Advanced Energy Systems and Intelligent Geoscience (AESIG)
https://www.mospbs.com/journal/aesig

• *Article* •

Prediction of Mineral Content and Petrophysical Parameters in Lacustrine Fine-Grained Mixed Sedimentary Rocks Based on a PhysicsInformed Hybrid Deep Learning Framework

Yiming Zhu, Yukun Liu*, Shuya Chen, Yulin Du , Xiang Cheng, Xiaolong Wang

Key Laboratory of Exploration Technologies for Oil and Gas Resources (Yangtze University), Ministry of Education, Wuhan 430100, China.

***Corresponding Author:** Yukun Liu **Email:** yukunliu@yangtzeu.edu.cn
Received: 16 May 2026 **Accepted:** 12 June 2026

Abstract: To address the challenges of severe logging response overlap and quantitative mineral prediction caused by the strong heterogeneity of lacustrine carbonate reservoirs, this study proposes a hybrid framework integrating a physics-informed Gradient Boosting Decision Tree (LightGBM) and a Temporal Convolutional Network (PINN-TCN). The framework employs a sliding-window TCN to capture the deep contextual information of logging sequences, and introduces a truncation strategy along with a physics-constrained loss function in the prediction phase to force the network outputs to comply with the laws of mineral volume closure and density conservation. Coupled with a cosine annealing dynamic weight scheduling strategy, the model achieves a smooth transition from physics-constrained to data-driven optimization, establishing a closed-loop optimization system. Validation using 14,104 data records from the Qianjiang Depression in the Jianghan Oilfield demonstrates that the framework achieves excellent prediction accuracy across three major mineral components, yielding a closure root-mean-square error (RMSE) of only 0.0124, with 99.7% of the predicted values falling within a 5% deviation. Furthermore, by adaptively learning mineral density parameters, the RMSE of solid density prediction is reduced by 73.3%. SHAP (Shapley Additive exPlanations) attribution analysis further confirms that the physical constraints induce a synergistic adjustment of feature contributions within the decision space. This study provides a novel approach for the quantitative evaluation of minerals in complex lithofacies reservoirs, balancing both physical consistency and interpretability.
Keywords: Lacustrine carbonate;Mineral content prediction;Physics-informed constraints; Temporal Convolutional Network (TCN);Well logging interpretation;Model interpretability

1 Introduction

Lacustrine carbonate reservoirs have attracted increasing attention due to their massive resource potential. Represented by the Qianjiang Depression in the Jianghan Oilfield, lacustrine carbonates develop a wide variety of complex

DOI: https://doi.org/10.58244/aesig.263907

lithofacies types characterized by frequently alternating vertical mineral compositions and intense reservoir heterogeneity. Compared with marine carbonates, lacustrine carbonates are more significantly influenced by high-frequency fluctuations in the depositional environment. Their mineral assemblages exhibit rapid vertical variations at centimeter-to-meter scales, resulting in highly overlapping well logging responses and severe non-uniqueness. In addition, factors such as poor borehole conditions and instrument failures during actual logging operations frequently lead to missing or distorted logging curves(Kang et al., 2026),further exacerbating the difficulty of quantitative mineral content prediction. Accurate prediction of reservoir mineral compositions is not only the foundation for logging interpretation and reservoir evaluation, but also a critical prerequisite for petrophysical modeling, diagenesis research, and development plan formulation(Pothana and Ling, 2025; Guo et al., 2026).

Traditional mineral content prediction methods primarily rely on empirical formulas and multivariate statistical analysis. Traditional empirical equations and linear inversion methods struggle to characterize the complex non-linear relationships of coexisting multiple minerals(Guo et al., 2026) and heavily depend on subjective parameter assignments (Pothana and Ling, 2025). Meanwhile, point-to-point machine learning algorithms such as support vector machines and random forests (Pang et al., 2026), as well as deep learning methods (Haritha et al., 2025; Hussain et al., 2025), typically ignore the vertical contextual dependencies of well logging sequences along the depth direction(Sun et al., 2026). Although deep learning methods such as convolutional neural networks (CNNs) and long short-term memory (LSTM) networks have demonstrated advantages in feature extraction, the fixed receptive field of conventional CNNs limits their capacity to represent multi-scale sequence information, while standard LSTMs still face bottlenecks regarding gradient vanishing and computational efficiency when processing long sequences (Wu et al., 2024; Park et al., 2026). More critically, the vast majority of existing deep learning models treat mineral content prediction as a standard multi-output regression problem, directly applying the Softmax function in the output layer to normalize the concentrations of individual mineral components. However, mineral content is a typical form of compositional data subjected to closure constraints, meaning that the sum of all components must strictly equal 1. Mapping the outputs to a simplex space via the Softmax transformation introduces a forced negative correlation among components. This induces a severe "gradient vanishing" phenomenon during gradient optimization: when one component approaches 1, the gradients of the remaining components approach zero, hindering effective network parameter updates. This mathematical flaw not only degrades prediction accuracy but also prevents the reasonable integration of the physically vital density constraint equation, $\rho_{solid} = \sum V_i \cdot \rho_i$(where $\rho_{solid}$is the rock solid equivalent density, V_i is the volume fraction of the *i* -th mineral, and ρ_i is the grain density of the *i* -th mineral), into the model framework.

The aforementioned analysis indicates that existing methods suffer from critical deficiencies in three areas: deep sequence contextual modeling, mathematical treatment of compositional data, and the integration of physical constraints. To address these limitations, this paper proposes a deep learning mineral prediction framework based on sliding-window deep contextual modeling and physical constraint fusion. This approach fully captures local contextual features

and depositional sequence information from well logging curves using a sliding window mechanism. Inspired by the concepts of Physics-Informed Neural Networks (PINNs) (Bodong et al., n.d.), petrophysical constraints such as the density conservation equation are embedded into the loss function to ensure the physical consistency of the predictions. Concurrently, a robust approximation of mineral component normalization is achieved through an output truncation strategy combined with a closure constraint loss term. Taking the lacustrine carbonate reservoirs in the Qianjiang Depression of the Jianghan Oilfield as the study object, this paper systematically validates the effectiveness of the proposed framework for mineral content prediction under complex lithofacies conditions.

2 Methodology

To address the aforementioned limitations of traditional methods in deep sequence modeling, compositional data processing, and physical constraint integration, this paper constructs a deep learning mineral prediction framework that integrates sliding-window contextual modeling with physics-informed constraints. Using the preliminary predictions of a baseline tree model as auxiliary inputs, the framework extracts multi-scale contextual features of logging sequences via a Temporal Convolutional Network (TCN). Subsequently, it simultaneously achieves residual correction of mineral content and density physical constraints through dual output heads, ensuring that the prediction results possess both data-driven fitting accuracy and petrophysical consistency. The model takes a 3D tensor $X \in \mathbb{R}^{B\times C\times W}$ as input, where B is the batch size, W is the sliding window length, and C is the number of input feature channels. The feature channels are constructed by concatenating two parts: $C = C_{\text{feat}} + C_{\text{base}}$, where C_{feat} is the number of standardized original logging feature channels (including depth, natural gamma, deep/shallow laterolog resistivity, compensated neutron, density, acoustic transit time,

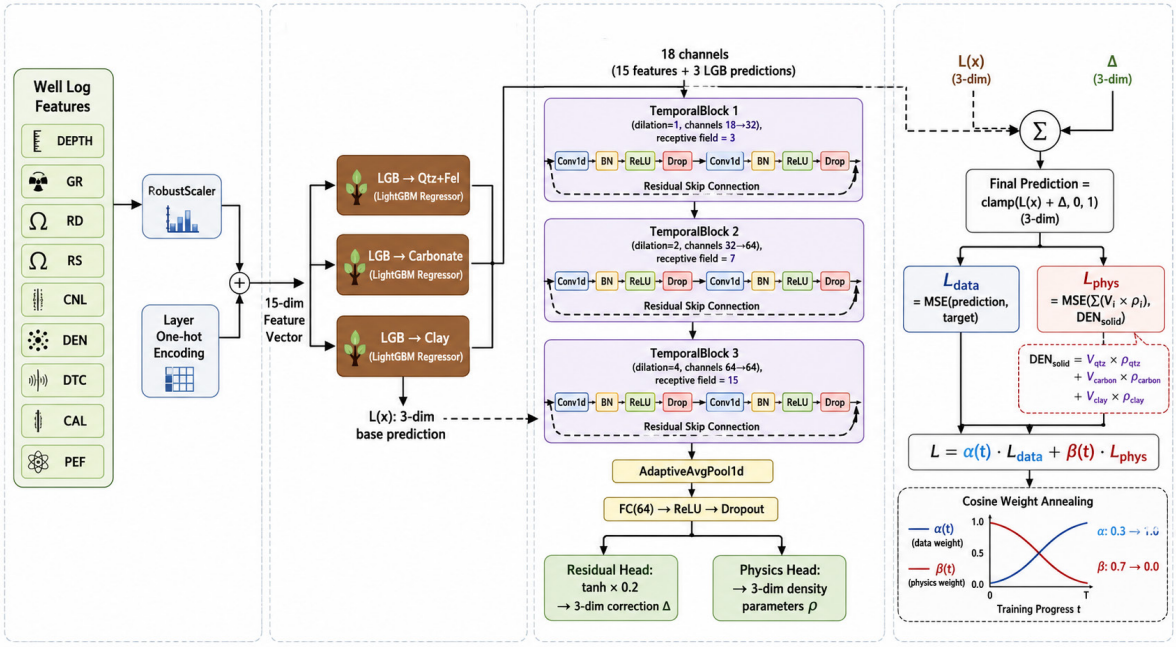

Figure 1. Schematic diagram of the hybrid deep learning mineral prediction framework structure.

caliper, photoelectric absorption cross-section index, and their cross-features), and C_{base} is the number of preliminary prediction channels for each mineral component generated by the baseline tree model. The prediction results of the baseline model are concatenated with the original features as additional input channels, enabling the subsequent deep learning model to focus solely on learning the residual correction for the baseline predictions, thereby circumventing the complex process of learning the complete mineral-logging mapping from scratch. This hybrid architecture significantly reduces the fitting burden on the deep model and accelerates the convergence process.

The prediction results of the baseline model (LightGBM) are concatenated with the original logging features as additional channels. After the sequence features are extracted by the TCN, the residual correction and density physical constraints are accomplished separately through dual output heads. The construction of sliding windows is grouped by well to strictly maintain the continuity of intra-well sequences. At the boundaries of each well, when the center point of the window approaches the beginning or end of the sequence and the window cannot fully cover it, a mirror padding strategy is adopted—the portion of the window exceeding the sequence boundary is padded with a mirror reflection of the internal sequence data, ensuring that a complete window tensor can be generated for every depth sampling point.

This study employs a three-layer 1D-TCN architecture as the feature extractor. By setting the dilation factor $d \in \{1,1,2,2,4,4\}$, the final receptive field of the network is exponentially expanded to 29 sampling points without increasing the number of parameters, thus fully covering the depositional sequence features within the sliding window, and symmetric padding is applied to maintain the consistency of the sequence dimensions.

Within each temporal block, the outputs of the two dilated convolutional layers, following batch normalization and ReLU activation, are added to the block's input via a residual connection:

$$o = ReLU\left(\mathcal{F}(x) + \mathcal{R}(x)\right) \tag{1}$$

where $\mathcal{F}(x)$ is the transformation path of the feature variable, in which the input signal sequentially undergoes feature extraction by two layers of dilated convolutions, numerical scaling by batch normalization, and regularization constraints by a Dropout layer. $\mathcal{R}(x)$ is the shortcut path of a 1×1 convolution, used for dimensional alignment when the number of input channels and output channels is inconsistent. When the number of input and output channels is identical, $\mathcal{R}(x)$. degrades to an identity mapping. The introduction of residual connections effectively mitigates the gradient decay problem in deep convolutional networks, allowing gradients to be directly propagated back to shallow layers via the shortcut path, thereby supporting the stable training of deeper networks.

The output of the TCN backbone network compresses the temporal dimension to 1 via global adaptive average pooling, and is subsequently mapped to the hidden feature space through a fully connected layer. Based on this, the model is equipped with two parallel output heads.

The first output head is the residual correction head, which outputs a *D*-dimensional vector (*D* is the number of mineral components) through a fully connected layer. After the output range is constrained by the tanh activation function, it is multiplied by a preset scaling factor α ($0 < \alpha \ll 1$) for scaling:

$$\ddot{\boldsymbol{u}} = \alpha \cdot tanh\left(W_r s + b_r\right) \tag{2}$$

Where *s* is the output of the fully connected hidden layer, and W_r and b_r are the weights and biases of the residual head. The final mineral

content prediction is the superposition of the baseline model prediction $\hat{v}_{\text{base}}$ and the residual correction amount, constrained to $[0,1]$ via a clamp operation:

$$\hat{v} = clamp\left(v_{base} + \ddot{a}, 0, 1\right) \tag{3}$$

The tanh activation restricts the correction amount to the range $(-\alpha, \alpha)$, which, combined with the small scaling factor α, ensures that the TCN only applies minor fine-tuning to the baseline predictions, avoiding drastic deviations and guaranteeing prediction stability.

The second output head is the physically constrained density head, which outputs the density parameters of the D mineral components in parallel. Because mineral densities have explicit physical range constraints, the raw output of this head is mapped into the preset physical boundaries $\left[\rho_j^{min}, \rho_j^{max}\right]$ via a sigmoid function:

$$\hat{\rho}_j = \rho_j^{min} + \left(\rho_j^{max} - \rho_j^{min}\right) \cdot \sigma\left(w_j\, s + b_j\right) \tag{4}$$

Where $\sigma(\cdot)$ is the sigmoid function, and $\rho_j^{\min}$ and $\rho_j^{\max}$ are the physical density boundaries of the j -th mineral. This design enables the density parameters to be adaptively learned and adjusted during the training process while consistently satisfying physical feasibility constraints. The density balance equation constitutes the core of the physics-informed loss term:

$$\mathcal{L}_{\text{phys}} = \frac{1}{B} \sum_{i=1}^{B} \left(\sum_{j=1}^{D} \hat{v}_{i,j} ? \hat{\rho}_{i,j} - \rho_{\text{solid},i}^{\text{meas}} \right)^2 \tag{5}$$

Combining the two output heads mentioned above, the total loss function is a weighted combination of the data-driven loss and the physics-constrained loss:

$$\mathcal{L}_{\text{total}} = w_{\text{data}} \cdot \mathcal{L}_{\text{data}} + w_{\text{phys}}\, \mathcal{L}_{\text{phys}} \tag{6}$$

Where $\mathcal{L}_{\text{data}}$ is the weighted mean squared error of the mineral content. The regulation of loss weights operates synergistically through a two-layer mechanism: the first layer is a staged scheduling strategy based on cosine annealing, which smoothly and monotonically increases w_{data} from its initial value to its final value according to a cosine curve based on the training epoch t. Correspondingly, $w_{\text{phys}} = 1 - w_{\text{data}}$ decreases from a high value to a low value, achieving a progressive transition from “physics-constraint dominant in early training” to “data-fitting dominant in later stages.” The second layer involves a learnable uncertainty weight, which adaptively corrects the magnitude difference between the two losses by introducing a trainable log-variance parameter $\log\sigma^2$, avoiding optimization bias caused by the differing numerical scales of $\mathcal{L}_{\text{data}}$ and $\mathcal{L}_{\text{phys}}$.

3 Experiments and Results Analysis

3.1 Dataset construction and feature engineering

Experimental data were collected from five wells (W-1, W-2, W-3, W-4, and W-5) within the lacustrine carbonate strata of the study area, totaling 16,648 depth sampling records. The prediction targets are three mineral component end-members: quartz-feldspar content ($v_{\text{Qtz+Fel}}$), carbonate mineral content (v_{Carbon}), and clay mineral content (v_{Clay}), which satisfy the closure constraint $v_{\text{Qtz+Fel}} + v_{\text{Carbon}} + v_{\text{Clay}} = 1$. The input features consist of nine conventional logging curves, including GR, RD, RS, CNL, DEN, DTC, CAL, and PEF, as well as three physical cross-features (PEF×DEN, CNL×DEN, and GR×DEN) and the one-hot encoding of stratigraphic horizons, totaling 15 dimensions. During the preprocessing stage, well W-5 was entirely excluded due to the complete absence of target mineral contents and was only involved in the auxiliary modeling for imputing missing CNL values. For the remaining four wells, after eliminating samples with missing target values or unfulfilled physical rationality ($\rho_{\text{solid}} \le 0$ or exceeding the $\left[Q_{0.01}, Q_{0.99}\right]$ range), a total of 14,104 valid samples from wells W-1, W-2, W-3, and

W-4 were ultimately retained. After standardizing the data using RobustScaler, it was divided into a training set (12,693 records) and a validation set (1,411 records) at a ratio of 90%:10%.

The correlation analysis in Figure 2 demonstrates that the input features and mineral targets possess clear petrophysical correspondence: DEN and PEF show strong positive correlations with carbonate rocks (*r* values of 0.63 and 0.45, respectively), while GR is strongly positively correlated with clay (0.58).Crucially, the correlation between the physical cross-feature PEF×DEN and carbonate minerals jumps to 0.68, significantly higher than that of single curves, proving the effectiveness of cross-features in amplifying the sensitive signals of specific minerals.

3.2 Model training parameters and evaluation metrics

The hybrid model adopts a two-stage training strategy: first, independent LightGBM regressors are trained for each mineral component to obtain baseline predictions; subsequently, the ResidualTCN model is trained to predict the residual correction amounts for the LightGBM outputs. The key training parameters are as follows: batch size of 512, learning rate of 3×10^{-4} (Adam optimizer, weight decay of 10^{-4}), maximum

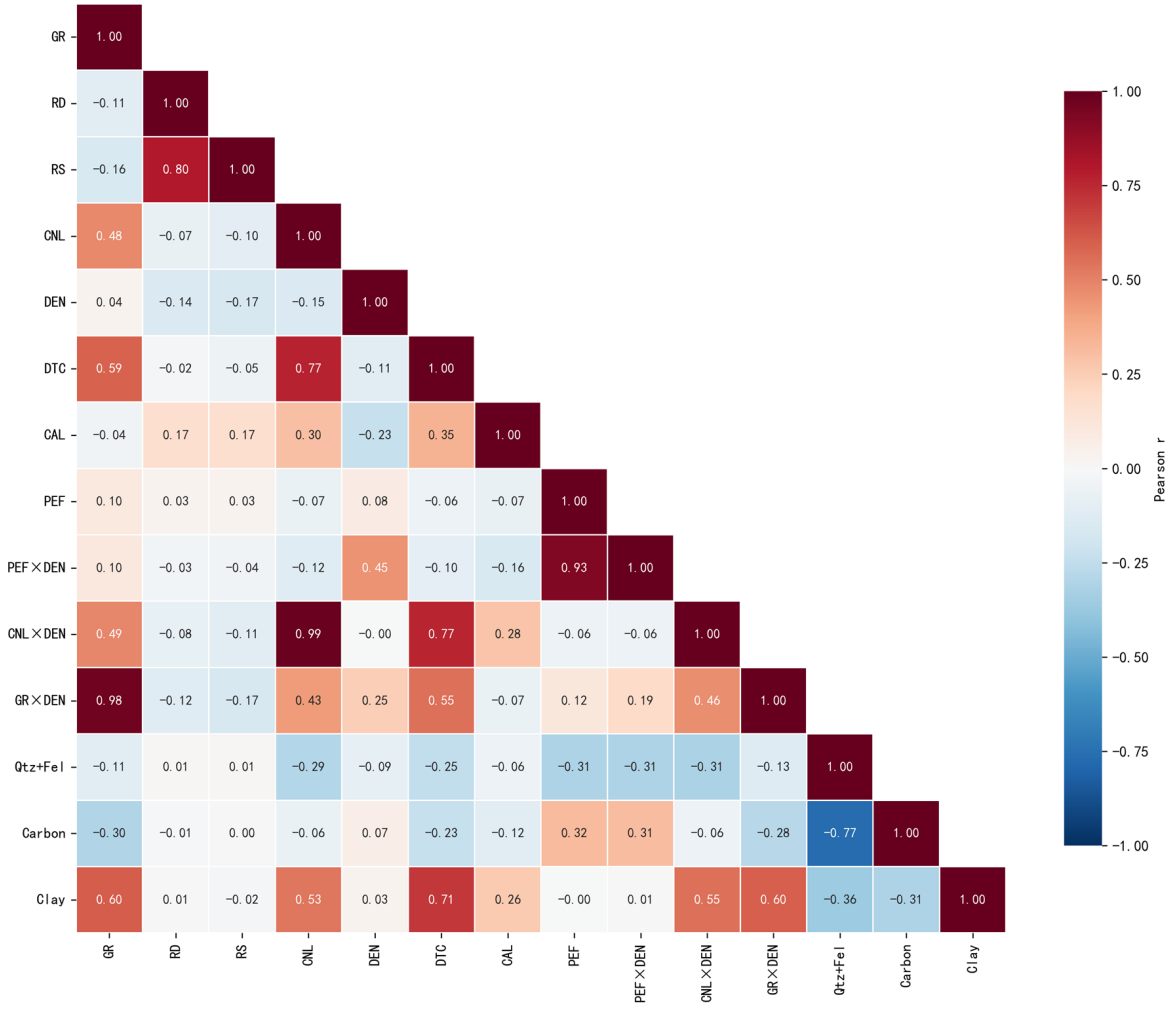

Figure 2. Pearson correlation coefficient heatmap between input features and mineral targets.

training epochs of 1000 (early stopping patience of 150), sliding window size of 21, and learning rate scheduling using ReduceLROnPlateau (decay factor of 0.5, patience of 30). The model converged stably after 1,000 training epochs, with the optimal validation loss appearing at the 893rd epoch. The contrast baselines consist of two point-wise gradient boosting models: LightGBM (n_estimators=500, lr=0.05, max_depth=6) and CatBoost (iterations=500, lr=0.05, depth=6), both utilizing all 15-dimensional features without introducing sliding-window contexts. Utilizing point-wise inputs is the standard practice for tree models, whereas sequence contextual modeling is an inherent advantage of deep learning architectures; the hybrid model did not surpass LightGBM in terms of R^2, indicating that its core improvement stems from physical constraints rather than contextual information volume.

Figure 3 illustrates the evolutionary trajectories of six key metrics during the training process of the hybrid model. As shown in subplot (a), both the data MSE and the physics MSE exhibited a healthy trend of continuous decline as training progressed. Subplot (b) further demonstrates that the adaptive total loss for both the training and validation sets maintained stable convergence without divergence, which strongly confirms that no overfitting occurred throughout the entire training cycle. Regarding the execution of the dynamic optimization strategy, the previously detailed weight scheduling mechanism is intuitively reflected in subplot (c), successfully guiding the data weight through a smooth transition from 0.3 to 1.0. Accompanying this process, subplot (d) reveals the stable convergence states of the three dynamic density parameters within their physical boundaries, while subplots (e) and (f) record the

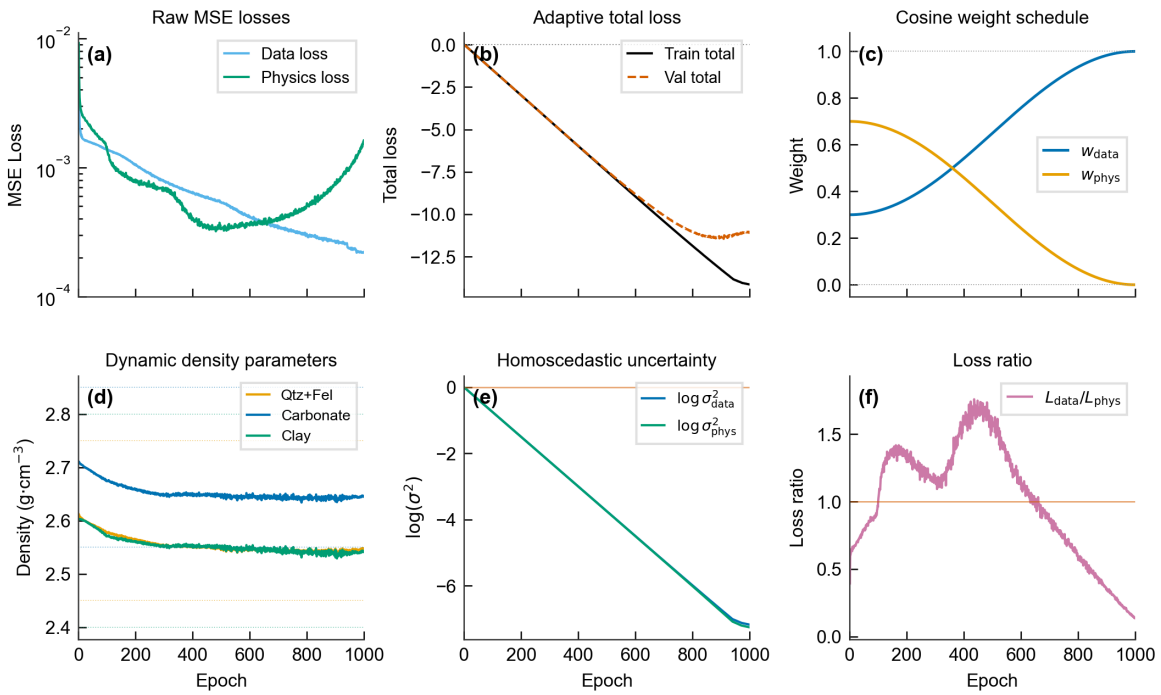

Figure 3. Visualization of the hybrid model training process. (a) Original MSE loss (logarithmic scale); (b) Adaptive total loss; (c) Weight scheduling curve; (d) Evolutionary trajectories of dynamic density parameters; (e) Homoscedastic uncertainty parameter; (f) Data/physics loss ratio.

Table 1. Comparison of prediction performance across different models (R2 and RMSE).

Model	Qtz+Fel R^2	Qtz+Fel RMSE	Carbon R^2	Carbon RMSE	Clay R^2	Clay RMSE
LightGBM	0.8583	0.0452	0.8513	0.0450	0.8822	0.0267
CatBoost	0.8017	0.0534	0.7967	0.0526	0.8505	0.0301
Hybrid (LGB+TCN)	0.8368	0.0485	0.8360	0.0472	0.8713	0.0279

continuous decline of the homoscedastic uncertainty parameter and the synchronous increase in the ratio of data loss to physics loss in tandem with weight scheduling, respectively. This series of detailed training dynamics fully indicates that the cosine annealing weight scheduling and the adaptive uncertainty mechanism produced highly efficient synergistic effects. This synergy not only drove the model to achieve an optimal balance between data fitting accuracy and petrophysical consistency but also enabled the physical prior information to gradually internalize into the deep inductive bias of the model's predictions, rather than merely remaining as a superficial external penalty term.

This study employs three metrics to evaluate model performance: Root Mean Square Error (RMSE), Mean Absolute Error (MAE), and Coefficient of Determination (R^2). RMSE is sensitive to large errors, and a lower value indicates a smaller overall dispersion of the predictions; MAE measures the average deviation magnitude linearly and is more robust to outliers; R^2 measures the model's ability to explain the variance of the target variable, with values closer to 1 indicating a higher degree of fitting.

3.3 Comparative analysis of mineral composition prediction accuracy

Observing the comprehensive performance of each model on the validation set in Table 1 reveals that conventional well logging data contain sufficient effective information for mineral content prediction. Among the purely data-driven baselines, LightGBM achieves the highest statistical accuracy. In contrast, the overall coefficient of determination of the hybrid model falls between the two baseline tree models, which is due to the strong regularization effect introduced by the physics-informed loss term. This

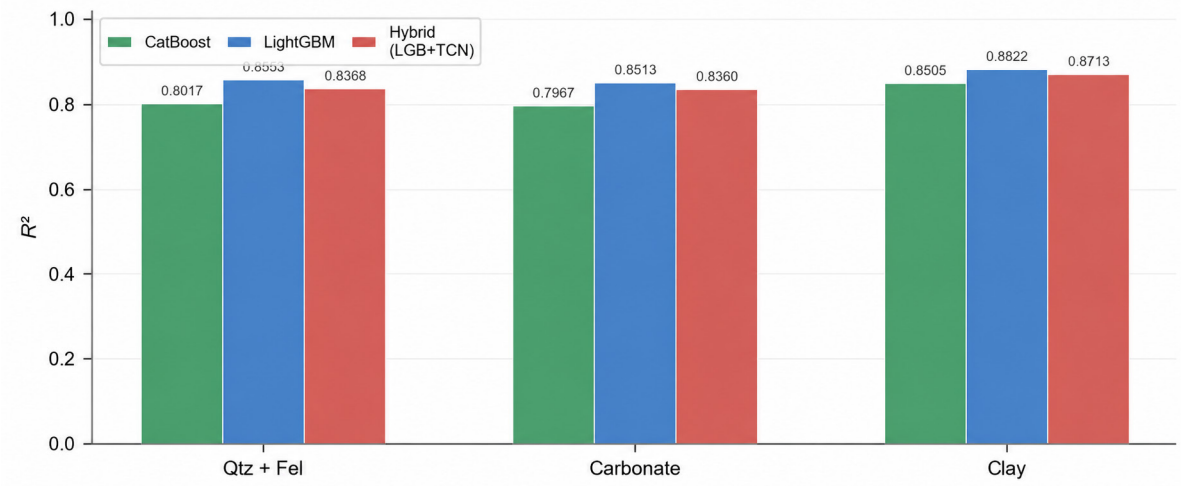

Figure 4. Bar chart comparing R^2 across different models for the three mineral components.

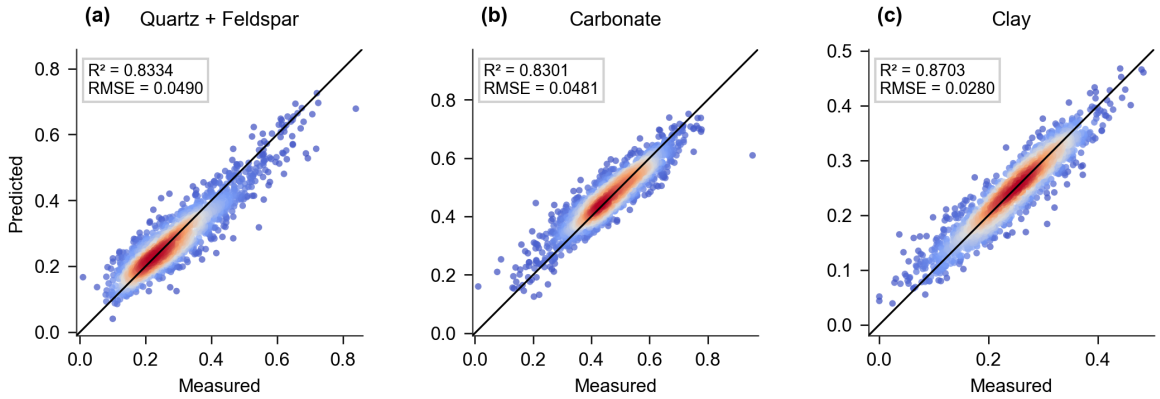

Figure 5. Cross-scatter plot of predicted versus measured values for the hybrid model.

mechanism directs the prediction trajectory toward a petrophysically consistent manifold space at the expense of minor purely statistical fitting metrics. (Figure 4.)

The bar chart in Figure 4 intuitively reflects the accuracy differences across models and components. Among them, the carbonate component is the most difficult to predict; however, the hybrid model achieves a significant improvement over CatBoost for this component. Combined with the cross-scatter plot in Figure 5, it can be seen that although the overall R^2 improvement is limited, the sequence contextual modeling of TCN effectively corrects local systematic deviations. Particularly in the high- and low-value regions of quartz-feldspar and clay, the hybrid model (blue scatter points) clusters more tightly along both sides of the ideal 1:1 reference line than the baseline model (green scatter points), confirming the corrective effect of contextual information on extreme value predictions.

3.4 Evaluation of physical constraints and petrophysical consistency

In quantitative mineral analysis, closure constraints and density consistency are the core metrics for evaluating the geological rationality of a model (Tables 2 and 3). Although the purely data-driven LightGBM achieved the highest R^2, its closure consistency is the weakest (Table 2), which profoundly reveals the trade-off contradiction between purely pursuing statistical fitting and petrophysical consistency in well logging inversion. The hybrid model performs best on the closure RMSE metric, with nearly 100% of the predicted values falling within a 5% deviation (Figure 6a).

Furthermore, because the traditional fixed-density scheme fails to adapt to the complex diagenetic alterations of lacustrine carbonates, the solid density (DEN_solid) prediction R^2 of all baseline models fell into negative values (Table 3). The dynamic density optimization mechanism introduced in this study successfully breaks this deadlock, enabling the hybrid model to

Table 2. Evaluation of closure constraints for each model.

Model	Closure Mean	Closure Std. Dev.	Closure RMSE	Proportion of $\Delta < 5\%$ (%)
LightGBM	1.0008	0.0189	0.0190	98.4
CatBoost	0.9996	0.0140	0.0140	99.6
Hybrid (LGB+TCN)	0.9981	0.0122	0.0124	99.7

Table 3. Evaluation of density consistency (DEN_solid).

Model	RMSE (g/cm³)	R^2	MAE (g/cm³)
LightGBM（Fixed ρ）	0.1458	−4.083	0.1235
CatBoost（Fixed ρ）	0.1396	−3.662	0.1195
Hybrid（Fixed ρ）	0.1295	−3.009	0.1145
Hybrid（Dynamic ρ）	0.0390	0.636	0.0293

significantly reduce the density RMSE (a relative reduction of 73.3%, Figures 6b and 6c). Crucially, the carbonate equivalent density (2.638 g/cm³) obtained through the model's adaptive learning is lower than the theoretical value of pure calcite. This highly aligns with the actual lithofacies characteristics of the lacustrine strata in the study area, which are influenced by detrital inputs and rich in argillaceous impurities, strongly confirming that the physical output head has captured true geological significance.

The above results confirm the dual advantages of the hybrid model in statistical accuracy and physical consistency, and indicate that physical cross-features and stratigraphic horizon encoding played a key role. However, the nonlinear contributions of each input feature to the prediction of specific mineral components and their interactive mechanisms still require further analysis, which will be detailed in Chapter 4 through SHAP analysis in the discussion of model interpretability.

4 Model Interpretability and Feature Contribution Analysis

Global feature attribution based on SHAP (Figures 7 and 8) indicates that the model automatically evolved a feature-specific allocation mechanism: CNL×DEN and PEF×DEN suppressed pore fluid interference through nonlinear combinations, becoming the core for identifying the sandstone matrix and carbonates; meanwhile, DTC and GR×DEN relied on their sensitivity to hydration and high-gamma/low-density properties to influence the discrimination of clay components. This synergistic discrimination pattern perfectly matches known petrophysical response mechanisms.

Notably, Figure 8 shows that the hybrid model did not adopt a uniform global weight strategy,

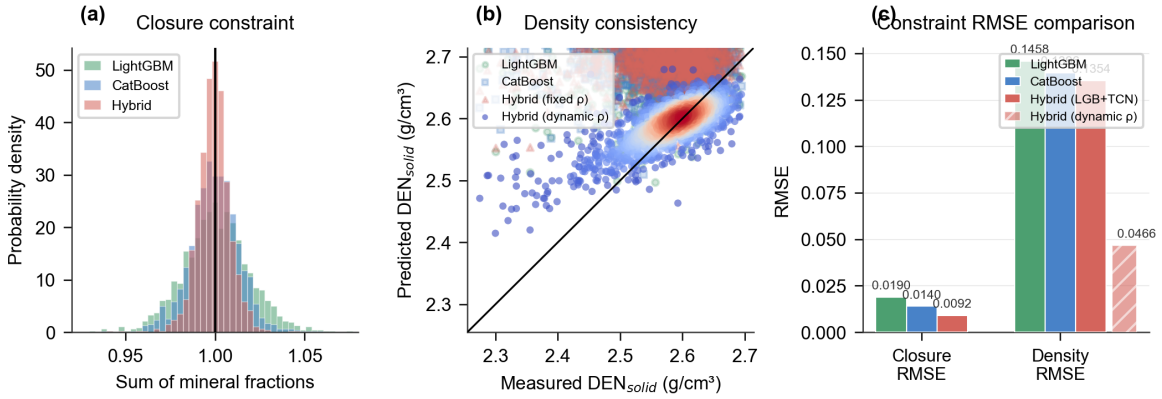

Figure 6. Evaluation of physical constraint effectiveness. (a) Probability density distribution of the sum of mineral contents; (b) Scatter plot of predicted versus measured DEN_solid (purple diamonds: hybrid model + dynamic density); (c) Bar chart comparing closure and density RMSE.

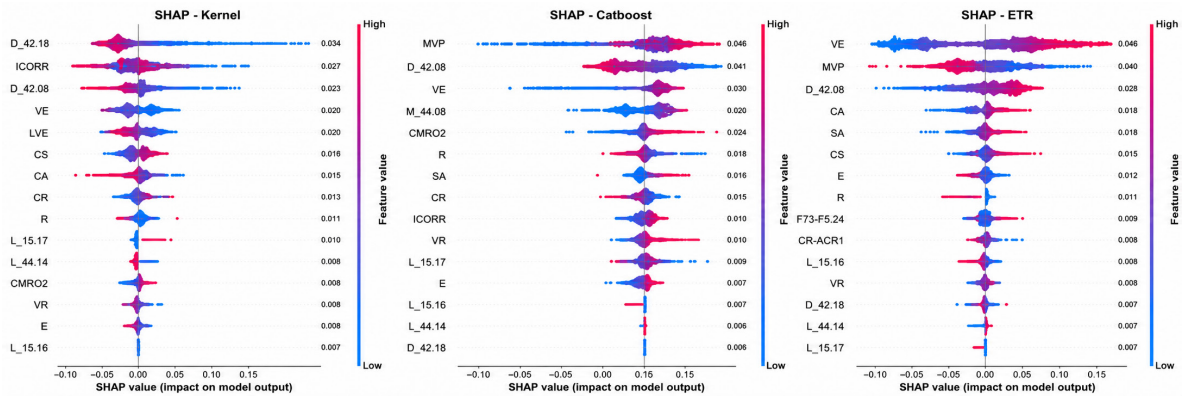

Figure 7. SHAP feature importance analysis of the LightGBM baseline model.

but automatically evolved a feature-specific allocation mechanism: DTC serves as the core anchor for identifying clay, utilizing its sensitivity to compaction and clay hydration to achieve high-precision discrimination; CNL×DEN and PEF×DEN serve as the core discriminative features for quartz+feldspar and carbonate, respectively, fully utilizing the cross-coupled information of mineral physical responses. This differentiated allocation of internal weights is the key intrinsic reason why the model can simultaneously predict multiple components with high precision. The model does not rely on a single logging curve, but has learned a synergistic discrimination pattern among multiple logging parameters. The physical output head forces the model at the architectural level to adhere to $\sum V_i = 1$. This strong constraint is objectively projected into the SHAP attribution space: when a certain feature drives an increase in the predicted value of one component, the model inevitably and synergistically suppresses other components to maintain total balance. This penalty mechanism is accurately captured by SHAP as a negative contribution. This confirms that the baseline model is not merely performing statistical fitting, but is executing strict decision logic within the physically constrained manifold space.

5 Discussion

5.1 Petrophysical implications

The prediction of well logging mineral content is essentially a compositional data analysis problem constrained by closure effects. Through a physics-informed loss and output truncation, the proposed model achieves a robust closure approx-

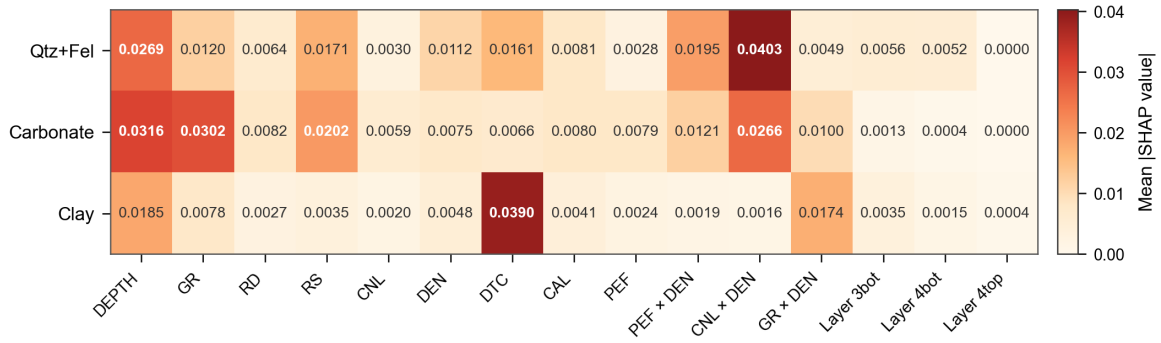

Figure 8. Normalized heatmap of SHAP feature importance for the three mineral components.

imation (with a closure RMSE of only 0.0124), circumventing the fundamental mathematical flaw of treating the three major mineral components as independent regression targets. Furthermore, the dilated causal convolutions of the TCN structurally align with the hierarchical nesting of sedimentary microcycles, equipping the model with the ability to distinguish whether "high density originates from carbonate minerals" or "high density merely reflects physical compaction."

In typical saline lacustrine basins such as the Qianjiang Depression in the Jianghan Basin, the mineral assemblages of lacustrine tight carbonates alternate rapidly in the vertical direction at centimeter-to-meter scales (Pothana and Ling, 2025)particularly in hydrocarbon exploration, CO2 sequestration, and geothermal energy development. Current techniques, such as multimineral petrophysical analysis, offer details into mineralogical distribution. However, it is inherently time-intensive and demands substantial geological expertise for accurate model evaluation. Furthermore, traditional machine learning techniques often struggle to predict mineralogy accurately and sometimes produce estimations that violate fundamental physical principles. To address this, we present a new approach using Physics-Integrated Neural Networks (PINNs. Such deep-seated rocks develop rigid frameworks, and their pore evolution and compaction behaviors are primarily controlled by intense diagenesis. This causes empirical models based on the normal compaction trend (NCT) of clastic rocks to completely fail in these scenarios. The proposed framework discards the reliance on empirical compaction assumptions and directly embeds the law of density conservation into the network as a learnable constraint, opening a novel path for the evaluation of carbonate rocks under complex diagenetic alterations.

5.2 Limitations and future prospects

Despite achieving a favorable balance, this framework still presents certain limitations. The vertical resolution is constrained by the sliding window size (W=21). For extremely thin beds (e.g., laminated argillaceous limestone), the local averaging effect of the TCN may cause signals to be diluted by the surrounding rock. Meanwhile, the current equivalent equation does not adequately account for the impact of secondary spaces, such as dissolution pores and vugs, on the framework density. Future research could introduce leave-one-well-out cross-validation (Lai et al., 2024) to evaluate the extrapolation and generalization capabilities of the model, and utilize the framework density parameters learned by the physical output head (Li et al., 2025) as indirect indicators of diagenetic intensity. Additionally, seismic attributes could be introduced as low-frequency spatial constraints to smoothly extend single-well predictions into three-dimensional space (Sun et al., 2026), ultimately constructing an end-to-end comprehensive reservoir evaluation system (Yang et al., 2016; Bagheri et al., 2024; Cao et al., 2026).

6 Conclusions

Aiming at the quantitative mineral content prediction challenge in strongly heterogeneous lacustrine carbonate reservoirs, this paper proposes a hybrid deep learning framework integrating a Temporal Convolutional Network with physics-informed constraints and a Gradient Boosting Decision Tree baseline. The framework achieves the unification of prediction accuracy and physical consistency through three synergistic mechanisms: a sliding-window TCN captures the local depositional context of logging sequences using dilated causal

convolutions; an output truncation strategy, combined with a physics-constrained loss function, jointly constrains the closure conditions of mineral components and the density balance equation; and a cosine annealing dynamic weight scheduling strategy progressively transitions the training objective from physics-constraint dominant to data-fitting dominant. SHAP interpretability analysis further reveals the internal decision logic of the model: different mineral components are driven by differentiated core features, with acoustic transit time playing a dominant discriminative role for clay content, and the neutron-density cross-feature for quartz-feldspar content, indicating that the model has learned a feature-specific allocation pattern. The introduction of physical constraints causes the multi-component predictions to exhibit a synergistic suppression effect within the attribution space, ensuring the intrinsic unification of interpretability and physical consistency. Validation using four wells from the Qianjiang Depression in the Jianghan Oilfield demonstrates that the closure consistency and density prediction accuracy of the hybrid model are significantly superior to those of the baseline models, and the equivalent density parameters learned by the physical output head highly align with actual lithofacies characteristics. In summary, the proposed framework provides a novel approach for lacustrine carbonate mineral prediction that balances fitting accuracy, physical consistency, and interpretability.

Acknowledgement

The authors would like to thank all individuals and institutions that contributed to this study.

Funding Statement

This work was supported by the Natural Science Foundation of Hubei Province, China (Grant No. 2024AFB271).

Author Contributions

All authors contributed to the study conception and design, data collection, analysis and interpretation of the results, and manuscript preparation, and take responsibility for the integrity of the work.

Availability of Data and Materials

None.

Conflicts of Interest

The authors declare that they have no conflicts of interest to report regarding the present study.

References

[1] Bagheri, H., Mohebian, R., Moradzadeh, A., & Mehmandost Olya, B. A. (2024). Pore size classification and prediction based on distribution of reservoir fluid volumes utilizing well logs and deep learning algorithm in a complex lithology. *Artificial Intelligence in Geosciences, 5*, 100094. doi:10.1016/j.aiig.2024.100094

[2] Bodong, Z., Maojin, T., Guohua, Z., & Chenglin, L. (n.d.). An integrating physics-informed neural networks and attention mechanism based intelligent permeability prediction method for shale reservoirs.

[3] Cao, Z., Qin, Z., Luo, S., Chen, F., Huang, K., Dong, X., Meng, L., Su, K., & Wei, K. (2026). Research and insights into the impact of different sampling strategies on machine learning-based lithology identification using well logging data. *Artificial Intelligence in Geosciences, 7*(2), 100224. doi:10.1016/j.aiig.2026.100224

[4] Guo, J., Wang, C., Wang, Y., Zhang, W., Zhao, Q., Wei, H., & Nie, X. (2026). A shale-oil reservoir multi-mineral inversion method based on well-logging data using an improved random-forest approach: Methodology and application. *Ap-*

plied Computing and Geosciences, 29, 100334. doi:10.1016/j.acags.2026.100334

[5] Haritha, D., Satyavani, N., & Ramesh, A. (2025). Generation of missing well log data with deep learning: CNN-bi-LSTM approach. *Journal of Applied Geophysics, 233*, 105628. doi:10.1016/j.jappgeo.2025.105628

[6] Hussain, W., Luo, M., Ali, M., Sadiq, I., Kasala, E. E., Aziz, T., & Batool, Z. (2025). Hybrid modeling of deep neural networks and unsupervised machine learning algorithms for missing well log prediction based on geological lithofacies similarities. *Journal of Applied Geophysics, 241*, 105846. doi:10.1016/j.jappgeo.2025.105846

[7] Kang, Z., Liu, Y., Wang, X., & Du, Y. (2026). A deep learning-based method for well log data reconstruction in marine carbonate reservoirs. *Journal of Applied Geophysics, 246*, 106106. doi:10.1016/j.jappgeo.2026.106106

[8] Lai, J., Su, Y., Xiao, L., Zhao, F., Bai, T., Li, Y., Li, H., Huang, Y., Wang, G., & Qin, Z. (2024). Application of geophysical well logs in solving geologic issues: Past, present and future prospect. *Geoscience Frontiers, 15*(3), 101779. doi:10.1016/j.gsf.2024.101779

[9] Li, D., Li, X., Liu, L., He, W., Li, Y., Li, S., Shi, H., & Fan, G. (2025). Prediction on rock strength by mineral composition from machine learning of ECS logs. *Energy Geoscience, 6*(2), 100386. doi:10.1016/j.engeos.2025.100386

[10] Pang, Q., Chen, C., Li, W., & Pang, S. (2026). Lithology identification from missing well log data: A multi-constraint guided self-supervised diffusion framework. *Advanced Engineering Informatics, 74*, 104697. doi:10.1016/j.aei.2026.104697

[11] Park, J., Jeong, J., & Holden, E.-J. (2026). Deep-learning-based hybrid model with iterative lithology constraints for the enhanced prediction of missing well-logs. *Computers & Geosciences, 209*, 106106. doi:10.1016/j.cageo.2026.106106

[12] Pothana, P., & Ling, K. (2025). Physics-integrated neural networks for improved mineral volumes and porosity estimation from geophysical well logs. *Energy Geoscience, 6*(2), 100410. doi:10.1016/j.engeos.2025.100410

[13] Sun, Y., Pang, S., Li, H., Wang, W., & Qiao, S. (2026). Architectural advancements in dual-contrastive learning: A framework for precision lithology identification from geophysical well logging data. *Journal of Applied Geophysics, 251*, 106310. doi:10.1016/j.jappgeo.2026.106310

[14] Wu, X., Xu, H., Zhou, H., Wang, L., Jiang, P., & Wu, H. (2024). Improving lithofacies prediction in lacustrine shale by combining deep learning and well log curve morphology in sanzhao sag, songliao basin, China. *Computers & Geosciences, 193*, 105735. doi:10.1016/j.cageo.2024.105735

[15] Yang, Y.-Q., Qiu, L.-W., Gregg, J., Shi, Z., & Yu, K.-H. (2016). Formation of fine crystalline dolomites in lacustrine carbonates of the eocene sikou depression, bohai bay basin, east China. *Petroleum Science, 13*(4), 642–656. doi:10.1007/s12182-016-0128-z