ACCSCIENCE PUBLISHING

Journal of Seismic Exploration

ARTICLE

A two-branch convolutional neural network for pre-stack elastic parameter inversion of coal-bearing gas reservoirs

Fei Li[1,2], Mengbo Zhang[1,2], Qiang Liang[1], Xiaojie Cui[1,2], Na Ni[1], Qingzhou Zhang[2], and Yongheng Zhang[3]*

[1]Exploration and Development Research Institute, Changqing Oilfield Company, PetroChina, Xi'an, Shaanxi, China

[2]National Engineering Laboratory for Exploration and Development of Low-Permeability Oil and Gas Fields, College of Petroleum Engineering, Xi'an, Shaanxi, China

[3]College of Geophysics, Chengdu University of Technology, Chengdu, Sichuan, China

Abstract

Accurate estimation of Poisson's ratio is essential for the characterization of coal-bearing gas reservoirs, particularly in thin-bed and low signal-to-noise ratio environments where conventional elastic impedance (EI) inversion suffers from wavelet interference, limited resolution, and reliance on linearization assumptions. To address these limitations, we develop a physics-guided two-branch convolutional neural network (TB-CNN) that directly predicts Poisson's ratio by jointly integrating EI-inverted P-wave velocity, S-wave velocity, and density with small-, medium-, and large-angle seismic stacks. The first branch provides geologically consistent, physics-informed background trends, while the second branch captures thin-bed-sensitive reflectivity features and amplitude tuning effects. The fused latent representation is explicitly regularized using empirical rock-physics relationships to ensure physical plausibility and enhanced generalization. Field validation on the 8# coal seam of the Ordos Basin demonstrates that the proposed TB-CNN improves vertical resolution, sharpens seam boundary delineation, and better preserves thickness variations compared with EI inversion and a single-branch CNN. Near-well comparisons show higher correlation with log-derived Poisson's ratio, while lateral slices reveal improved continuity and thin-layer detectability. These results confirm that combining physics-guided stability with data-driven resolution provides a robust and interpretable framework for Poisson's ratio inversion in thin coal seams and holds promise for broader applications in unconventional gas reservoir prediction.

Keywords: Coalbed methane; Elastic parameter inversion; Pre-stack seismic data; Convolutional neural network

***Corresponding author:**
Yongheng Zhang
(zhangyhzjk@163.com)

Citation: Li F, Zhang M, Liang Q, *et al*. A two-branch convolutional neural network for pre-stack elastic parameter inversion of coal-bearing gas reservoirs. *J Seismic Explor.* 2026;35(2):025390076.
doi: 10.36922/JSE025390076

Received: September 22, 2025

Revised: November 10, 2025

Accepted: January 28, 2026

Published online: April 23, 2026

Copyright: © 2026 Author(s). This is an Open-Access article distributed under the terms of the Creative Commons Attribution License, permitting distribution, and reproduction in any medium, provided the original work is properly cited.

Publisher's Note: AccScience Publishing remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

1. Introduction

Elastic parameters—especially Poisson's ratio (ν) and Young's modulus (E)—are fundamental for assessing the mechanical behavior, brittleness, and gas-bearing potential of coal reservoirs. In coalbed methane (CBM) development, reliable ν and

E volumes support *in situ* stress prediction, hydraulic-fracturing design, and dynamic reservoir simulation, and thus are crucial for identifying sweet spots and evaluating completion strategies in thinly bedded, laterally heterogeneous coal seams, such as those in the Ordos Basin, China.[1,2] However, deriving these parameters from seismic data remains challenging due to insufficient resolution and stability.

Conventional pre-stack inversion frameworks—grounded in amplitude versus offset (AVO)/amplitude variation with angle theory and linearized reflectivity equations—are well established.[3,4] Yet, they are often sensitive to noise, wavelet, and angle calibration, and acquisition footprint, and typically rely on smooth initial models and strong regularization. As a result, the inverted P-wave velocity (Vp), S-wave velocity (Vs), and density (ρ) volumes tend to be band-limited and over-smoothed, which limits thin-bed detection and blurs sharp elastic contrasts. These limitations are particularly restrictive in coal-bearing settings where the target seams are thin and embedded within a single reflection lobe, as is the case for the 8# coal seam in the northern Yishan Slope of the Ordos Basin.[5,6] In the standard workflow, ν is conventionally derived from the Vs/Vp ratio, which inevitably causes errors from intermediate elastic parameter inversions to propagate and degrade the final prediction accuracy. Hence, any bias from intermediate inversions propagates and degrades the final attribute used for brittleness and stress analysis.[7-10]

Recent advances in deep learning (DL) offer an alternative by directly learning nonlinear mappings between seismic observations and rock properties.[6,11-14] In particular, convolutional neural networks (CNNs) have shown improved recovery of high-frequency components and robustness to moderate noise relative to purely descent-based or strongly linearized inversions.[7-11] At the same time, physics-guided DL—via architectural design or loss regularization, where additional physics-based penalty terms are introduced into the loss function to enforce consistency with rock-physics or wave-propagation principles—has emerged to couple data-driven representations with geophysical constraints, improving interpretability and generalization.[15-19] Physics-guided DL incorporates physical priors into data-driven models through architectural design or loss regularization, whereas physics-constrained inversion directly embeds forward modeling into the inversion. Generative and uncertainty-aware models further help capture non-uniqueness and quantify reliability.[20] Nevertheless, many DL applications still rely on post-stack inputs or treat pre-stack traces independently, underutilizing the angle-dependent reflectivity and thin-bed interference information critical to coal reservoirs. Moreover, most existing physics-guided CNNs or generative adversarial network (GAN)-based inversions are developed for generalized impedance prediction or post-stack attributes, with limited adaptability to thin coal seams where reflections overlap within a single lobe, and signal-to-noise ratio (SNR) is intrinsically low. Similarly, attention-fusion strategies enhance multi-scale features but often lack explicit mechanisms to enforce rock-physics plausibility, which constrains their reliability in complex CBM settings.

To bridge these gaps, we develop a two-branch CNN (TB-CNN) for direct ν prediction in coal-bearing gas reservoirs. The first branch ingests elastic impedance (EI)-inverted Vp, Vs, and ρ to provide stable, physics-consistent large-scale trends that anchor the solution space, while the second branch processes small-, medium-, and large-angle partial stacks that preserve high-resolution waveform variations sensitive to thin-bed effects. The two streams are fused in a shared latent space with soft rock-physics regularization derived from well-log relationships, enabling a nonlinear mapping directly to ν and mitigating error propagation from intermediate inversions. Conceptually, this design differs from prior physics-guided CNNs or GAN-based inversions in three respects: (i) it explicitly integrates angle-dependent reflectivity information often ignored in post-stack or single-trace DL workflows; (ii) it embeds empirical rock-physics constraints as soft priors, improving interpretability and avoiding geologically implausible predictions common in unconstrained GAN outputs; and (iii) it is tailored to thin-bed coal seams where conventional inversions fail due to wavelet interference. These design choices allow TB-CNN to leverage the complementary strengths of physics-derived attributes and multi-angle data-driven features: the former stabilizes training by providing low-frequency, physically consistent background trends, while high-resolution information is primarily recovered from multi-angle seismic data through nonlinear feature learning, thereby enhancing vertical resolution and lateral detectability beyond what linearized inversions typically deliver.[7,11,21]

We validate the framework on field data from the northern Yishan Slope of the Ordos Basin, where the 8# coal seam (average thickness ~8.3 m) exhibits a strong impedance contrast but a thin-bed seismic response embedded within a single lobe. Using three-dimensional (3D) pre-stack seismic data and a suite of wells for supervision and blind testing, we show that TB-CNN-predicted ν volumes delineate seam boundaries more sharply, better capture gradual thickness variations and localized thinning, and correlate more closely with log-

derived ν than conventional EI-based three-parameter inversions followed by ratio computation. Horizon-slice analysis further indicates superior lateral continuity and sensitivity to subtle heterogeneity. These results substantiate that integrating physics-informed trends with multi-angle data-driven features can reduce dependence on linearized approximations and smooth initial models, thereby enhancing the reliability of ν prediction for CBM reservoir characterization in low SNR, thin-bed environments.[1,2]

2. Methods

This study focuses on predicting elastic parameters in coal-bearing gas reservoirs, which are often characterized by thin beds, strong impedance contrasts, and low SNR. These geological and seismic conditions make inversion particularly challenging. Conventional pre-stack inversion methods rely on linearized forms of the Zoeppritz equations, where high-frequency information is frequently suppressed due to smoothing constraints and approximation errors. In practice, the smooth initial model, together with the regularization applied during inversion, further attenuates thin-bed responses and weak reflections. As a result, conventional approaches typically struggle to preserve fine-scale heterogeneity and fail to accurately delineate thin coal seams. This limitation is especially evident when predicting ν, which is highly sensitive to subtle variations in both Vp and Vs.

To mitigate these limitations, we propose the TB-CNN model that integrates physics-constrained inversion outputs with high-resolution seismic data in a unified learning framework, shown in Figure 1. The first branch (Branch 1) ingests the Vp, Vs, and ρ volumes obtained through EI inversion, providing stable, large-scale trends consistent with rock-physics principles. The second branch (Branch 2) directly utilizes stacked pre-stack seismic data from three incidence-angle ranges (small, medium, and large), thereby retaining high-frequency waveform details and thin-bed sensitivity without dependence on inversion linearization.

Feature representations extracted by the two branches are fused within a shared latent space, where empirical rock-physics constraints derived from well-log analysis are implicitly embedded to ensure geological plausibility. The fused features are subsequently mapped to the target elastic parameters through a sequence of fully connected (FC) layers. By jointly leveraging the stability and interpretability of physics-based trends from Branch 1 and the fine-scale, noise-resilient information from Branch 2, the proposed TB-CNN reduces reliance on linearized forward models, mitigates bias introduced by smooth initial models, and enhances the detectability of subtle elastic variations in complex coal-bearing gas reservoirs.

2.1. Branch 1 design: Physics-constrained elastic parameter input

The first branch of the proposed TB-CNN incorporates physics-based prior information obtained from conventional EI inversion, using the inverted Vp, Vs,

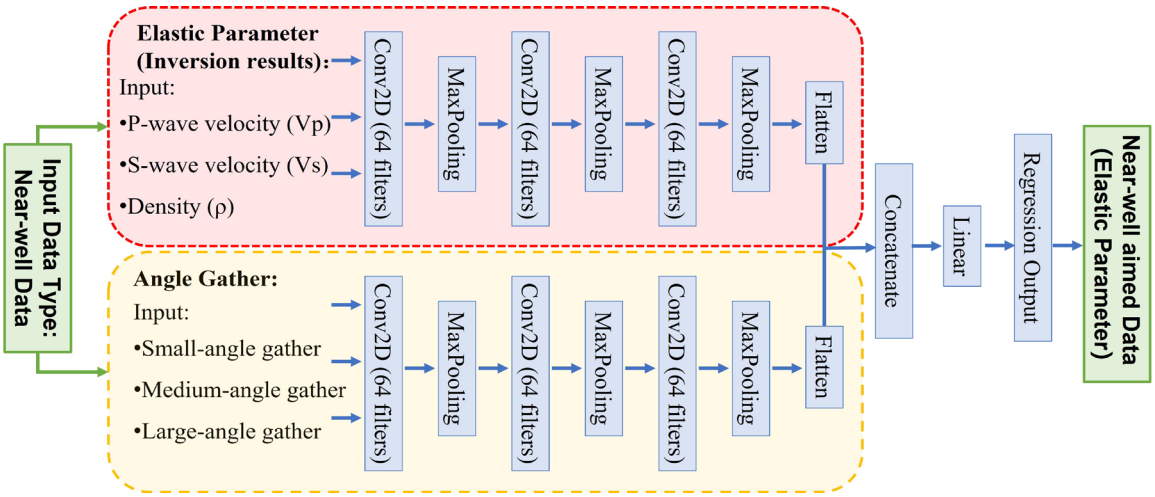

Figure 1. Schematic workflow of the proposed two-branch convolutional neural network-based inversion methodology
Abbreviation: Conv2D: Two-dimensional convolution.

and ρ volumes as inputs (Figure 2). These parameters are estimated through a multi-angle inversion procedure grounded in the linearized Zoeppritz equations and the EI formulation, thereby ensuring consistency with established rock-physics relationships.

To extend acoustic impedance to non-normal incidence, the concept of EI was introduced, which expresses the angle-dependent reflectivity R(θ) in an impedance-like form consistent with the Aki–Richards approximation.[22] Given pre-stack seismic reflectivity R(θ) at incidence angle θ, the linearized Aki–Richards equation can be written as **Equation 1**:

$$R(\theta) \approx A + B\sin^2\theta + C(\tan^2\theta - \sin^2\theta) \tag{1}$$

where A, B, and C depend on the fractional contrasts in Vp, Vs, and ρ across an interface, as shown in **Equations 2–4**:

$$A = \frac{1}{2}\left(\frac{\Delta V_P}{V_P} + \frac{\Delta\rho}{\rho}\right) \tag{2}$$

$$B = \frac{\Delta V_p}{V_p} - 2\left(\frac{V_s}{V_p}\right)^2\left(\frac{\Delta V_s}{V_s} + \frac{\Delta\rho}{\rho}\right) \tag{3}$$

$$C = \frac{1}{2}\frac{\Delta V_p}{V_p} \tag{4}$$

For a reference angle θ_0, the EI (θ_0) is defined as **Equation 5**:

$$EI(\theta_0) = V_p^{a(\theta_0)} V_s^{b(\theta_0)} \rho^{c(\theta_0)} \tag{5}$$

where K = $(Vs/Vp)^2$ is the shear-to-compressional velocity ratio, and by matching the logarithmic contrast of EI (θ_0) with the Aki–Richards expression, the exponents are obtained as **Equations 6–8**:

$$a(\theta_0) = 1 + \tan^2\theta_0 \tag{6}$$

$$b(\theta_0) = -8K\sin^2\theta_0 \tag{7}$$

$$c(\theta_0) = 1 - 4K\sin^2\theta_0 \tag{8}$$

Performing EI inversion over multiple reference angles enables the joint recovery of Vp, Vs, and ρ volumes that preserve large-scale geological trends and conform to rock-physics constraints derived from well logs. However, such inversions remain susceptible to smoothing effects and linearization errors, particularly in thin coal seams and highly heterogeneous zones. The EI formulation thus preserves the physical consistency between angle-dependent reflectivity and elastic parameters, providing a stable and interpretable input for the physics-constrained branch of the TB-CNN.

In the TB-CNN framework, these three inverted volumes are concatenated into a three-channel input tensor, as shown in **Equation 9**:

$$X_1 \in \mathrm{R}^{H\times W\times 3} = [V_p, V_s, \rho] \tag{9}$$

where H and W denote the spatial dimensions in inline–crossline or time–depth coordinates. Branch 1 applies

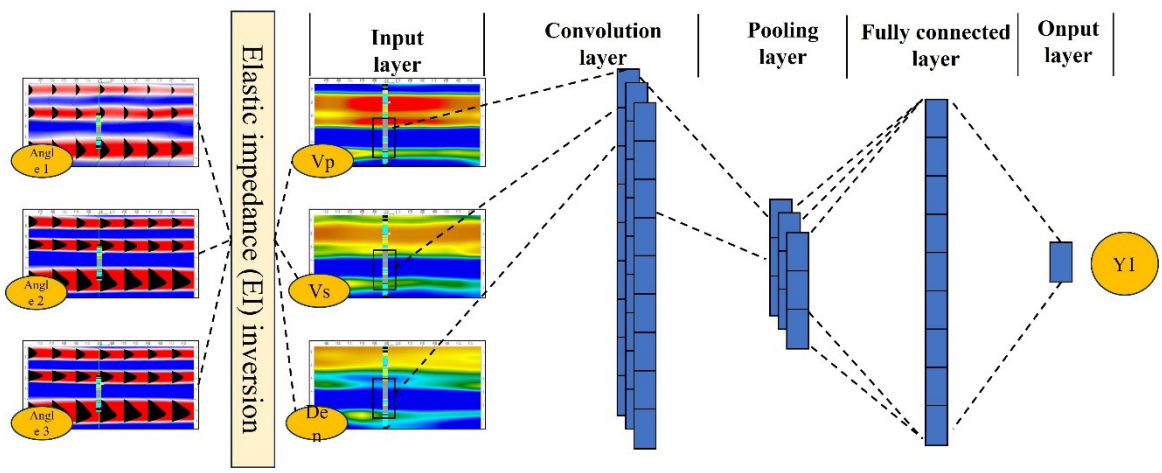

Figure 2. Branch 1 convolutional structure diagram
Abbreviations: ρ: Density; Vp: P-wave velocity; Vs: S-wave velocity.

a sequence of two-dimensional convolutional layers, as shown in **Equation 10**:

$$F_1^{(l)} = \sigma\left(W^{(l)} * F_1^{(l-1)} + b^{(l)}\right), l = 1, \ldots, L_1 \tag{10}$$

where $F_1^{(l)}$ = X1, W (1) and b (1) are the trainable convolution kernels and biases, and σ(·) is a nonlinear activation function (e.g., Rectified Linear Unit [ReLU]). These operations progressively extract multi-scale spatial features that encode the physics-consistent elastic property patterns from the inversion-derived parameters.

After the convolutional feature extraction, the final feature maps of Branch 1 are flattened and projected into a latent representation, as shown in **Equation 11**:

$$z_1 = \text{Flatten}\left(F_1^{(L_1)}\right) W_z + b_z \tag{11}$$

where W_z and b_z are trainable parameters. This latent vector z_1 encapsulates the physics-consistent information from inversion-derived Vp, Vs, and ρ, and serves as one of the two complementary inputs to the feature-fusion module.

By leveraging EI-derived Vp, Vs, and ρ as inputs, Branch 1 preserves physically meaningful structural trends while allowing subsequent network layers to learn data-driven corrections for systematic inversion biases. The resulting feature embeddings are later fused with high-resolution, waveform-sensitive features from Branch 2 to improve thin-bed detectability and robustness under low-SNR conditions.

2.2. Branch 2 design: Multi-angle seismic input (data-driven)

The second branch of the proposed TB-CNN directly ingests three incidence-angle seismic volumes and learns a purely data-driven mapping from raw seismic amplitudes to the target elastic parameter(s), without invoking any explicit AVO formulation. The three input volumes are obtained from partial-angle stacks (small, medium, and large incidence angles) extracted after standard pre-stack processing, including velocity analysis, NMO correction, and angle-domain binning. Each angle stack captures complementary reflectivity patterns arising from different incidence angles, enabling the network to exploit both near-ofset continuity and far-ofset contrast sensitivity. These three stacks are spatially co-registered with the inversion volumes in Branch 1, ensuring consistent supervision windows (Figure 3).

The three partial-angle stacks are concatenated to form a three-channel input tensor, as shown in **Equation 12**:

$$X_2 \in \mathrm{R}^{H \times W \times 3} = [S_g, S_m, S_l] \tag{12}$$

where S_g, S_m, and S_l denote the small-, medium-, and large-angle stacks, respectively, and $H \times W$ corresponds to the spatial sampling in inline/crossline or trace/time dimensions.

Branch 2 applies a series of two-dimensional convolutional layers designed to extract short-lag waveform features and thin-bed interference signatures. Let $F_2^{(0)} = X_2$ be the input to the branch. The convolutional

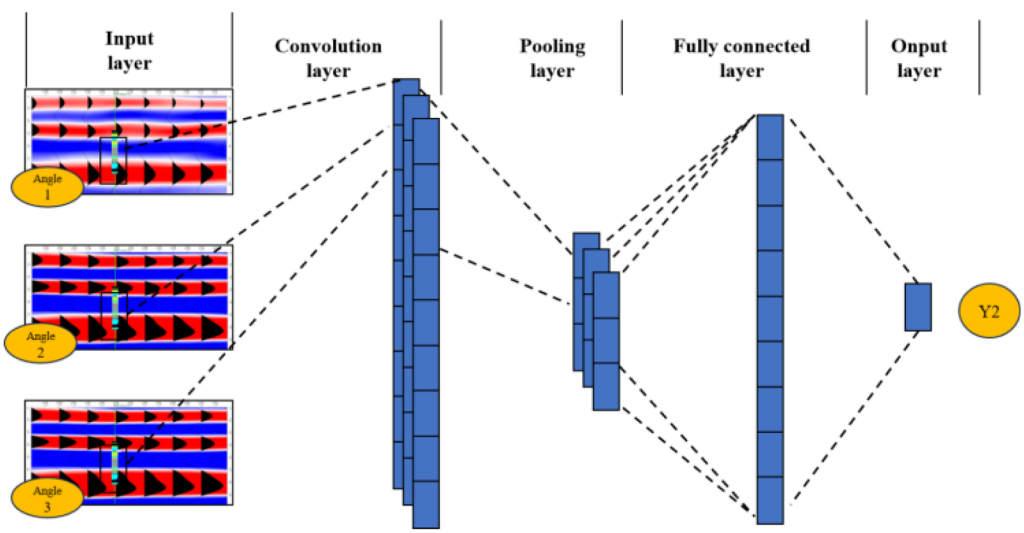

Figure 3. Branch 2 convolutional structure diagram

feature extraction proceeds as **Equation 13**:

$$F_2^{(l)} = \sigma\left(W^{(l)} * F_2^{(l-1)} + b^{(l)}\right), l = 1, \ldots, L_2 \tag{13}$$

where $W^{(l)}$ and $b^{(l)}$ are the learnable kernels and biases, * denotes the two-dimensional convolution operator, and σ(·) is a nonlinear activation (e.g., ReLU). Kernel sizes are chosen to emphasize short-period oscillations along the temporal axis (e.g., 1 × 2 or 3 × 3), which are particularly sensitive to thin-bed tuning effects. Pooling layers may be interleaved to reduce temporal redundancy and improve robustness to random noise.

After the final convolutional block, the feature maps are flattened and linearly projected into a latent representation, as shown in **Equation 14**:

$$z_2 = \text{Flatten}(F_2^{(L_2)})W_z + b_z \tag{14}$$

where W_z and b_z are trainable parameters. This vector z_2 encapsulates high-resolution reflectivity information derived directly from the multi-angle seismic data and complements the physics-consistent z_1 produced by Branch 1. Both z_1 and z_2 are subsequently fused in the shared latent space to generate the final elastic parameter predictions.

2.3. Feature fusion and output prediction

After the independent processing of the physics-informed branch (Branch 1) and the purely data-driven seismic branch (Branch 2), their respective embeddings z_1 and z_2 are concatenated to form a joint latent representation, as shown in **Equation 15**:

$$z = \text{Concat}[z_1, z_2] \in \mathrm{R}^{d_1+d_2} \tag{15}$$

where d_1 and d_2 denote the dimensionalities of z_1 and z_2, respectively. This concatenation preserves complementary information: z_1 encodes large-scale geological trends and physically interpretable elastic responses from inversion results, whereas z_2 provides fine-scale seismic patterns and thin-bed interference signatures extracted directly from waveform data.

2.3.1. Joint feature refinement

The fused latent vector z is subsequently passed through a sequence of FC layers to enable nonlinear cross-branch interaction and hierarchical abstraction. Each FC layer is followed by batch normalization (BN) and a nonlinear activation function (ReLU) to improve convergence stability and prevent gradient saturation. Dropout layers (drop probability 0.2–0.3) are optionally applied after BN to enhance generalization. The transformation is expressed as **Equation 16**:

$$h^{(l)} = \sigma\left(\text{BN}\left(w^{(l)}h^{(l-1)} + b^{(l)}\right)\right), l = 1, \ldots, L \tag{16}$$

with $h^{(0)} = z$, $W^{(l)}$ and $b^{(l)}$ are trainable parameters, and σ(·) denotes the activation.

2.3.2. Prediction layer

The final FC layer maps the refined feature vector h (L) to the predicted elastic parameter(s) $\hat{y}$, as shown in **Equation 17**:

$$\hat{y} = w^{(o)}h^{(L)} + b^{(o)} \tag{17}$$

where $\hat{y}$ can be a scalar (e.g., a single elastic modulus) or a vector (e.g., Vp, Vs, and ρ jointly). Model training minimizes the mean squared error (MSE) loss between network predictions and well-log-derived ground truth, as shown in **Equation 18**:

$$\mathcal{L}_{\text{MSE}} = \frac{1}{N}\sum_{i=1}^{N}\left\|\hat{y}^{(i)} - y^{(i)}\right\|_2^2 \tag{18}$$

where N is the number of training samples.

2.3.3. Advantages of dual-branch fusion

The proposed fusion architecture leverages the complementary nature of the two branches:

(i) Physics-informed branch (Branch 1): Provides large-scale structural context and physically interpretable elastic attributes, ensuring that predictions remain geophysically consistent.

(ii) Seismic-driven branch (Branch 2): Supplies high-resolution, noise-resilient features that capture subtle waveform and stratigraphic variations not recoverable by conventional inversion.

By unifying these representations within a shared latent space, the network achieves enhanced robustness to noise, improved thin-bed resolution, and increased predictive accuracy under low SNR conditions.

2.4. Network architecture details

To ensure reproducibility, the complete TB-CNN architecture is summarized in Table 1. Both branches adopted a symmetric design with two convolutional blocks, each followed by batch normalization and ReLU activation. Pooling layers were interleaved to enhance generalization, while the fusion stage employed two FC layers prior to the final regression output.

Although the TB-CNN employs two parallel branches, each branch contains only two convolutional blocks with moderate kernel dimensions and channel counts, forming a lightweight architecture with fewer than 1.2 million trainable parameters. In addition, the network operates on fixed-size sliding windows centered on the target coal-bearing interval rather than the entire 3D volume. As only three representative partial-angle stacks (small, medium, and large) were used, the effective data volume is about three times that of a post-stack dataset but far smaller than the full pre-stack gathers. Together, these design choices ensure that the computational cost remains moderate and the network can be efficiently trained and deployed on standard graphics processing unit hardware.

2.5. Soft rock-physics regularization

To ensure geological plausibility, we embedded empirical rock-physics constraints into the learning process as a soft regularization term. Specifically, we constructed a penalty L_{phys} based on the deviation of predicted elastic relationships from well-log-derived regressions, such as **Equation 19**:

$$L_{phys} = \frac{1}{M}\sum_{j=1}^{M}\left(\alpha\left(\frac{V_s}{V_p}\right)_j - \beta\rho_j - \gamma\right)^2 \tag{19}$$

where $(Vs/Vp)_j$ and ρ_j denote the predicted parameters at location j, while α, β, and γ are the regression coefficients fitted from the well-log data. This formulation penalizes predictions that violate empirical Vp–Vs–ρ relationships, thereby reducing the likelihood of nonphysical solutions.

The final loss function is a weighted sum of data misfit and physics-based regularization, as shown in **Equation 20**:

$$L = L_{MSE} + \lambda L_{phys} \tag{20}$$

where L_{MSE} is the conventional mean squared error, L_{phys} is the soft rock-physics penalty, and λ is a hyperparameter (0.1–0.5 in our experiments) that balances accuracy and physical consistency. It should be noted that the proposed regularization represents an empirical rock-physics consistency constraint rather than a hard physical constraint derived from forward modeling or wave-equation operators.

This strategy differs from hard constraints in that deviations from empirical trends are not strictly forbidden, but gradually penalized. As a result, the network retains flexibility to adapt to local heterogeneities while preserving global consistency with rock-physics knowledge.

3. Field data test

3.1. Study area and data

The study area is located in the northern segment of the Yishan Slope, Ordos Basin, spanning Wushen Banner in Inner Mongolia and Yuyang District and Mizhi County in Shaanxi Province (Figure 4). The structural setting is characterized by a relatively stable slope belt, where Carboniferous–Permian coal-bearing strata are widely developed. The target reservoir was the 8# coal seam of the Shanxi Formation, with a thickness ranging from 5.8 m to 9.7 m (average 8.3 m). This seam serves as the primary production horizon for CBM development in the region and is distinguished by a pronounced acoustic impedance

Table 1. Detailed architecture of the proposed TB-CNN

Layer	Branch	Input dimension	Kernel/Stride	Output dimension
Conv1	Branch 1 (EI inputs)	H × W × 3	3 × 3/1	H × W × 32
Conv2	Branch 1	H × W × 32	3 × 3/1	H × W × 64
Pooling	Branch 1	H × W × 64	2 × 2/2	$\frac{H}{2}\times\frac{W}{2}\times 64$
Conv1	Branch 2 (Angle stacks)	H × W × 3	3 × 3/1	H × W × 32
Conv2	Branch 2	H × W × 32	3 × 3/1	H × W × 64
Pooling	Branch 2	H × W × 64	2 × 2/2	$\frac{H}{2}\times\frac{W}{2}\times 64$
Flatten	Both branches	-	-	$1 \times (d_1 + d_2)$
FC1	Fusion	$d_1 + d_2$	-	128
FC2	Fusion	128		64
Output	Fusion	64		1 (Poisson's ratio)

Note: H × W denotes the input patch dimension, which is consistent across both branches.
Abbreviations: Conv: Convolution; FC: Fully connected; TB-CNN: Two-branch convolutional neural network.

contrast with surrounding lithologies.

To illustrate the seismic response and geological continuity of the target seams, well–seismic ties from three representative boreholes—Shuang50–45, Shuang133, and Shuang48–45—are shown in Figure 5. In all three wells, the 8# coal seam exhibited a strong impedance contrast with the surrounding mudstone–sandstone–limestone assemblage, producing a high-amplitude composite reflection with polarity reversal indicative of thin-bed tuning. This consistent response across multiple wells validates the seam interpretation and provides key calibration data for pre-stack inversion and TB-CNN training.

The regional continuity of the coal-bearing sequence is further demonstrated by the cross-well seismic section in Figure 6. Four key composite reflection events were identified. Given that the 5# and 8# coal seams are sub-wavelength thin beds (typically thickness < λ/8), their seismic responses are dominated by tuning interference rather than by isolated interface reflections. Consequently, the observed positive and negative phases represent composite waveforms generated by constructive and destructive superposition of closely spaced boundaries, and should not be interpreted as discrete reflections from the seam roof and floor. In this study, these composite events were used only to delineate a consistent inversion window for subsequent elastic parameter prediction, rather than to infer detailed coal–seam interface geometries.

Within this interval, the 8# seam maintains remarkable lateral continuity among the three wells, though thin-bed interference blurs the roof and floor reflections, limiting vertical resolution. These features highlight both the geological importance of the 8# seam as a major CBM reservoir and the inherent limitations of conventional pre-stack inversion in resolving thin coal seams.

To address these challenges, this study applied the proposed TB-CNN to the field dataset. By integrating physics-constrained inversion products with multi-angle seismic inputs, the TB-CNN was designed to enhance thin-bed resolution and improve ν inversion for CBM reservoir characterization.

3.2. Network training and direct Poisson's ratio prediction

Poisson's ratio is a highly diagnostic elastic parameter for coal-bearing formations, as it reflects both lithological variations and fluid content. In rock physics, ν can be expressed in terms of the Vs/Vp ratio as shown in **Equation 21**:

$$\nu = \frac{V_p^2 - 2V_s^2}{2(V_p^2 - V_s^2)} = \frac{1}{2} \cdot \frac{1 - 2\left(\frac{V_s}{V_p}\right)^2}{1 - \left(\frac{V_s}{V_p}\right)^2} \tag{21}$$

which provides the physical basis for deriving ν from well-log measurements. In this study, however, ν was not computed from post-inverted seismic velocities. Instead, it was directly predicted by the TB-CNN, with the above relationship serving only to define the supervision labels and guide physical interpretation.

In the context of CBM exploration, accurate estimation of ν is particularly important because it governs rock brittleness, fracture propagation behavior, and gas-bearing potential. Thin coal seams present a persistent challenge for conventional EI inversion: their seismic responses are commonly embedded within a single tuning lobe, and the resulting Vp and Vs estimates are often over-smoothed. Consequently, ν computed from EI-derived velocities suffers from blurred seam boundaries and limited sensitivity to thickness variations and local thinning. To address these limitations, we adopted a direct-inversion strategy in which TB-CNN is trained to predict ν directly from multi-source seismic inputs, thereby avoiding error accumulation associated with sequential velocity estimation.

The TB-CNN integrates two complementary input domains: (i) EI-inverted elastic parameters (Vp, Vs, and ρ) as physics-consistent prior information (Branch 1), and (ii) partial-angle seismic stacks as high-resolution reflectivity inputs (Branch 2). Figure 7A–C shows the small-, medium-, and large-angle stacks used as Branch 2 inputs, while Figure 8A–C displays the EI-inversion products serving as Branch 1 inputs. This dual-input design enables the network to preserve large-scale elastic trends while exploiting high-frequency waveform variations associated with thin-bed tuning.

The training dataset was constructed using two wells with reliable log coverage across the 8# coal seam, located at approximately 1,400 ms two-way travel time. Around each well, fixed-size two-dimensional patches were extracted in both inline and crossline directions, with patch size and stride selected to balance sample diversity and spatial continuity. Each patch was centered on the well trajectory to ensure alignment between seismic inputs and the supervision target. The Branch 1 input $X1 \in RH \times W \times 3$ contains Vp, Vs, and ρ, while the Branch 2 input $X2 \in RH \times W \times 3$ consists of the three angle stacks.

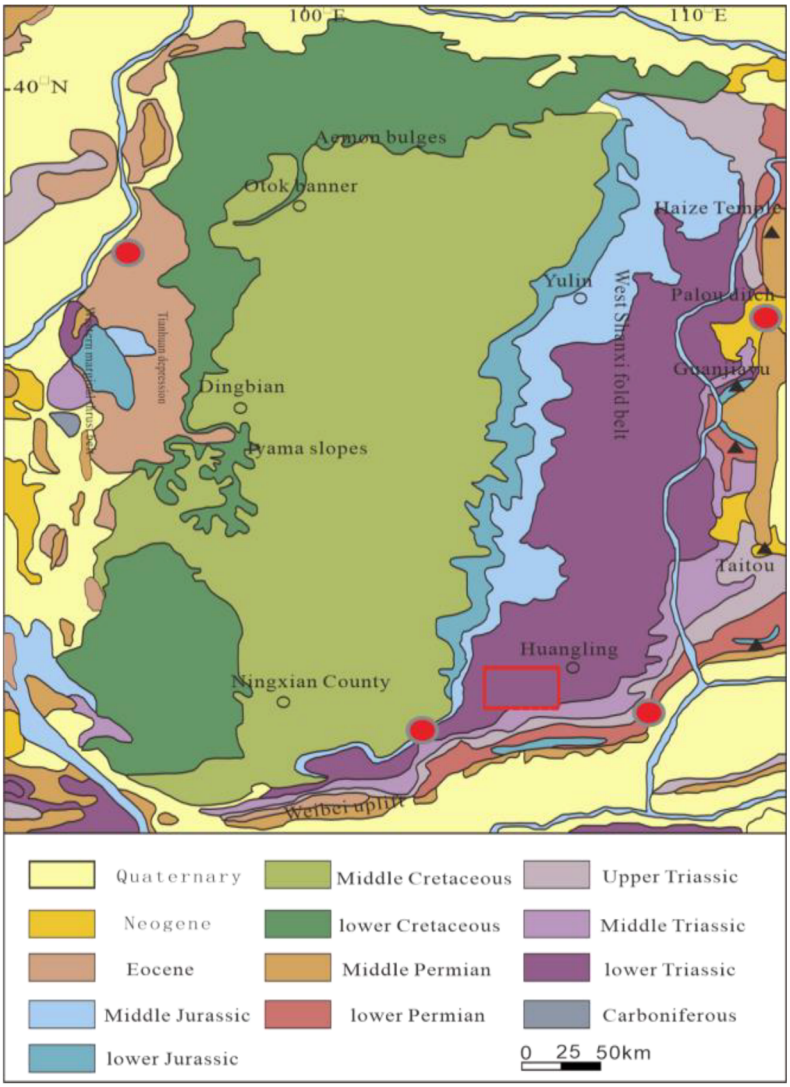

Figure 4. Location map of the study area in the northern Yishan Slope of the Ordos Basin. The red rectangle marks the three-dimensional seismic survey boundary, and red circles indicate the distribution of wells used for inversion and two-branch convolutional neural network training.

The supervision label for each patch is the log-derived ν, as shown in **Equation 22**,

$$\nu^{\log} = \frac{1}{2}\left(1 - 2\left(\frac{\nu^{\log}}{\nu_p^{\log}}\right)^2\right) \tag{22}$$

which is depth–time converted and wavelet-aligned to the seismic domain. All input channels are standardized using the mean and variance of the training set to ensure numerical stability.

The TB-CNN was trained to directly regress ν via a fused latent representation, as shown in **Equation 23**:

$$\hat{\nu} = f_\theta(X_1, X_2) \tag{23}$$

with the network output constrained to a geologically realistic interval of ν ∈ (0.30, 0.38) based on well-log statistics and regional lithology. This constraint was applied during training to stabilize regression rather than enforcing uniform or constant values in the prediction stage.

The objective function combines the MSE loss with a soft physics-guided regularization term, as shown in **Equation 24**,

$$L = L_{\text{MSE}} + \lambda L_{\text{phys}} \tag{24}$$

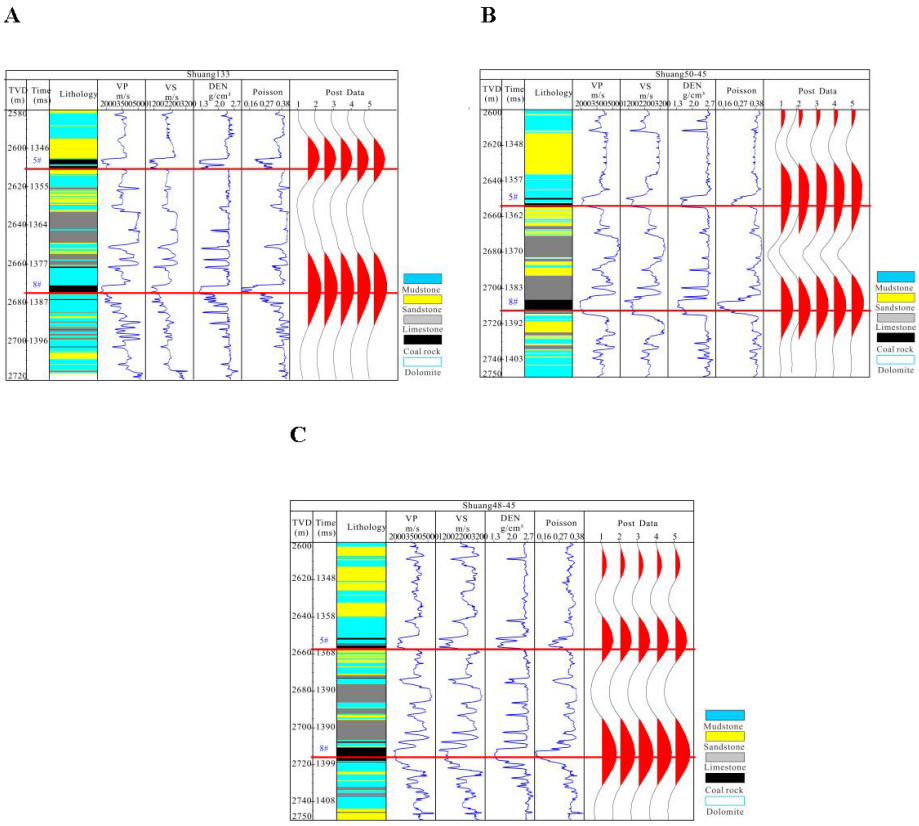

Figure 5. Well-to-seismic ties of three representative boreholes in the study area. (A) Shuang50–45, (B) Shuang133, and (C) Shuang48–45. The 8# coal seam, delineated by the red horizontal line, shows a pronounced impedance contrast with the adjacent mudstone, sandstone, and limestone layers, producing a strong composite reflection with polarity reversal indicative of thin-bed tuning. The consistent seismic signature across the three wells confirms the geological reliability of the 8# seam interpretation and provides essential calibration for pre-stack inversion and two-branch convolutional neural network training. The positions of both the 5# and 8# coal seams are consistently indicated across all figures to ensure unified stratigraphic reference.

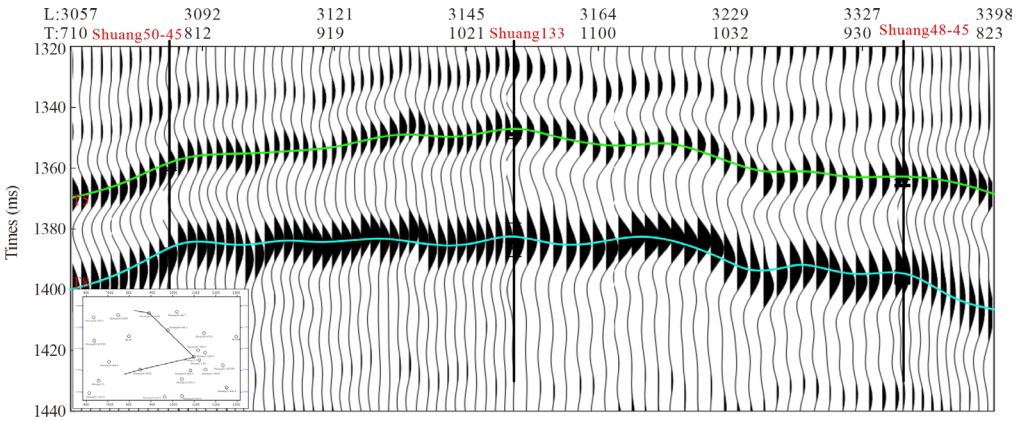

Figure 6. Cross-well seismic section illustrating the lateral continuity of the coal-bearing sequence. Four key composite reflection events were identified: T5 and T8 correspond to the dominant tuning-controlled composite responses associated with the 5# and 8# coal seams. These reflections do not represent discrete roof or floor interfaces, and their positive/negative phases should not be interpreted as top/bottom seam reflections. Instead, they arise from constructive and destructive waveform superposition caused by thin-bed tuning (thickness < λ/8 at the dominant frequency). The figure demonstrates the regional continuity of tuning responses and defines a physically reasonable inversion window for subsequent elastic parameter prediction.

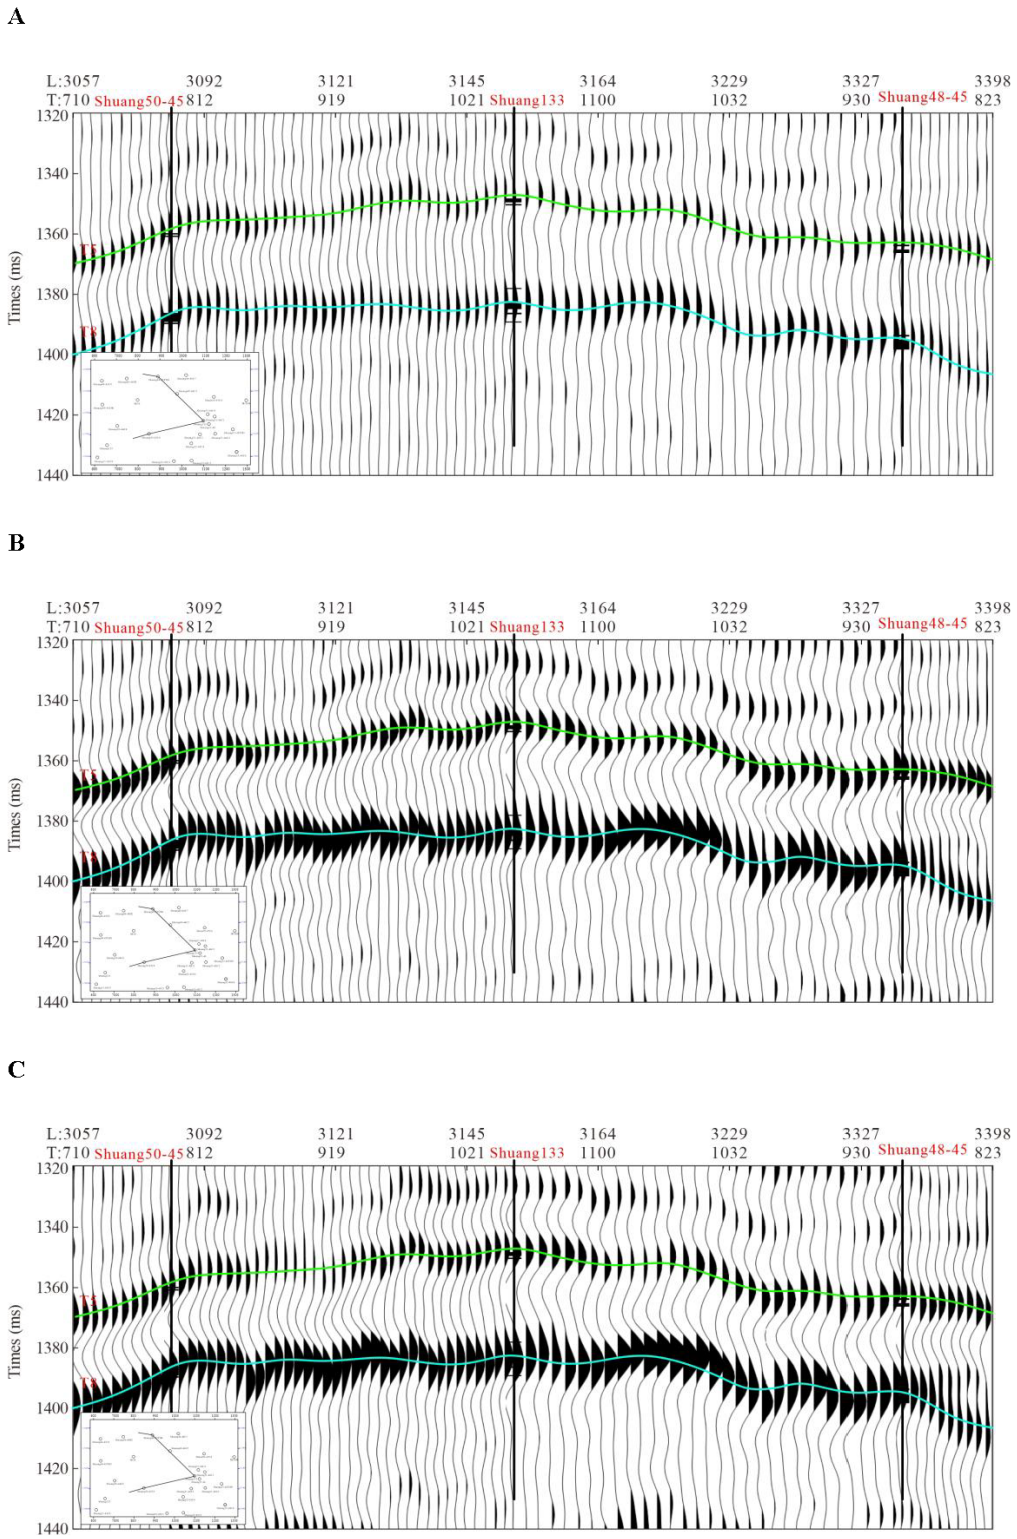

Figure 7. Partial-angle stacked seismic profiles used as Branch 2 inputs. (A) Small-angle stack, (B) medium-angle stack, and (C) large-angle stack.

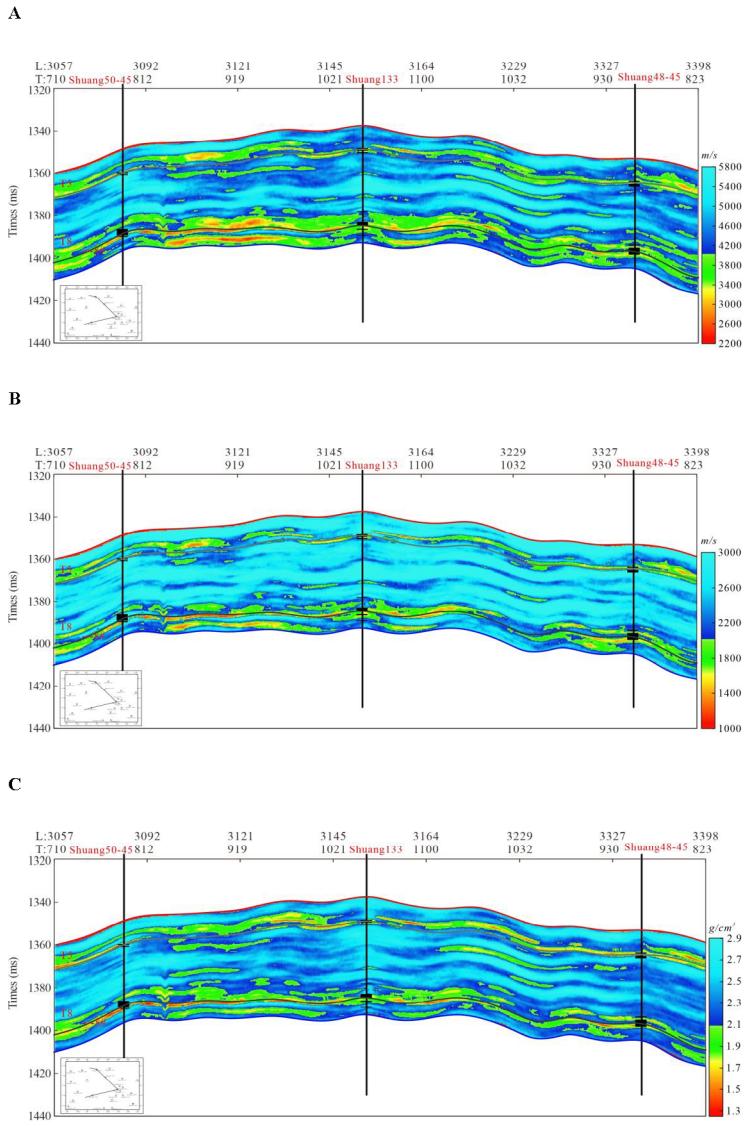

Figure 8. Elastic impedance-inversion products used as Branch 1 inputs. (A) P-wave velocity (Vp), (B) S-wave velocity (Vs), and (C) density (ρ).

which enforces consistency between the predicted ν and empirical rock-physics relationships implied by the input Vp, Vs, and ρ volumes. Rather than constraining velocities explicitly, this regularization penalizes ν predictions that violate lithology-consistent elastic trends inferred from well logs, thereby improving stability and interpretability.

Optimization was performed using the Adam algorithm with an initial learning rate of 1×10^{-4}, a batch size of 32, and a weight decay of 1×10^{-5}. Batch normalization and dropout were applied to mitigate overfitting, and early stopping was employed based on validation loss. Data splitting was performed strictly at the well level to avoid spatial leakage, so that blind wells provided an independent assessment of generalization. Convergence was typically achieved within 80–100 epochs, with stable training and validation loss curves indicating robust learning (Figure 9).

Once trained, the TB-CNN was applied to the full 3D seismic volume in a sliding-window manner to generate a continuous ν cube for reservoir characterization. An example inline ν section predicted by TB-CNN is shown in Figure 10, which illustrates the direct network output used in subsequent analyses.

The TB-CNN result (Figure 10) reveals laterally continuous composite reflection bands associated with the 5# and 8# coal seams. These bands represent

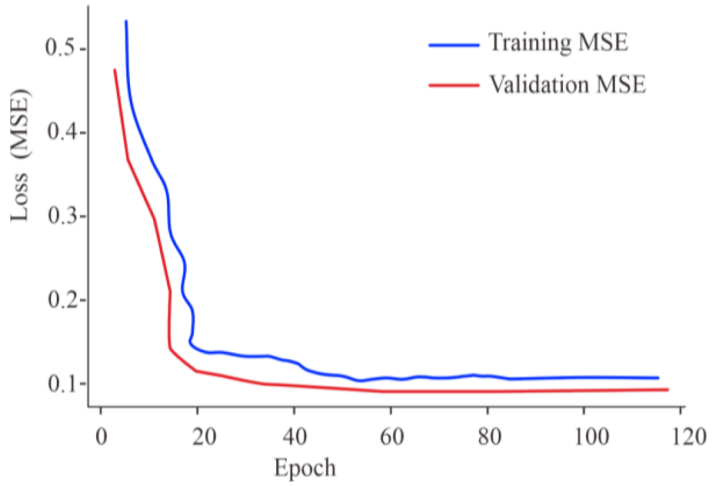

Figure 9. Training and validation loss curves of the two-branch convolutional neural network, showing stable convergence within approximately 100 epochs
Abbreviation: MSE: Mean squared error.

tuning-dominated seismic responses generated by thin coal layers, rather than isolated reflections from individual roof or floor interfaces. Despite the presence of tuning interference, the TB-CNN preserves the relative stratigraphic position and lateral continuity of the coal-bearing intervals, allowing thickness variations and local thinning trends to be recognized. Within the coal seams, ν values were generally elevated in this study area, from approximately 0.30 to 0.38, reflecting the reduced shear rigidity and higher compressibility of coal. Between the 5# and 8# coal seams, ν values showed partial overlap with those of the surrounding sandstone–limestone interbeds, which is expected in interbedded sedimentary sequences. The spatial distribution and relative contrasts of ν were consistent with the log-derived elastic properties, indicating that the TB-CNN prediction captured the first-order lithological variations within the coal-bearing sequence. Based on well-log data from three representative wells in the study area, the ν of the 8# coal seam typically ranged from 0.30 to 0.38. This range aligns with typical

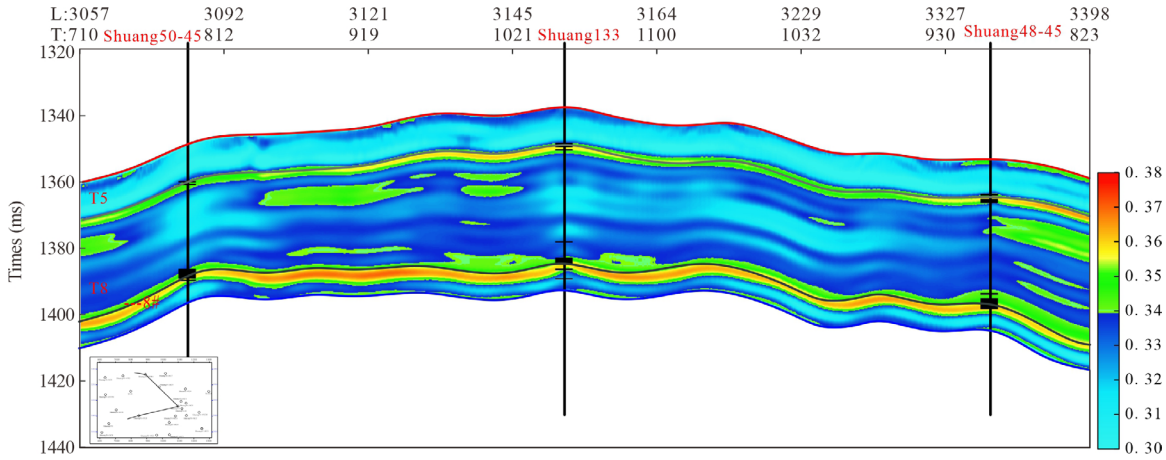

Figure 10. Inline profile of Poisson's ratio predicted by the proposed two-branch convolutional neural network. Coal-bearing intervals are characterized by relatively elevated Poisson's ratio values, forming laterally continuous composite bands controlled by thin-bed tuning effects. These responses do not represent isolated roof or floor reflections, but rather the effective elastic behavior of thin coal seams. In contrast, comparatively lower or transitional Poisson's ratio values were observed in the interbedded sandstone–limestone units between coal seams.

values reported for coal seams under similar geological conditions, characterized by low shear rigidity and higher compressibility.

3.3. Results comparisons

To further demonstrate the effectiveness of the proposed TB-CNN, inline sections of the inverted ν were compared with results obtained from two reference methods: (i) conventional EI inversion followed by Vp–Vs–based ν computation (Figure 11), and (ii) a purely data-driven single-branch CNN (SB-CNN) that directly predicts ν from seismic data (Figure 12). The corresponding TB-CNN result is shown in Figure 10.

The side-by-side comparison highlights pronounced differences in the ability of the three approaches to characterize thin coal seams under strong impedance contrast and tuning-dominated conditions. Relative to the EI- and SB-CNN-based results, the TB-CNN section (Figure 10) provided a stable and geologically coherent reference for seam continuity and thickness variation.

The EI-derived Poisson's ratio section (Figure 11) was dominated by over-smoothing effects, which blurred coal–seam boundaries and suppressed thin-bed signatures as a result of linearized inversion assumptions and band-limited seismic input. In contrast, the SB-CNN result (Figure 12) partially enhanced thin-bed visibility but suffered from local amplitude fluctuations, discontinuous patches, and reduced lateral continuity, reflecting the inherent instability of purely data-driven inversion when physics-based constraints are absent.

To further validate the inversion reliability, near-well comparisons were performed at three representative wells (Shuang50–45, Shuang133, and Shuang48–45), where each profile included a lithology column, log-derived ν, and inversion results from EI inversion, SB-CNN, and the proposed TB-CNN (Figure 13).

The TB-CNN predictions (red curves) showed the highest agreement with the log-derived ν. They accurately delineated the top and bottom of the coal seams, preserved thin-bed variations, and reproduced both the magnitude and trend of ν across lithological transitions. In contrast, EI inversion yielded overly smooth curves with weak sensitivity to interbedded coal layers and limited correspondence to lithological boundaries. The SB-CNN improved local resolution but still exhibited unstable amplitudes and diminished contrast, resulting in poorer alignment with both the log-derived ν and the lithology column.

A comparative analysis of the horizon slices (Figures 14–16) extracted along the 8# coal seam further highlights the advantages of the proposed TB-CNN in characterizing lateral heterogeneity and continuity of coal-bearing reservoirs. The TB-CNN prediction (Figure 14) exhibited a laterally coherent ν distribution with clearly resolved localized anomalies, effectively capturing thin-seam variability and subtle lateral changes that are consistent with the known geological framework and gas-producing well distribution. In contrast, the horizon slice derived from conventional EI inversion (Figure 15) appeared overly smoothed, with weakened lateral contrasts and blurred

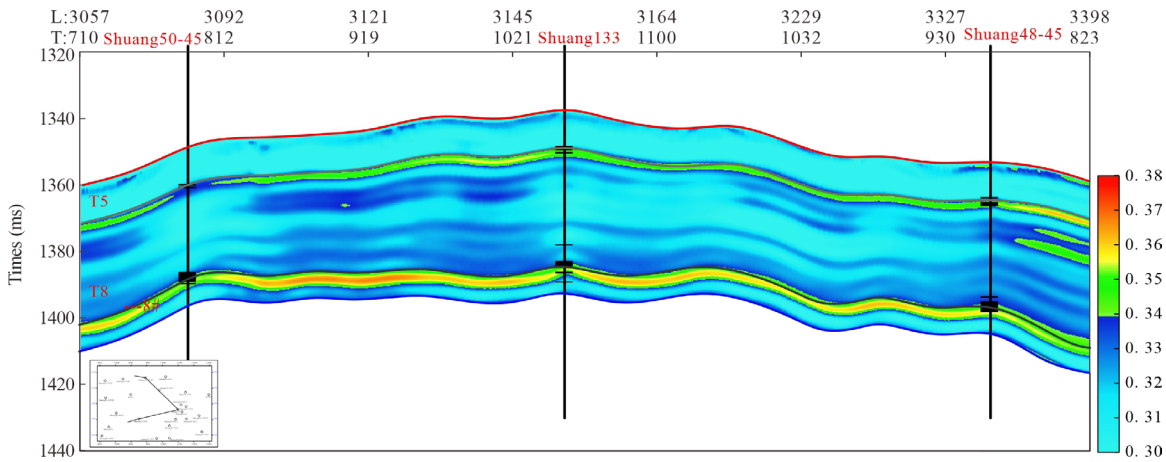

Figure 11. Inline profile of Poisson's ratio derived from conventional elastic impedance inversion. Although the overall stratigraphic trend was preserved, the resulting Poisson's ratio section exhibited pronounced spatial smoothing. Thin coal seams were not sharply delineated, and tuning-dominated responses were significantly attenuated, reflecting the inherent limitations of linearized elastic impedance inversion when applied to thin-bed coal-bearing sequences.

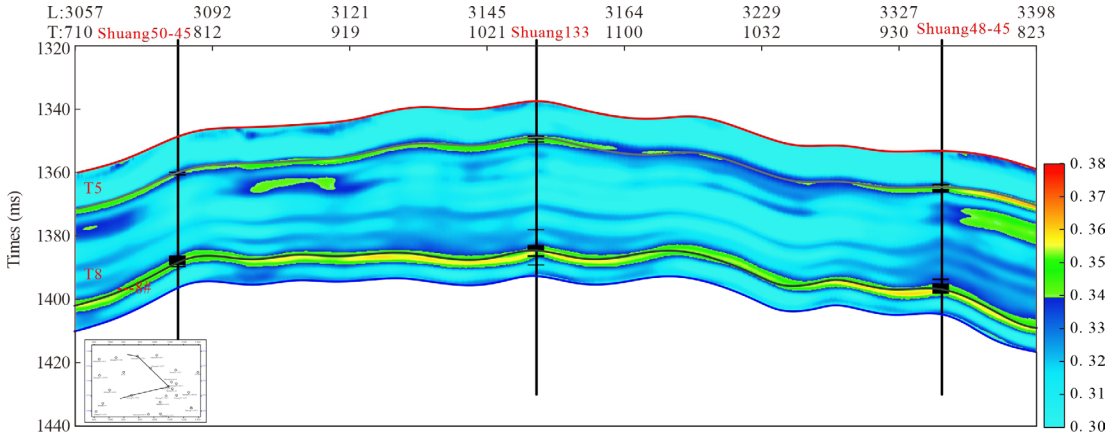

Figure 12. Inline profile of Poisson's ratio derived from single-branch convolutional neural network (SB-CNN) inversion. Although the overall Poisson's ratio range remained physically reasonable, the section showed localized lateral fluctuations and discontinuous features. The lack of physics-guided constraints led to reduced boundary coherence and limited the robustness of SB-CNN in resolving thin coal seams.

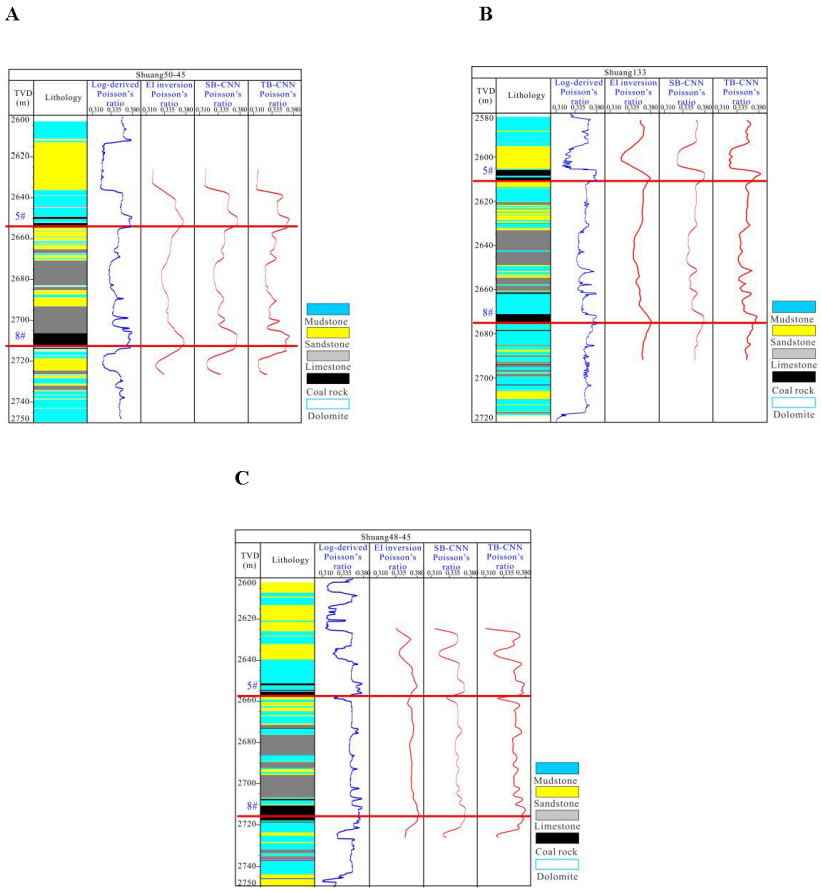

Figure 13. Comparison of Poisson's ratio profiles along three representative wells. (A) Shuang50–45, (B) Shuang133, and (C) Shuang48–45. Each well profile includes a lithology column, a log-derived Poisson's ratio computed from acoustic and density logs, and inversion results from elastic impedance (EI) inversion, single-branch convolutional neural network (SB-CNN), and the proposed two-branch (TB) CNN. The TB-CNN predictions (red curves) showed the highest consistency with the log-derived Poisson's ratio, particularly in reproducing the low Poisson's ratio coal seams and the moderate Poisson's ratio interbedded carbonate–clastic layers, and in delineating coal–seam boundaries and thin-bed variations, whereas the EI and SB-CNN results appeared smoother and less sensitive to thin interbeds.

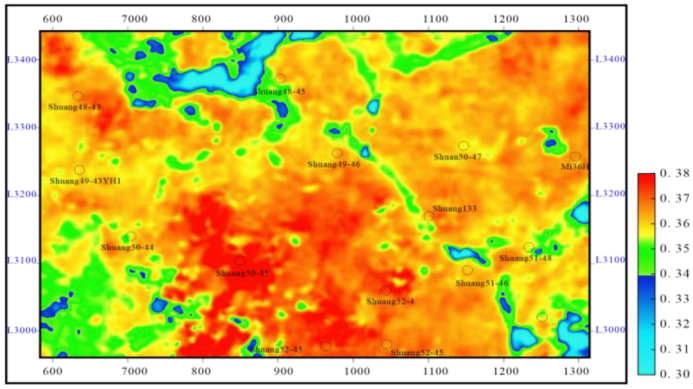

Figure 14. Horizon slice of Poisson's ratio predicted by a two-branch convolutional neural network along the 8# coal seam. Thin-seam variations and localized heterogeneity are clearly resolved.

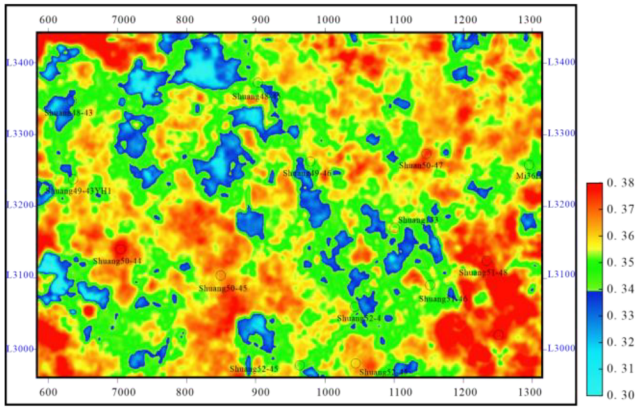

Figure 15. Horizon slice of Poisson's ratio derived from conventional EI inversion. Boundaries appear smoothed, and lateral variations are poorly captured

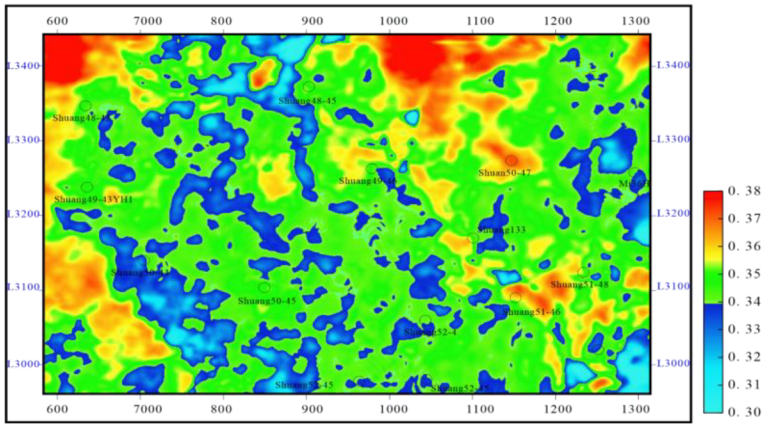

Figure 16. Horizon slice of Poisson's ratio derived from single-branch convolutional neural network inversion. Resolution is moderately improved compared with elastic impedance inversion, but lateral continuity and geological consistency remain limited.

seam boundaries, reflecting the inherent limitations of linearized inversion under thin-bed tuning and strong impedance interference. The SB-CNN result (Figure 16) showed moderate improvement in spatial resolution compared with EI inversion; however, discontinuous anomaly patterns and reduced geological coherence were still observed due to the absence of physics-guided constraints. Overall, the horizon-slice comparison demonstrated that the TB-CNN provided the most geologically consistent and spatially stable representation of ν along the 8# coal seam, achieving a superior balance between lateral resolution and continuity compared with both EI inversion and SB-CNN.

Three systematic advantages of the proposed TB-CNN were consistently observed from the comparative analyses: (i) boundary resolution: the upper and lower interfaces of coal seams were sharply delineated, significantly reducing vertical uncertainty induced by tuning effects; (ii) thickness sensitivity: gradual lateral thickening and thinning trends were preserved rather than flattened, as commonly observed in conventional EI inversion; (iii) Noise robustness: random seismic noise was effectively suppressed without excessive smoothing, while spurious local fluctuations were markedly reduced compared with the purely data-driven SB-CNN.

Overall, the TB-CNN effectively integrates the interpretability and stability of physics-guided inversion with the flexibility and resolution of data-driven learning. It provides a geologically consistent and physically reliable framework for ν inversion and CBM reservoir characterization, achieving consistent improvements in boundary sharpness, heterogeneity detection, and predictive robustness.

4. Discussion

This study presents a physics-constrained TB-CNN for direct ν inversion in thin coal-bearing gas reservoirs, addressing key limitations of conventional EI methods, including wavelet interference, bandwidth loss, and weak sensitivity to thin-bed tuning. By combining EI-derived Vp, Vs, and ρ with multi-angle seismic stacks, the proposed framework preserved large-scale geological consistency while enhancing thin-layer detectability through high-frequency seismic attributes. This dual-branch fusion strategy enabled clearer seam boundary delineation, improved thickness continuity, and more accurate tracking of localized thinning.

The two-branch architecture was central to the model's effectiveness. Branch 1 introduced physics-guided stability and provided rock-physics-consistent background trends, while Branch 2 focused on high-resolution reflectivity patterns sensitive to AVO behavior and thin-bed interference. Their latent fusion mitigated the limitations of linearized Zoeppritz approximations and reduced error propagation typically associated with sequential EI inversion. Direct inversion of ν avoided error propagation associated with sequential estimation of Vp and Vs, thereby improving prediction stability, while physical interpretability was maintained through rock-physics-guided regularization.

The robustness of the TB-CNN arises from both its data formulation and architectural design: (i) partial-angle stacking inherently suppressed random noise while retaining AVO-sensitive reflectivity, and (ii) the fusion of high-SNR, physics-consistent EI inputs with data-driven seismic features mitigated strong-reflection shielding at coal–rock interfaces. As a result, the model produced stable, high-resolution predictions even under low SNR conditions without excessive smoothing.

Unlike fully data-driven end-to-end approaches, the proposed TB-CNN leverages EI inversion as a physically meaningful foundation, enhancing transferability to geologically similar CBM blocks. By learning nonlinear residuals between EI-derived parameters and log-derived ν, the model corrects for angle inconsistencies, tuning effects, and local SNR deterioration. This EI-based feature space reduces dependency on the number of training wells and naturally supports transfer learning. In practice, the EI branch and early convolutional layers may be retained or frozen during domain adaptation, while only the fusion and regression layers are fine-tuned, making the approach practical for few-well scenarios. This feature is particularly valuable for CBM development in areas with sparse well control.

While the results are promising, some limitations remain. The model assumes relatively consistent wavelet characteristics across the study area, which may not hold in strongly heterogeneous or anisotropic environments. Additionally, the number of training wells remains limited, potentially affecting generalization in more complex lithological settings. Future work should explore anisotropy-aware rock-physics constraints,multi-scale feature fusion, and domain-adaptive transfer learning strategies.

5. Conclusion

This study introduced a practical and interpretable TB-CNN framework for ν inversion in thin coal-bearing gas reservoirs. The key findings are as follows:

(i) The proposed dual-branch architecture effectively fused EI-derived elastic attributes with multi-angle

seismic stacks, overcoming resolution loss, wavelet interference, and strong-reflection shielding typical of EI inversion in thin-bed and low SNR conditions.

(ii) The TB-CNN produced higher-resolution ν volumes with sharper seam boundaries, improved thin-layer detectability, and stronger agreement with well-log observations compared with conventional EI and single-branch CNN approaches.

(iii) Sensitivity and uncertainty analyses confirmed that the model is robust with respect to noise levels, tuning effects, and seam thickness, and maintains geological interpretability through rock-physics–guided regularization.

(iv) The hybrid EI–CNN design enhanced transferability, offering a scalable and adaptable framework suitable for few-well scenarios and lateral extension across geologically similar CBM blocks.

Overall, the TB-CNN provides a reliable tool for thin-bed elastic parameter inversion and CBM reservoir evaluation. Its ability to balance physics-guided stability with data-driven resolution makes it particularly promising for broader applications in unconventional reservoir characterization, including tight sandstone, shale, and carbonate formations.

Acknowledgments

The authors gratefully acknowledge this support and thank all project collaborators for their assistance with data processing and interpretation.

Funding

This study was supported by the commissioned project "Development of Seismic Prediction Methods for Coal and Rock Gas Reservoirs Based on Wave Theory" (Contract No. KFKT2024-29), funded by the Exploration and Development Research Institute of Changqing Oilfield Branch, CNPC, and Chengdu University of Technology.

Conflict of interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Author contributions

Conceptualization: Fei Li, Mengbo Zhang, Qiang Liang
Formal analysis: Yongheng Zhang, Mengbo Zhang, Qingzhou Zhang
Investigation: Fei Li, Xiaojie Cui, Na Ni
Methodology: Mengbo Zhang, Qingzhou Zhang
Visualization: Xiaojie Cui ,Yongheng Zhang
Writing–original draft: Yongheng Zhang, Qiang Liang, Na Ni
Writing–review & editing: Fei Li, Mengbo Zhang, Qiang Liang

Availability of data

Not applicable.

References

1. Shuey RT. A simplification of the Zoeppritz equations. *Geophysics*. 1985;50(4):609-614.

   doi: 10.1190/1.1441936

2. Widess MB. How thin is a thin bed? *Geophysics*. 1973;38(6):1176-1180.

   doi: 10.1190/1.1440403

3. Ristow D, Rühl T. Fourier finite-difference migration. *Geophysics*. 1994;59(12):1882-1893.

   doi: 10.1190/1.1443575

4. Buland A, Omre H. Bayesian linearized AVO inversion. *Geophysics*. 2003;68(1):185-198.

   doi: 10.1190/1.1543206

5. Jardani A, Revil A, Slob E, Söllner W. Stochastic joint inversion of 2D seismic and seismoelectric signals in linear poroelastic materials: A numerical investigation. *Geophysics*. 2010;75(1):N19-N31.

   doi: 10.1190/1.3279833

6. Hampson DP, Schuelke JS, Quirein JA. Use of multiattribute transforms to predict log properties from seismic data. *Geophysics*. 2001;66(1):220-236.

   doi: 10.1190/1.1444899

7. Ge Q, Cao H, Yang Z, *et al*. High-resolution seismic impedance inversion integrating the closed-loop convolutional neural network and geostatistics: an application to the thin interbedded reservoir. *J Geophys Eng*. 2022;19(3):550-561.

   doi: 10.1093/jge/gxac035

8. Zhang B, Pu Y, Dai R, Cao D. Seismic poststack impedance inversion using geophysics-informed deep-learning neural network. *Interpretation*. 2025;13(2):T219-T232.

   doi: 10.1190/int-2024-0041.1

9. Chen Y, Fomel S. Random noise attenuation using local signal-and-noise orthogonalization. *Geophysics*. 2015;80(6):WD1-WD9.

   doi: 10.1190/geo2014-0227.1

10. Xue YJ, Zhang H, Zhang JQ, Wang XJ, Cao JX, Liu ZG. Local quantum filtering and denoising of seismic data. *Geophysics*. 2025;90(5):V455-V472.

doi: 10.1190/geo2024-0723.1

11. Wu X, Yan S, Bi Z, Zhang S, Si H. Deep learning for multidimensional seismic impedance inversion. *Geophysics*. 2021;86(5):R735-R745.

doi: 10.1190/geo2020-0564.1

12. Fang W, Fu L, Zhang M, Li Z. Seismic data interpolation based on U-net with texture loss. *Geophysics*. 2021;86(1):V41-V54.

doi: 10.1190/geo2019-0615.1

13. Zhu G, Chen X, Li J, Guo K. Data-Driven Seismic Impedance Inversion Based on Multi-Scale Strategy. *Remote Sens*. 2022;14(23):6056.

doi: 10.3390/rs14236056

14. Jiang X, Cao J, You C, Wang X, Du Z. Deep Neural Networks for Direct Hydrocarbon Detection in Prestack Seismic Data Based on AVO Analysis. *IEEE Trans Geosci Remote Sens*. 2023;61:1-9.

doi: 10.1109/tgrs.2023.3333910

15. Yang N, Li G, Zhao P, Zhang J, Zhao D. Porosity prediction from pre-stack seismic data via a data-driven approach. *J Appl Geophys*. 2023;211:104947.

doi: 10.1016/j.jappgeo.2023.104947

16. Rasht-Behesht M, Huber C, Shukla K, Karniadakis GE. Physics-Informed Neural Networks (PINNs) for Wave Propagation and Full Waveform Inversions. *J Geophys Res Solid Earth*. 2022;127(5):e2021JB023120.

doi: 10.1029/2021jb023120

17. Sun J, Innanen K, Zhang T, Trad D. Implicit Seismic Full Waveform Inversion With Deep Neural Representation. *J Geophys Res Solid Earth*. 2023;128(3):e2022JB025964.

doi: 10.1029/2022jb025964

18. Jiang X, Cao J, Yang J, Liu J, Zhou P. AVO Analysis Combined With Teager-Kaiser Energy Methods for Hydrocarbon Detection. *IEEE Geosci Remote Sens Lett*. 2022;19:1.

doi: 10.1109/LGRS.2020.3029170

19. Xudong J, Junxing C, Jiangtao H, Xuanchen X, Jie L. Pre-stack gather optimization technology based on an improved bidimensional empirical mode decomposition method. *J Appl Geophys*. 2020;177:104026.

doi: 10.1016/j.jappgeo.2020.104026

20. Jiang X, Cao J, Zu S, Xu H, Wang J. Detection of hidden reservoirs under strong shielding based on bi-dimensional empirical mode decomposition and the Teager-Kaiser operator. *Geophys Prospect*. 2021;69(5):1086-1101.

doi: 10.1111/1365-2478.13073

21. Li P, Wang Y, Liu J, Liang M. Evaluation of Carbon Emission Efficiency and Analysis of Influencing Factors of Chinese Oil and Gas Enterprises. *Energy Sci Eng*. 2025;13(3):1156-1170.

doi: 10.1002/ese3.2055

22. Castagna JP, Sun S, Siegfried RW. Instantaneous spectral analysis: Detection of low-frequency shadows associated with hydrocarbons. *Lead Edge*. 2003;22(2):120-127.

doi: 10.1190/1.1559038