage is a stable property of the method rather than a single-seed coincidence, and that the synthetic-to-Marmousi transfer initialization contributes a measurable and consistent improvement in both accuracy and run-to-run stability.

Table 6: Reconstruction metrics on noisy data over five independent training seeds (mean ± sample standard deviation). Classical inversions are deterministic.

Crop	Method	SNR (dB)	RMSE
Fault zone	Reweighted ℓ_1 ADMM (trace-wise)	29.01	278.3
	2D-coupled ADMM (TV + RW-ℓ_1)	31.48	209.5
	ResCNN (random init)	32.73 ± 0.09	181.2 ± 1.9
	U-Net (random init)	33.15 ± 0.27	172.9 ± 5.4
	ResCNN (transfer init)	33.72 ± 0.01	161.8 ± 0.2
	U-Net (transfer init)	35.01 ± 0.06	139.5 ± 0.9
Second	Reweighted ℓ_1 ADMM (trace-wise)	33.37	82.1
	2D-coupled ADMM (TV + RW-ℓ_1)	33.63	79.7
	ResCNN (random init)	38.04 ± 1.26	48.4 ± 7.3
	U-Net (random init)	39.94 ± 2.22	39.6 ± 11.1
	ResCNN (transfer init)	43.13 ± 0.06	26.7 ± 0.2
	U-Net (transfer init)	44.71 ± 0.15	22.3 ± 0.4

7. Computation and Reproducibility

The computational measurements are reported to document reproducibility and feasibility in a low-resource CPU environment, not to make a hardware-specific performance claim. The machine used for the reported timing measurements was a MacBook Pro with an Apple M3 Pro chip, 11 CPU cores (5 performance and 6 efficiency cores), and 18 GB memory, running macOS Tahoe 26.5. Python reported `os.cpu_count()=11`. PyTorch reported five intra-op CPU threads and eleven inter-op threads. In this environment, PyTorch 2.10.0 reported `cuda_available=False` and `mps_available=False`, so the verified timing values below are CPU timings. The numerical environment used Python 3.11.15, NumPy 2.4.6, SciPy 1.17.1, and Matplotlib 3.10.9.

The deterministic ADMM timing was measured on the benchmark data using the same parameters. The clean ADMM run required 21.04 s for 400 traces, with a mean of 33.08 ADMM iterations per trace and a maximum of 65. The noisy ADMM run required 25.56 s, with a mean of 40.25 iterations and a maximum of 85. These runs exactly reproduced the reported ADMM SNR and RMSE values.

Full-run wall-clock time was recorded only for the Attention ResUNet training runs: 186.53 s for clean data and 184.35 s for noisy data. Because full-run timing for the U-Net and ResCNN was not separately recorded, a one-epoch CPU benchmark was performed on the same dataset.

The benchmark is reported as measured per-epoch timing, with a linear 80-epoch equivalent only as a hardware-specific estimate, not as a historical log. Table 7 collects these verified timing values.

Table 7: Verified computational timing on a low-resource CPU environment.

Method	Data	Time	Notes
Reweighted ℓ_1 ADMM	clean	21.04 s	Full deterministic rerun.
Reweighted ℓ_1 ADMM	noisy	25.56 s	Full deterministic rerun.
U-Net	clean	2.69 s/epoch	One-epoch benchmark; 80-epoch equivalent ≈215.16 s.
U-Net	noisy	1.85 s/epoch	One-epoch benchmark; 80-epoch equivalent ≈148.14 s.
ResCNN	clean	2.70 s/epoch	One-epoch benchmark; 80-epoch equivalent ≈216.05 s.
ResCNN	noisy	2.72 s/epoch	One-epoch benchmark; 80-epoch equivalent ≈217.89 s.
Attention ResUNet	clean	186.53 s	Full 80-epoch training time.
Attention ResUNet	noisy	184.35 s	Full 80-epoch training time.

8. Discussion

The results support a balanced reading. On clean synthetic data, the reweighted ℓ_1 ADMM baseline of He et al. (2022) is already strong, because its sparse reflectivity constraint matches impedance sections dominated by discrete interfaces; the neural gains there are real but modest. The benefit of the neural approach is clearest under noise: by exploiting lateral spatial context, the 2D U-Net and ResCNN avoid interpreting random trace-wise perturbations as impedance boundaries, reducing noisy-data RMSE by 44.89% and 41.03% relative to ADMM while preserving the interpretability of sparse inversion.

Architectural complexity is not automatically beneficial. The Attention ResUNet improves over ADMM under noise but does not surpass the simpler U-Net on this compact benchmark, likely because the limited number of training patches and the repetitive layered geometry give attention gates too few distinct spatial contexts. The transfer learning experiment is more decisive: features learned on the controlled synthetic model generalize to the geologically distinct Marmousi-2 fault-zone crop, where the U-Net reduces clean- and noisy-data RMSE by 55.19% and 43.74% relative to trace-wise ADMM. This cross-model consistency indicates transferable physical priors rather than memorization of a single geometry.

The comparison is deliberately conservative. The networks use supervised and gradient-domain terms that the unsupervised classical baselines do not, so the study measures what learned priors achieve against what is achievable without them. Two observations show the gap is not merely an artifact of this supervision: the pretrained weights transfer to the Marmousi-2 crops without refitting and still dominate both baselines, and the spatially-coupled 2D baseline, despite explicit lateral context, remains 6–13 dB short of the networks. The advantage therefore reflects a transferable, geometry-aware learned prior rather than supervision leakage or spatial coupling alone.

9. Limitations

This study validates the framework on two synthetic benchmarks generated from known forward models and wavelets, so the reported gains are controlled numerical evidence rather than proof of field performance; real surveys add uncertainty from wavelet estimation, phase errors, multiples, attenuation, anisotropy, and acquisition footprints. The wavelet-mismatch analysis of Section 6.7 addresses one such factor and shows that the neural refinement is only as reliable as the wavelet estimate it is given, whereas the spatially-coupled classical inversion degrades most gracefully. The training objective also uses impedance-supervision and gradient terms that rely on the known true model; these must be replaced by field-compatible constraints for blind deployment, so the present results are not a fully blind field inversion.

The implementation targets reproducible research rather than production-scale inversion: the trace-wise ADMM baseline uses dense linear solves, and the patch-based networks would need retuning of memory, batch size, and training time for server-scale or 3D surveys. The classical comparison includes both the trace-wise reweighted ℓ_1 inversion and a spatially-coupled 2D baseline (lateral total variation solved by consensus ADMM, Sections 2.1 and 6.4); the latter isolates spatial coupling and does not close the gap to the networks. Other lateral regularizers aligned with local dip remain natural comparators for future work.

The noise-robustness sweep of Section 6.6 uses matched training at each noise level; a single deployed model retains its advantage at higher input SNR but degrades more steeply than the deterministic 2D-coupled baseline under heavy, unseen noise, so out-of-distribution robustness, for example through noise-augmented training, is left for future work. Reported

timings are hardware-specific. The architecture ranking is conditional on model size, noise level, and training schedule, although the five-seed analysis of Section 6.8 confirms that the neural advantage is stable across initializations. Remaining directions include field data, broader wavelet and noise regimes, per-term loss ablation, uncertainty quantification, and transfer to larger benchmark models.

10. Conclusion

This paper presents a 2D acoustic impedance inversion study that combines reweighted ℓ_1 ADMM with physics-informed convolutional neural refinement, validated on both a controlled synthetic benchmark and a cropped fault-zone section from the Marmousi-2 model. The classical ADMM solver provides a sparse and physically interpretable baseline following the formulation of He et al. (2022). The neural models then use the ADMM result as a prior and refine it using wavelet-convolution physics, sparse reflectivity, lateral coherence, and synthetic benchmark supervision.

On the controlled synthetic model, the clean-data experiment shows that ADMM is already strong, with ResCNN and U-Net producing only modest gains. The noisy-data experiment demonstrates the main value of the proposed approach: the 2D U-Net raises SNR by 5.18 dB and reduces RMSE by 44.89% relative to the trace-wise ADMM baseline. Transfer learning to the Marmousi-2 fault-zone crop confirms that these improvements generalize: the U-Net reduces RMSE by 55.19% on clean data and 43.74% on noisy data, with a 5.00 dB SNR gain over ADMM. The results therefore support a balanced claim: ADMM-guided physics-informed neural inversion is most compelling when seismic noise and lateral geological complexity make independent trace-wise sparse inversion insufficient, and the learned features transfer effectively across geologically distinct synthetic benchmarks.

Acknowledgements

The authors gratefully acknowledge the Institute of Geophysics, Polish Academy of Sciences (IGF PAN), for the institutional support and research environment that made this work possible.

Conflict of Interest

The authors declare that they have no conflicts of interest.

Code and Data Availability

The source code and numerical data that support the findings of this study are openly available on GitHub at `https://github.com/kumarDeepak-su/Physics-Informed-Neural-Network`.

References

Boyd, S., Parikh, N., Chu, E., Peleato, B., Eckstein, J., 2011. Distributed optimization and statistical learning via the alternating direction method of multipliers. Foundations and Trends in Machine Learning 3 (1), 1–122. https://doi.org/10.1561/2200000016.

Candès, E. J., Wakin, M. B., Boyd, S. P., 2008. Enhancing sparsity by reweighted ℓ_1 minimization. Journal of Fourier Analysis and Applications 14 (5), 877–905. https://doi.org/10.1007/s00041-008-9045-x.

Cooke, D. A., Schneider, W. A., 1983. Generalized linear inversion of reflection seismic data. Geophysics 48 (6), 665–676. https://doi.org/10.1190/1.1441497.

He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition, in: Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), pp. 770–778. https://doi.org/10.1109/CVPR.2016.90.

He, L., Wu, H., Wen, X., You, J., 2022. Seismic acoustic impedance inversion using reweighted ℓ_1-norm sparse constraint. IEEE Geoscience and Remote Sensing Letters 19, art. no. 8025505, 1–5. https://doi.org/10.1109/LGRS.2022.3168015.

Lindseth, R. O., 1979. Synthetic sonic logs: a process for stratigraphic interpretation. Geophysics 44 (1), 3–26. https://doi.org/10.1190/1.1440922.

Martin, G. S., Wiley, R., Marfurt, K. J., 2006. Marmousi2: An elastic upgrade for Marmousi. The Leading Edge 25 (2), 156–166. https://doi.org/10.1190/1.2172306.

Oktay, O., Schlemper, J., Le Folgoc, L., Lee, M., Heinrich, M., Misawa, K., Mori, K., McDonagh, S., Hammerla, N. Y., Kainz, B., Glocker, B., Rueckert, D., 2018. Attention U-Net: Learning where to look for the pancreas, in: Proc. Medical Imaging with Deep Learning (MIDL). arXiv:1804.03999.

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S., 2019. PyTorch: An imperative style, high-performance deep learning library, in: Advances in Neural Information Processing Systems, vol. 32.

Raissi, M., Perdikaris, P., Karniadakis, G. E., 2019. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. Journal of Computational Physics 378, 686–707. https://doi.org/10.1016/j.jcp.2018.10.045.

Rudin, L. I., Osher, S., Fatemi, E., 1992. Nonlinear total variation based noise removal algorithms. Physica D: Nonlinear Phenomena 60 (1–4), 259–268. https://doi.org/10.1016/0167-2789(92)90242-F.

Ronneberger, O., Fischer, P., Brox, T., 2015. U-Net: Convolutional networks for biomedical image segmentation, in: Medical Image Computing and Computer-Assisted Intervention, MICCAI 2015 (Lecture Notes in Computer Science, vol. 9351), Springer, pp. 234–241. https://doi.org/10.1007/978-3-319-24574-4_28.

Taylor, H. L., Banks, S. C., McCoy, J. F., 1979. Deconvolution with the ℓ_1 norm. Geophysics 44 (1), 39–52. https://doi.org/10.1190/1.1440921.

Wang, Y., 2010. Seismic impedance inversion using ℓ_1-norm regularization and gradient descent methods. Journal of Inverse and Ill-Posed Problems 18 (7), 823–838. https://doi.org/10.1515/jiip.2011.005.

Wu, B., Meng, D., Wang, L., Liu, N., Wang, Y., 2020. Seismic impedance inversion using fully convolutional residual network and transfer learning. IEEE Geoscience and Remote Sensing Letters 17 (12), 2140–2144. https://doi.org/10.1109/LGRS.2019.2963106.