

## ARTICLE

## GeoSSL: A geology-aware self-supervised framework for fault detection in 3D seismic data

Mengge Wang<sup>1\*</sup>, Xinrong Hu<sup>2</sup>, Junyu Zhang<sup>3</sup>, and Haofa Lin<sup>4</sup><sup>1</sup>School of Mathematics and Computer Science, Hanjiang Normal University, Shiyan, Hubei, China<sup>2</sup>School of Computer and Artificial Intelligence, Wuhan Textile University, Wuhan, Hubei, China<sup>3</sup>School of Computer Science & Technology, Beijing Institute of Technology, Beijing, China<sup>4</sup>School of Artificial Intelligence, Wuhan Technology and Business University, Wuhan, Hubei, China

## Abstract

Fault identification is a critical task in 3D seismic interpretation, directly influencing the efficiency and accuracy of reservoir characterization and drilling decisions. However, traditional methods rely heavily on manual experience and high-quality annotated data, making it difficult to adapt to the demands of processing massive amounts of seismic data. To address this, we introduce a novel self-supervised learning (SSL) paradigm specifically designed for geological feature learning, which for the first time unifies masked voxel reconstruction, 3D patch contrastive learning, and multimodal attribute joint contrast into a coherent multi-task pre-training framework. This framework is uniquely tailored to leverage the spatial continuity and physical attributes of seismic data, enabling the model to learn transferable, geologically meaningful representations without any manual labels. It leverages unlabeled seismic data to learn geologically meaningful feature representations and leverages a transfer learning mechanism to achieve high-precision fault identification with small sample sizes. Experiments on multiple public and field-measured datasets, including F3 and SEAM Phase I, demonstrate that this method achieves key metrics such as intersection over union (IoU) and F1-scores of 0.76 and 0.87, respectively, significantly outperforming traditional attribute analysis (IoU = 0.49) and supervised deep learning models (IoU = 0.72). Furthermore, the method remains robust in areas with low signal-to-noise ratios (average confidence > 85%) and consistently estimates fault strike and dip (average absolute error  $\leq 3.5^\circ$ ). This research provides an effective solution for reducing reliance on manual annotation and improving the reliability of automated fault interpretation in complex tectonic areas.

**\*Corresponding author:**Mengge Wang  
(wangmengge@hynu.edu.cn)

**Citation:** Wang M, Hu X, Zhang J, Lin H. GeoSSL: A geology-aware self-supervised framework for fault detection in 3D seismic data. *J Seismic Explor.* 2026;35(2):025510127.  
doi: 10.36922/JSE025510127

**Received:** December 20, 2025**Revised:** February 24, 2026**Accepted:** March 10, 2026**Published online:** April 22, 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.

**Keywords:** Self-supervised learning; 3D seismic interpretation; Fault identification; Contrastive learning; Multimodal pre-training

## 1. Introduction

Commonly used in oil and gas exploration and geological hazard prediction, 3D seismic data provide rich subsurface structural information, making them highly valuable for resource assessment and risk management.<sup>1</sup> Fault identification and interpretation are crucial components of seismic data interpretation. Accurate characterization of fault systems not only impacts the effectiveness of reservoir evaluation but also directly impacts

the success rate of drilling deployment and the profitability of exploration and development.<sup>2</sup> With advances in seismic acquisition technology and the rapid growth of data volumes, traditional manual interpretation methods are no longer able to meet the demands for efficient and detailed interpretation.<sup>3</sup> There is an urgent need to develop automated and intelligent fault identification technologies.

The present mainstream fault identification methods can be classified into three categories: manual interpretation, traditional attribute analysis, and supervised deep learning. Manual interpretation relies on the empirical judgment of geologists. While highly accurate, it is inefficient and highly subjective, often resulting in significant inter-interpreter variability. Traditional seismic attribute analysis methods—such as coherence volumes and curvature attributes—enhance fault responses through mathematical transformations, improving interpretation efficiency to a certain extent.<sup>4</sup> However, these methods often rely on manually designed features, have limited generalization capabilities under complex geological conditions, and are ineffective in identifying weak faults. Supervised deep learning methods can automatically learn fault features and perform well on large amounts of labeled data. However, their performance is highly dependent on high-quality training samples. The high cost and time required to label earthquake faults severely restrict the widespread application of these methods in practical applications.

Self-supervised learning (SSL), an emerging paradigm in machine learning, can automatically learn effective representations from unlabeled data, offering a new approach to addressing the scarcity of labeled data.<sup>5</sup> In geophysics, SSL holds enormous potential. 3D seismic data, with its rich spatial structure and physical laws, provide an ideal data foundation for SSL.<sup>6–8</sup> However, the application of this technology in seismic interpretation is still in its infancy. In fault identification, many research gaps remain, including how to design effective self-supervised pre-training tasks, how to transfer learned features to downstream tasks, and how to ensure the reliability of models in practical applications.

The key innovations and contributions of this work are fourfold, directly addressing the aforementioned gaps:

- (1) A novel, multi-task SSL framework is proposed that synergistically combines masked reconstruction (for local texture), 3D contrastive learning (for global context), and multimodal attribute joint comparison (for physical consistency). While individual SSL tasks have been explored in geophysics, to the best of our knowledge, this work represents one of the first attempts to synergistically combine them

within a coherent framework specifically tailored for 3D seismic fault interpretation. The originality of our approach lies particularly in the multimodal joint contrast mechanism, which, unlike generic contrastive learning applied to a single data type, explicitly leverages the complementary geophysical information from multiple seismic attributes (such as amplitude, phase, coherence). By contrasting these different views of the same subsurface location, the model is forced to learn representations that are invariant to attribute-specific noise and more faithful to the underlying geology, thereby enhancing feature robustness and transferability.

- (2) A boundary-sensitive feature enhancement module is introduced driven by seismic gradient information. This module, a significant departure from generic SSL architectures that treat all voxels equally, explicitly guides the model to attend to fault discontinuities by weighting features based on the local gradient magnitude. This design is not merely an architectural addition but a geology-aware inductive bias that prioritizes learning from regions most indicative of faults, thereby improving the delineation of subtle and complex fault boundaries.
- (3) An efficient transfer learning pipeline is developed, that demonstrates the superior versatility of our pre-trained features across multiple high-value interpretation tasks, such as fault segmentation, classification, and strike/dip estimation, and achieves high accuracy with very few labeled samples.
- (4) Extensive, rigorous validation is provided not only on standard metrics but also on geological consistency, uncertainty quantification, and practical application (such as fault seal analysis), establishing a new benchmark for reliability in automated seismic interpretation.

This study focuses on the identification of structural faults (such as normal, reverse, strike-slip) within post-stack 3D seismic amplitude data and derived attributes. The proposed framework is designed for regions with moderate to good data quality and is validated in both siliciclastic (F3) and complex salt-influenced (SEAM) settings. Its applicability to areas with extremely low signal-to-noise ratio (SNR), heavily karstified strata, or volcanoclastic sequences requires further investigation and potential adaptation.

This study is organized as follows: Section 2 reviews related research; Section 3 details the dataset, preprocessing pipeline, and experimental platform; Section 4 elaborates on the design and implementation of the self-supervised model; Section 5 discusses the fault identification and

interpretation method; Section 6 presents experiments and results analysis; Section 7 verifies the model's generalization and practical application; Section 8 discusses its innovations and limitations; and Section 9 concludes the paper and outlines future research directions. Through systematic theoretical exploration and experimental validation, this study aims to advance 3D seismic fault identification technology toward a more intelligent and efficient direction.

To ensure conceptual clarity and delineate the scope of application, we explicitly define the methodological boundaries and key terminology. The proposed framework is developed for structural interpretation of conventional post-stack 3D seismic amplitude data and derivative attributes, targeting seismically resolvable fault discontinuities. It is architecturally optimized for structural fault identification rather than lithological fault discrimination, sub-resolution fracture characterization, or direct fluid prediction. Critical computational constructs are operationally defined: the fault probability volume denotes a 3D confidence field where each voxel value represents the model's estimated likelihood of fault presence, while pseudo-labeling refers to the iterative process of generating training labels from high-confidence model predictions on unlabeled data.

## 2. Review of related research

The application of SSL techniques in 3D seismic data interpretation and fault identification has attracted considerable attention in recent research. These methods aim to improve the accuracy and generalization of fault detection without over-reliance on labeled datasets. Such as a recent study highlighted the role of self-supervised pre-training in earthquake fault identification, demonstrating improvements in 3D segmentation.<sup>9</sup> This approach exploits the inherent structure of seismic data to learn meaningful representations, which facilitate more accurate fault delineation.

The development of specialized models such as SeisCoDE further demonstrates progress in this area. SeisCoDE is a foundational model for 3D seismic interpretation that employs a novel SSL framework that combines seismic signal processing and attribute analysis to preserve structural integrity.<sup>10</sup> These models highlight the importance of SSL in capturing seismic features that are crucial for fault interpretation. Similarly, contrastive reconstruction-based methods utilize self-supervised representation learning to improve the generalization of fault detection networks, enabling more robust identification across diverse seismic datasets.<sup>11</sup> Attention-enhanced architectures have also contributed to this field,

with models like Attention FaultFormer combining 3D Convolutional Neural Network (CNN) and Transformer modules in a self-supervised pre-training paradigm to improve fault identification performance.<sup>12</sup> These models benefit from the Transformer's ability to capture long-range dependencies and contextual information, which is crucial for accurate fault interpretation in complex seismic volumes.

In addition, semi-supervised methods that combine SSL with lightweight modules have been proposed to improve existing Segment Anything Models for earthquake fault detection, further demonstrating the versatility of SSL techniques.<sup>13</sup> End-to-end neural networks like ResACEUnet embody the trend of models that can directly interpret 3D seismic data without human intervention, leveraging SSL to enhance feature extraction and fault delineation.<sup>14</sup>

The consensus in the literature suggests that machine learning, especially SSL, is significantly advancing seismic interpretation because it enables models to effectively learn from unlabeled data. This shift not only reduces the reliance on large, labeled datasets but also improves the robustness and accuracy of fault detection systems.<sup>15-18</sup> Collectively, these studies highlight the great potential of SSL-based methods in advancing the field of 3D seismic fault identification and interpretation.

Despite these promising advances, critical gaps remain in applying SSL to seismic fault identification. First, existing methods often rely on a single pre-training task (such as reconstruction or contrastive learning), failing to holistically capture both the detailed local texture (essential for fault edges) and the global semantic context (essential for fault continuity) of 3D seismic volumes. Second, the rich, physically meaningful information contained in multiple seismic attributes is underexploited in a joint, contrastive manner. Third, there is a lack of a streamlined pipeline that effectively transfers SSL-learned features to diverse downstream interpretation tasks (segmentation, parameter estimation) with minimal supervision. Finally, the robustness and geological plausibility of SSL-based predictions in low-SNR areas and complex structures require systematic validation.

## 3. Research data and experimental platform

### 3.1. Description of 3D seismic dataset

This study used three different types of 3D seismic datasets to train and validate the fault identification algorithm, covering both simulated data and real exploration data. The F3 Demo dataset comes from the North Sea region

of the Netherlands and contains typical inclined bed and delta sedimentary structures, providing rich seismic phase and reflection surface annotation information. SEAM Phase I is an authoritative numerical simulation dataset for deepwater saline areas with precise physical parameters and synthetic seismic data, which is suitable for performance testing of the algorithm in complex geological environments. The actual data in the exploration area comes from real oil and gas exploration projects, including manually annotated fault lines and wellbore verification records, which provide an important basis for evaluating the effectiveness of the algorithm in practical applications. The dataset comparison is shown in Table 1.

These three datasets have their own characteristics in spatial resolution, temporal sampling rate and data scale, and together constitute a comprehensive and representative test benchmark.

### 3.2. Data preprocessing process

Both the raw 3D seismic data and annotations require preprocessing to adapt to SSL and downstream fault

identification tasks. The process is as follows:

The volume data is normalized and augmented. The raw amplitude volume is  $D(i, j, k)$ , where  $i$  and  $j$  are the inline and crossline indices, respectively, and  $k$  is the time sampling point index.<sup>19</sup> Normalization is defined as:

$$\tilde{D}(i, j, k) = \frac{D(i, j, k) - \mu}{\sigma} \quad (1)$$

where  $\mu = \mathbb{E}[D]$  and  $\sigma = \sqrt{\mathbb{E}[(D - \mu)^2]}$ , and the statistics are calculated over the entire data volume or on training sub-volume patches.

This study employed multiple data augmentation techniques to improve the model's generalization and robustness. Random flipping reverses the data in either the inline or crossline directions, increasing the diversity of the training data. Random rotation transforms slices by multiples of 90 degrees, effectively simulating geological structures in different orientations. Random cropping samples patches of varying sizes ( $64 \times 64 \times 64$ ) from the original volume data, ensuring that the model can handle geological features at varying scales. Gaussian noise

**Table 1. Comparison of 3D seismic fault identification datasets**

Dataset name	Region/Tectonic type	Spatial resolution (inline $\times$ crossline spacing)	Time sampling rate/Time window	Attribute and annotation information	Number of samples/ voxels
F3 Demo	Netherlands North Sea / including inclined bed, deltaic	25 m $\times$ 25 m	1.848 ms sampling step/0–1200 ms	Contains multiple seismic phases, several reflection surfaces (horizons), wellbore data	~386.9 km <sup>2</sup> area; voxel size approximately 400 $\times$ 400 $\times$ 600
SEAM Phase I	Simulation models of deepwater saline areas and other complex structures	20 m $\times$ 20 m $\times$ 10 m (after depth/ time conversion)	8 ms/full- depth model	Density, P-S wave velocity, synthetic lithology facies, and synthetic seismic data	Area approximately 35 km $\times$ 40 km $\times$ 15 km depth
Actual data of the exploration area	Oil and gas exploration area	25 m $\times$ 25 m	2 ms/0–2000 ms	Fault lines have been marked, and some wellbore penetration records have been recorded.	Voxel size: 400 $\times$ 400 $\times$ 1000, approximately $1.6 \times 10^8$ voxels



addition, implemented after data normalization, enhances the model's robustness to data noise by introducing normally distributed noise  $n(i, j, k) \sim \mathcal{N}(0, \sigma_n^2)$ , with a noise intensity  $\sigma_n$  set to 0.01. These augmentation strategies work together to significantly expand the scope of the training dataset and improve the model's adaptability and stability under diverse geological conditions.

To rigorously evaluate model performance, we established a clear split for the three datasets. For the F3 Demo dataset, 80% of the patches were randomly selected for training, 10% for validation to tune hyperparameters, and the remaining 10% for testing. The SEAM Phase I dataset was held out entirely for cross-dataset generalization testing and did not participate in any training phase. The real-field exploration dataset was divided based on spatial location: a complete subsurface structural block (approximately 70% of the volume) was used for training, while an adjacent, independent structural block (approximately 30% of the volume) was used for testing. This spatial split ensures independence between the training and test sets, preventing data leakage. All reported performance metrics are calculated based on the respective test sets.

To address the prevalent class imbalance in fault detection tasks (where fault voxels are significantly outnumbered by non-fault voxels), we adopted a combined strategy during training. First, we incorporated Dice Loss into our loss function (as shown in **Equation 11**), which is inherently robust to class imbalance. Second, during data loading, we implemented an online hard example mining technique that dynamically increases the sampling weight of patches containing fault voxels, ensuring the model adequately learns from these scarce but critical features.

In addition to class imbalance, we also implemented systematic noise processing and quality control to address data quality issues in the original F3 data. The noise primarily stems from shallow seafloor spurious waves, out-of-band high-frequency noise, and amplitude drift. We employed dip-steered median filtering, with a filter radius set to the neighborhood of two inline/crossline lines surrounding the trace, effectively eliminating random noise while preserving fault edge features. A bandpass filter with a passband range of 5–60 Hz, adjusted based on the spectral analysis results of each dataset, effectively removed low-frequency seismic background trends and high-frequency noise interference. Edge trimming specifically addresses undersampled traces at volume boundaries or in areas of low coverage, ensuring uniform data quality. Wellbore alignment correction accurately converted seismic fault travel times into depth information through time-depth calibration of seismic-well data, precisely

aligning reflection surfaces with corresponding markers in the wellbore, significantly improving the spatial accuracy of the data.

For data annotation quality control, we established a rigorous quality assurance system. For actual fault line or surface annotations, experts conducted geometric consistency reviews to ensure accuracy. For horizon or facies labels, the system checked for gaps or overlaps between labels to ensure annotation integrity. Consistency verification was performed between pixel-level and voxel-level annotations to ensure consistency across different scales. In terms of signal enhancement, we employ a variety of techniques to improve the discernibility of weak faults, including the generation of auxiliary channels using fault-enhancing attributes such as amplitude differences, variance attributes, and similarity attributes. These auxiliary channels, as an important component of the model input, significantly enhance the model's ability to detect weak fault features. The final training input is fed into the network in the form of blocks, typically of size  $P \times P \times T$ , such as  $64 \times 64 \times 64$  or  $32 \times 32 \times 128$ , where  $T$  represents the number of sampling points in the time dimension.

### 3.3. Experimental platform and hardware environment

This study constructed a high-performance computing platform to ensure the reproducibility and computational efficiency of large-scale 3D seismic data fault identification experiments. This hardware platform utilizes a heterogeneous architecture design combining graphics processing units (GPUs) and central processing units (CPUs), specifically optimized for the memory-intensive and massively parallel computing requirements of seismic volume data processing. Four NVIDIA RTX 4090 GPUs provide powerful parallel computing capabilities suitable for self-supervised pre-training and large-scale volume training tasks. A dual-socket Intel Xeon Silver 4214 CPU system provides 48 logical processing cores for data pre-processing and auxiliary computing. 512GB of large-capacity memory ensures simultaneous loading and processing of multiple large seismic volume data blocks, while the high-speed NVMe SSD storage system significantly improves data read efficiency and reduces I/O bottlenecks. A 100 Gbps high-speed network interconnect supports rapid data transfer between multiple nodes, providing the necessary network bandwidth for distributed training. This hardware configuration provides a solid infrastructure for the efficient execution of deep learning algorithms on 3D seismic data.

This study employed advanced deep learning

frameworks and tool environments to ensure experimental reliability and reproducibility. The deep learning frameworks primarily used were PyTorch (Version  $\geq 1.11$ ) and TensorFlow (Version  $\geq 2.6$ ). Cross-validation was used to ensure algorithm correctness and consistent results. Code management and experimental logging were performed using the Git version control system, combined with visualization tools such as Weights & Biases or TensorBoard, enabling complete tracking of experimental progress and visual monitoring of results. Data formatting supported the industry-standard SEG-Y format (Via the segyio library) as well as the efficient HDF5 format, ensuring accurate reading and efficient storage of seismic data. Visualization employed a multi-layered toolkit: Matplotlib and Seaborn for 2D slices and statistical charts, and Mayavi and Paraview for interactive visualization of 3D volumetric data, providing comprehensive visual support for result analysis and interpretation.

The training process employed carefully optimized parameter configuration and run strategies. The batch size setting carefully considered the balance between computing resources and model performance. On a single GPU, batches of 2 were used to process  $64 \times 64 \times 64$  blocks, and multi-GPU parallel training enabled a total batch size of 8. The optimizer uses the Adam or AdamW algorithm, with an initial learning rate set to  $\text{lr}_0 = 1 \times 10^{-4}$ . Learning rate decay strategies such as cosine decay or step decay are

employed, dynamically adjusting the learning rate based on the number of training epochs or steps to ensure stable convergence during training. The training cycle is tailored to the specific task. Self-supervised pre-training typically requires 200–300 epochs to fully learn data features, while downstream tasks such as fault segmentation and classification typically achieve optimal performance with 100–150 epochs. To ensure reproducibility, a fixed random seed (such as  $\text{seed} = 42$ ) was used across all experiments to eliminate the influence of random factors and ensure the reliability and comparability of the research conclusions.

By combining public and real-world datasets, a systematic preprocessing strategy, and a powerful computing platform, this research provides a solid foundation for subsequent self-supervised pre-training and fault identification tasks.

#### 4. Model design based on SSL

To fully exploit the fault geometry and amplitude characteristics in 3D seismic data, a multi-task framework based on SSL is constructed. This framework first pre-trains large-scale unlabeled data through tasks such as mask reconstruction, 3D patches contrast learning, and joint comparison of multimodal attributes. The pre-trained features are then transferred to the downstream task of fault identification and interpretation. The overall concept

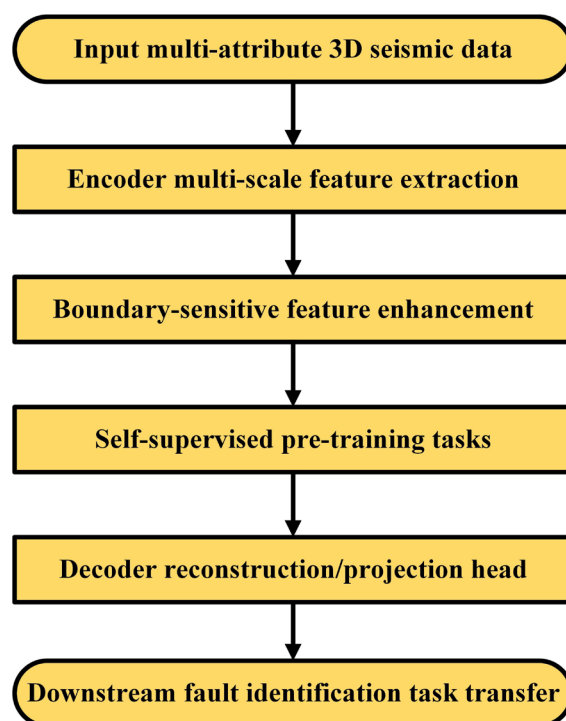


Figure 1. Model architecture flow chart

is shown in Figure 1.

To provide a clear technical blueprint of the proposed framework, we present a detailed schematic of the model

architecture and data processing flow in Figure 2. This figure illustrates the two-phase pipeline: (1) the multi-task self-supervised pre-training on unlabeled seismic

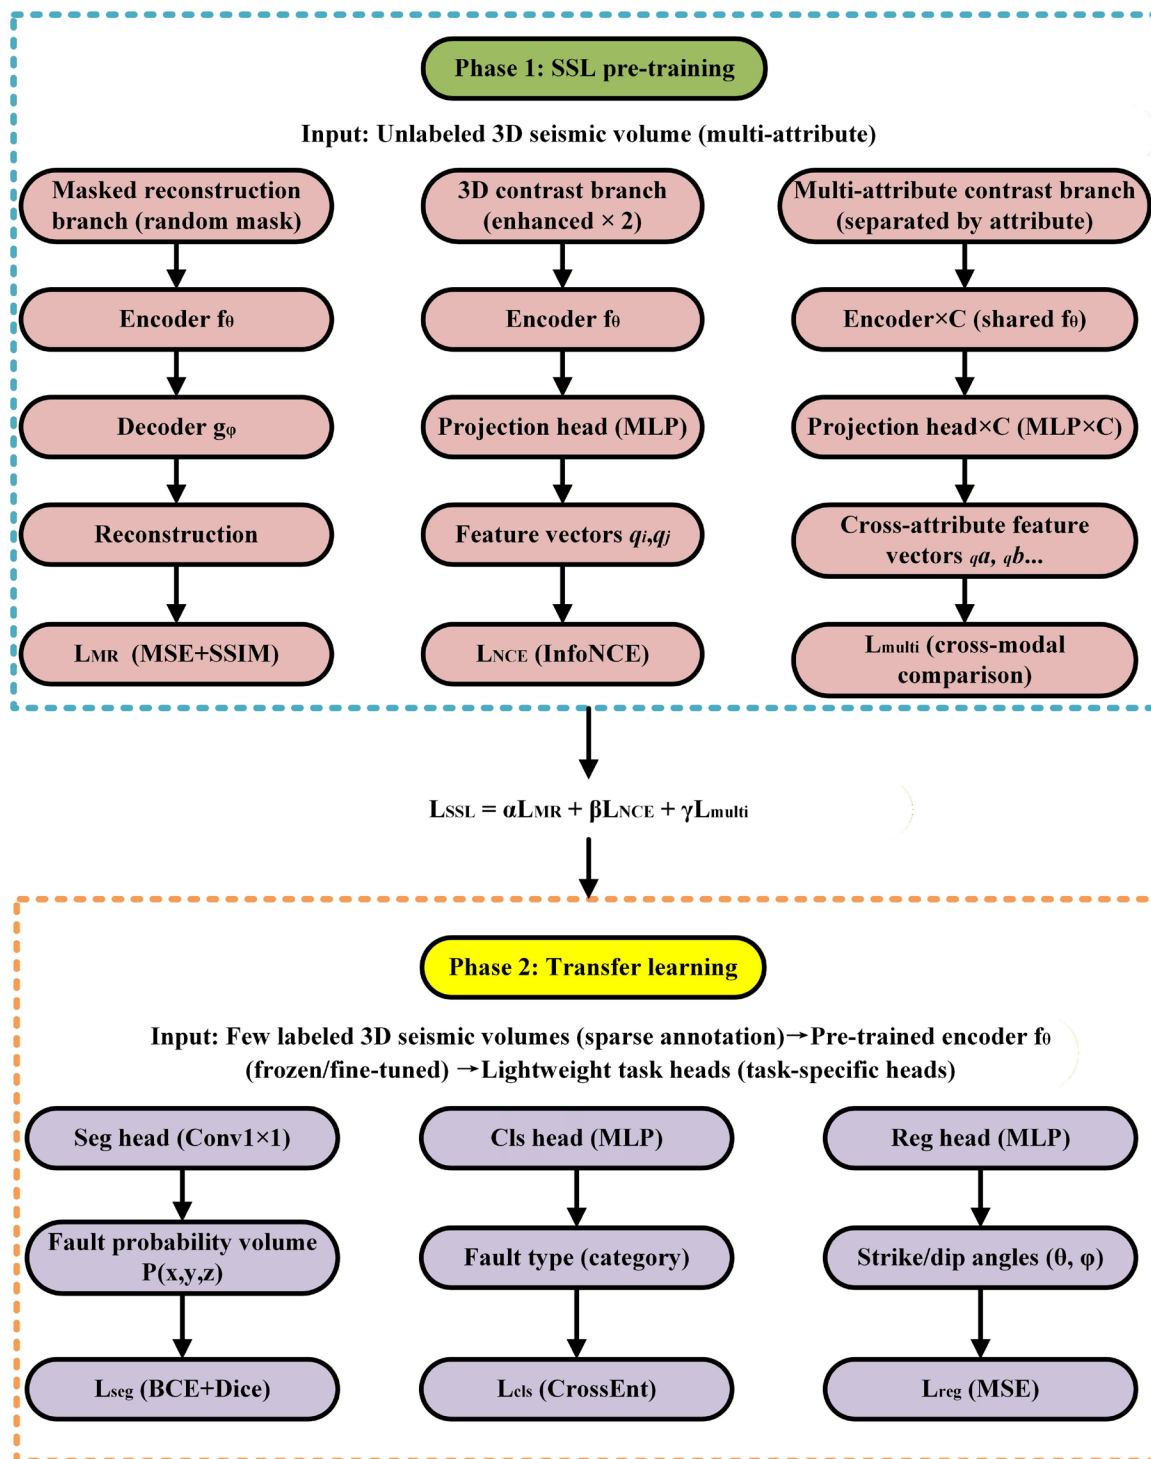


Figure 2. Detailed model architecture and data flow diagram

volumes, and (2) the transfer learning for downstream fault interpretation tasks.

Figure 2 detailed architecture and data flow of the proposed SSL-based fault interpretation framework. (1) Self-supervised pre-training phase: An unlabeled 3D seismic volume (or multi-attribute volume) is processed. It undergoes data augmentation and is fed into three parallel SSL tasks: Masked Reconstruction (with a masking module and encoder-decoder), 3D Patch Contrastive Learning (with two augmented views processed by a shared encoder and projection head), and Multimodal Joint Contrast (with different attribute channels processed and compared). Their losses are combined to update the encoder. (2) Transfer learning phase: The pre-trained encoder is frozen or fine-tuned, and connected to lightweight, task-specific heads (such as segmentation, classification, regression heads) for supervised training on a small set of labeled data, yielding fault probability volumes, classifications, or geometric attributes.

In terms of architectural design, a 3D convolutional encoder-decoder integrated with a multi-scale feature extraction module is employed to capture sensitive information about fault boundaries.<sup>20</sup> Let the original normalized input patch be  $X \in \mathbb{R}^{C \times D \times H \times W}$ , where  $C$  is the number of channels ( $C > 1$  for multi-attribute input), and  $D, H, W$  are the depth, row, and column dimensions, respectively. The encoder  $f_\theta$  maps this to latent features  $Z$ , which are then reconstructed by the decoder  $g_\phi$  or other self-supervised outputs  $\hat{X}$ .

$$Z = f_\theta(X), \hat{X} = g_\phi(Z) \quad (2)$$

The encoder consists of multiple 3D convolutional blocks (kernel size  $3 \times 3 \times 3$ ) with residual connections.<sup>21</sup> The multi-scale module introduces different receptive fields  $r$ :

$$Z_r = \text{Conv3D}_r(X), Z = \sum_{r \in \{1, 2, 4\}} Z_r \quad (3)$$

where  $r$  represents the dilation rate or pooling scale. Features of different scales are concatenated in the channel dimension and fed into the fault boundary sensitivity module. This module uses gradient magnitude or similarity as weights to enhance fault neighborhood features:

$$W(i, j, k) = \exp\left(-\frac{\|\nabla X(i, j, k)\|^2}{2\sigma^2}\right), Z_{\text{fault}} = W \odot Z \quad (4)$$

This allows the fault edge to be highlighted in the latent space while suppressing the background continuous reflection.

Table 2 lists the main layers of the model architecture

and parameter configuration.

This network utilizes an encoder-decoder architecture and is optimized for feature extraction and fault probability prediction from 3D volume data. The encoder layer progressively extracts feature representations at different levels through multi-scale convolution and downsampling operations. Block 2 employs downsampling to halve the spatial dimension while increasing the number of channels. The boundary-sensitive module specifically enhances the ability to extract fault edge features through a gradient weighting mechanism. The decoder layer progressively restores spatial resolution through upsampling and convolution operations, and the final output layer uses a sigmoid activation function to generate a fault probability map. Techniques such as residual connections and batch normalization are incorporated into the network design to ensure training stability and effective feature transfer.

During the self-supervised pre-training phase, three tasks are designed:

(1) Masked reconstruction:

Randomly mask a portion of the voxels in the input block (mask ratio  $m$ , such as 30%–50%) and only input the remaining portion  $X_{\text{visible}}$  into the encoder. The decoder needs to reconstruct the complete block  $\hat{X}$ . The loss function is the reconstruction means squared error (MSE):<sup>22–24</sup>

$$\mathcal{L}_{\text{MR}} = \frac{1}{|\Omega_m|} \sum_{(i, j, k) \in \Omega_m} (\hat{X}(i, j, k) - X(i, j, k))^2 \quad (5)$$

where  $\Omega_m$  is the set of masked voxels. We also introduce the structural similarity (SSIM) term to improve the reconstruction quality:<sup>25</sup>

$$\mathcal{L}_{\text{SSIM}} = 1 - \text{SSIM}(\hat{X}, X) \quad (6)$$

The total loss  $\mathcal{L}_{\text{MR+SSIM}} = \mathcal{L}_{\text{MR}} + \lambda \mathcal{L}_{\text{SSIM}}$ .

(2) 3D patch-wise contrastive learning:

The two views  $Z_i, Z_j$  generated at different positions or in different enhancement methods are fed into the projection head  $h(\cdot)$  to obtain vectors  $q_i, q_j$ . The InfoNCE loss is used:<sup>26</sup>

$$\mathcal{L}_{\text{contrast}} = -\log \frac{\exp(\cos(q_i, q_j) / \tau)}{\sum_{k=1}^N \exp(\cos(q_i, q_k) / \tau)} \quad (7)$$

where  $\cos(q_i, q_j) = \frac{q_i^\top q_j}{\|q_i\| \|q_j\|}$ , and  $\tau$  is the temperature

Table 2. Model architecture and parameter configuration

Module/Layer name	Input size	Convolution kernel/ stride	Number of channels	Activation function	Other notes
Encoder Block1	$64 \times 64 \times 64 \times C$	$3 \times 3 \times 3 / 1$	32	ReLU	3D convolution + batch normalization
Encoder Block2 (Down)	$32 \times 32 \times 32 \times 32$	$3 \times 3 \times 3 / 2$	64	ReLU	Multi-scale convolution/Dilated convolution
Encoder Block3	$16 \times 16 \times 16 \times 64$	$3 \times 3 \times 3 / 1$	128	ReLU	Residual connection
Boundary Sensitive Mod.	$16 \times 16 \times 16 \times 128$	-	128	Sigmoid	Gradient weighted enhancement of fault edges
Decoder Block1 (Up)	$32 \times 32 \times 32 \times 64$	$3 \times 3 \times 3 / 1$	64	ReLU	Upsampling + convolution
Decoder Block2 (Up)	$64 \times 64 \times 64 \times 32$	$3 \times 3 \times 3 / 1$	32	ReLU	Upsampling + convolution
Output Layer	$64 \times 64 \times 64 \times 1$	$1 \times 1 \times 1 / 1$	1	Sigmoid	Reconstruction/Prediction of fault probability

Abbreviation: ReLU: Rectified linear unit.

coefficient. This contrastive learning paradigm, which encourages the model to learn invariant representations by maximizing agreement between different views of the same data, has proven effective not only in seismic interpretation but also in other geoscience-related fields, such as remote sensing, where consistency reasoning has been explored to enhance object detection performance.<sup>27</sup>

### (3) Multi-attribute pretraining:

The selection of seismic attributes for multimodal input is grounded in their complementary geophysical roles in fault characterization. We deliberately choose four core

attributes: amplitude, instantaneous phase, instantaneous frequency, and similarity (coherence). This combination provides a balanced and physically meaningful representation for SSL.

Amplitude provides the primary signal of subsurface impedance contrast. Instantaneous phase delivers robust sensitivity to lateral discontinuities, crucial for identifying fault-related reflector breaks even in low-amplitude zones. Instantaneous frequency supplies lithological and stratigraphic context, as its abrupt variations are often associated with fault zones due to wave diffraction effects. Similarity (coherence) is a mathematically defined

attribute engineered to highlight discontinuities directly. Its inclusion creates a powerful bootstrapping effect, guiding the model to implicitly learn fault-sensitive features from the raw data. Collectively, this attribute suite balances direct structural signals with nuanced contextual cues. Their joint contrastive learning forces the model to discover the underlying geological structures that are consistently expressed across these diverse representations, thereby yielding more generalizable and robust feature embeddings.

The input contains  $C \geq 4$  attribute channels, such as amplitude, instantaneous phase, instantaneous frequency, and similarity, and the representations of different attributes are compared in the latent space. The loss function can be defined as multimodal cross-channel comparison:<sup>28</sup>

$$\mathcal{L}_{\text{multi}} = \sum_{a \neq b} -\log \frac{\exp(\sin(\mathbf{q}_i^{(a)}, \mathbf{q}_j^{(b)}) / \tau)}{\sum_k \exp(\sin(\mathbf{q}_i^{(a)}, \mathbf{q}_k^{(b)}) / \tau)} \quad (8)$$

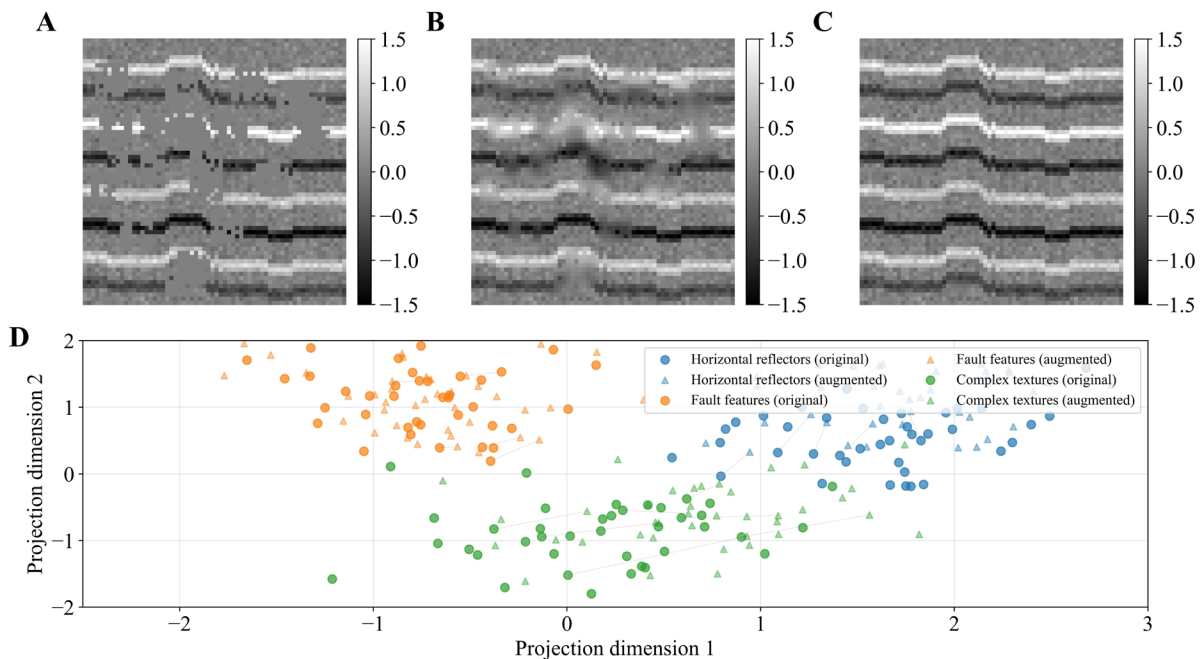
The total self-supervision loss is the weighted sum of three:

$$\mathcal{L}_{\text{SSL}} = \alpha \mathcal{L}_{\text{MR+SSIM}} + \beta \mathcal{L}_{\text{contrast}} + \gamma \mathcal{L}_{\text{multi}} \quad (9)$$

The hyperparameters  $\alpha$ ,  $\beta$ , and  $\gamma$  control the contribution of each self-supervised task to the overall pre-training objective. Through empirical experimentation on the validation set, these weights were set to  $\alpha = 1.0$ ,  $\beta = 0.5$ , and  $\gamma = 0.8$ . This configuration balances the three tasks, giving primary importance to masked reconstruction while maintaining significant contributions from contrastive learning and multimodal joint contrast. A sensitivity analysis was performed by varying each parameter within the range  $[0.2, 2.0]$  while holding others fixed; the results showed that the model performance (measured by downstream fault segmentation IoU) remained stable within  $\pm 10\%$  of the chosen values, demonstrating robustness to moderate hyperparameter variations.

The schematic diagram shown in Figure 3 clearly shows the key links of mask reconstruction and 3D patches comparison learning in the self-supervised pre-training process.

Figure 3A shows a randomly masked 3D seismic volume. A certain percentage of voxels in the original amplitude data are randomly masked. This forces the network to understand local structural and textural features with only partial information. Figure 3B shows the model's reconstruction of the masked region after



**Figure 3.** Self-supervised learning framework for seismic data: masked reconstruction and contrastive feature learning. (A) Input seismic section with patch-based masking preserving geological continuity, (B) reconstructed seismic data showing recovered reflector layers and fault features, (C) original complete seismic section with horizontal reflectors and structural fault for comparison, and (D) contrastive learning projection space where different geological feature types form distinct clusters, demonstrating effective feature representation learning.



training. The fault edges, amplitude transitions, and small-scale seismic phase features of the reconstructed volume are well restored, highly consistent with the original data in terms of geometry. **Figure 3C** shows the complete input volume, used to compare the consistency between the model reconstruction and the real data. **Figure 3D** illustrates the clustering of the projected space of the same seismic volume after different enhancements. The blue and red dots represent the two enhanced views, respectively, which naturally form clusters in the projected space. This demonstrates that the contrastive learning branch can map different views of the same semantic meaning to adjacent regions, thereby improving the discriminability and robustness of the features. This figure vividly demonstrates that after our multi-task joint optimization during the pre-training phase, the model not only fills in missing information but also clusters structural semantics in the latent space, facilitating subsequent fault identification.

The encoder features  $Z$  obtained through the above pre-training can be transferred to the downstream fault identification task, including fault classification, segmentation, and strike/dip prediction.<sup>29</sup> We introduce a lightweight segmentation head for the fault segmentation task:

$$\hat{Y} = \sigma(\text{Conv3D}_{1 \times 1 \times 1}(Z)) \quad (10)$$

where  $\hat{Y}$  is the fault probability volume. The loss function is binary cross entropy + Dice loss:

$$\mathcal{L}_{\text{fault}} = -\sum_{i,j,k} [y_{ijk} \log \hat{y}_{ijk} + (1 - y_{ijk}) \log (1 - \hat{y}_{ijk})] + 1 - \frac{2 \sum y \hat{y}}{\sum y + \sum \hat{y}} \quad (11)$$

For fault strike/dip prediction, we regress the fault strike  $\theta$  and dip  $\phi$  on  $Z$  to minimize the mean square error:<sup>30,31</sup>

$$\mathcal{L}_{\theta, \phi} = \frac{1}{N} \sum_{n=1}^N [(\hat{\theta}_n - \theta_n)^2 + (\hat{\phi}_n - \phi_n)^2] \quad (12)$$

Crucially, downstream tasks operate in two distinct modes, determined by label availability.

**Supervised Fine-tuning with Limited Labels:** For surveys with sparse manual interpretations (such as 5%–10% of the volume, as used in the main experiments described in Section 6), we directly fine-tune the pre-trained encoder using a lightweight task head and standard supervised losses (**Equations 11** and **12**). This setup leverages pre-training to achieve high performance with minimal annotation.

**Unsupervised Extraction via Pseudo-Labeling:** For completely unlabeled data, we employ a teacher-student

framework with consistency regularization. The teacher model generates an initial fault probability volume. High-confidence predictions (threshold  $\tau > 0.8$ ) are converted into pseudo-labels after morphological cleaning. The student model is trained on strongly augmented versions of the input to match these pseudo-labels via a consistency loss (such as MSE). The teacher's weights are updated via an exponential moving average of the student's weights. The total loss combines this consistency loss with the self-supervised loss (**Equation 9**) to prevent confirmation bias. For small samples or unlabeled areas, we use pseudo-labeling + consistency regularization to achieve automatic fault extraction, clustering self-supervised features or predicting fault probabilities on unlabeled data through a teacher-student network.

**Table 3** systematically demonstrates the self-supervised pre-training strategies and their transfer performance in the downstream fault identification task.

The study employed three different self-supervised pre-training methods: mask reconstruction pre-training (MSE = 0.012), 3D patches contrastive learning pre-training (InfoNCE = 0.45), and multimodal attribute joint contrastive pre-training (InfoNCE = 0.38). These methods all successfully learned meaningful feature representations to varying degrees. The pre-training phase fully utilized large amounts of unlabeled data (60,000–100,000 patches), while the downstream fine-tuning phase required only a small amount of labeled data (2,000–5,000 samples), significantly reducing the annotation cost. Multimodal attribute joint contrastive pre-training achieved the best InfoNCE loss (0.38) using four attribute channels, demonstrating that the fusion of multi-attribute information facilitates the learning of richer feature representations. The pre-trained features were successfully transferred to a variety of downstream tasks, including fault segmentation (IoU = 0.83), fault classification (Acc = 0.92), and geological parameter regression (MAE = 3.5°/2.1°), demonstrating the strong generalization capability of the pre-trained model. Particularly noteworthy is that even without annotations, pseudo-labeling combined with consistency training still achieves an IoU of 0.77, demonstrating the practicality of this approach in real-world applications.

**Figure 4** shows the effect of transferring pre-trained features to the downstream segmentation task


**Figure 4A** shows the original seismic section, which shows continuous reflectors, relatively weak fractures, and background noise. **Figure 4B** shows the manually interpreted fault annotations, with the fault locations highlighted. The outlines were drawn point by point by geological interpretation experts based on seismic

Table 3. Self-supervised pre-training and downstream task transfer configuration and performance

Task types	Number of training samples (patch)	Mask ratio m	Number of comparison views	Number of attribute channels	Migrating downstream tasks	Evaluation metrics
Mask reconstruction pre-training	100,000	0.4	-	1	Encoder pre-training	MSE = 0.012
3D patches contrastive learning pre-training	80,000	-	2	1	Encoder pre-training	InfoNCE = 0.45
Multimodal attribute joint contrastive pre-training	60,000	-	4	4	Encoder pre-training	InfoNCE = 0.38
Fault segmentation transfer (with annotations)	5,000	-	-	$\frac{1}{4}$	Segmentation head	IoU = 0.83
Fault classification transfer (with annotations)	3,000	-	-	$\frac{1}{4}$	Classification head	Acc = 0.92
Fault strike/Dip prediction (with annotations)	2,000	-	-	$\frac{1}{4}$	Regression head	MAE = 3.5°/2.1°
Small sample automatic extraction (without annotations)	-	-	-	$\frac{1}{4}$	Pseudo-labeling + consistency	IoU = 0.77

Abbreviations: Acc: Accuracy; InfoNCE: Information noise-contrastive estimation; IoU: Intersection over union; MAE: Mean absolute error; MSE: Mean squared error.





**Figure 4.** Transfer learning for fault segmentation using pre-trained seismic features. (A) Input seismic section displaying horizontal reflector sequences and fault-displaced layers; (B) Expert-annotated fault labels showing precise fault plane delineation; (C) Model-predicted fault probability map using pre-trained self-supervised features, demonstrating improved fault detection accuracy compared to traditional methods.

reflection discontinuities, which are highly subjective. Figure 4C shows the fault probability volume obtained after transferring the self-supervised pre-trained model. The color intensity indicates the model's confidence in the existence of the fault. Compared to manual annotations, the model predicts more consistent fault locations, effectively captures fractures in weak reflection areas, and displays a continuous geometric pattern of strike and dip. This demonstrates that the deep features extracted by pre-training are highly sensitive to fault boundaries. Even in the absence of sufficient labeled samples, high-precision fault probability volumes can be generated, providing more stable and automated support for seismic interpretation.

Through this design, our model acquires rich structural and amplitude characterization capabilities on large-scale unlabeled 3D seismic data, and then realizes fault classification, segmentation and geometric attribute prediction through a lightweight downstream head network. It can automatically extract fault structure even in small samples or unlabeled areas, providing a high-precision, automated tool for seismic interpretation.

## 5. Fault identification and interpretation methods

For fault extraction and segmentation, we propose a 3D fault imaging pipeline based on self-supervised features. The model first performs feature mapping on the seismic volume to obtain the fault-sensitive feature vector  $f_i \in \mathbb{R}^d$  for each voxel. The model then outputs the fault probability volume  $P(x, y, z)$  through a 3D convolutional decoder, which is expressed as:


$$P(x, y, z) = \sigma(W_2 * \phi(W_1 * F(x, y, z))) \quad (13)$$

where  $F(x, y, z)$  is the deep feature tensor output by the self-supervised model,  $W_1, W_2$  are trainable convolution kernels,  $\phi$  is a nonlinear activation function, and  $\sigma$  is a sigmoid function that constrains the output to  $[0, 1]$  to represent the fault probability.

Figure 5 shows a comparison of 3D tomographic images generated based on self-supervised features.

The original seismic section in Figure 5A primarily displays a continuous reflection layer system, where fault locations are often obscured by noise, impedance variations, and multi-stage sedimentary overlays, resulting in blurred boundaries that are difficult to accurately identify. The fault probability volume slice in Figure 5B utilizes a depth model to estimate the probability of fault presence for each voxel. The results show that the fault boundary is highlighted with a high probability value, the continuity of the background layer system is effectively suppressed, and the spatial location and geometric outline of the fault are more clearly defined. The 3D maximum projection in Figure 5C further integrates the high-confidence fault region of the entire volume data into a 2D image, making the fault network structure clear and facilitating macroscale interpretation. A comparison of the three models demonstrates that the SSL model can extract stable fault features in complex backgrounds, achieving continuous visualization from local slices to the entire 3D space.

Automatic estimation of multi-scale fault strike and dip utilizes principal direction analysis of the feature tensor.<sup>32</sup>



**Figure 5.** 3D fault imaging and visualization: from seismic data to fault network characterization. (A) Original seismic section showing continuous horizontal reflectors with fault-displaced layers and structural discontinuities, (B) fault probability slice displaying computed likelihood of fault presence with enhanced fault corridor delineation, and (C) 3D maximum projection visualization revealing complex fault network architecture including vertical, diagonal, and curved fault systems in three-dimensional space.

Structural tensor decomposition is performed on the local region of the probability volume:

$$T = \begin{bmatrix} \langle g_x^2 \rangle & \langle g_x g_y \rangle & \langle g_x g_z \rangle \\ \langle g_x g_y \rangle & \langle g_y^2 \rangle & \langle g_y g_z \rangle \\ \langle g_x g_z \rangle & \langle g_y g_z \rangle & \langle g_z^2 \rangle \end{bmatrix} \quad (14)$$

where  $g_x, g_y, g_z$  are the gradients of the probability volume in three directions. The notation  $\langle \cdot \rangle$  denotes local spatial averaging within a small analysis window centered at each voxel, typically implemented using a Gaussian or uniform kernel. By eigendecomposition, the main eigenvector  $v_{\max}$  is obtained, and the fault strike  $\theta$  and dip angle  $\phi$  are obtained:

$$\theta = \arctan2(v_y, v_x), \phi = \arctan\left(\frac{v_z}{\sqrt{v_x^2 + v_y^2}}\right) \quad (15)$$

Understanding the geometric properties of fault networks, such as their branching patterns and hierarchical organization, connects to fundamental research in graph theory and tree structures. For instance, exact combinatorial methods have been developed to analyze subtree properties in tree-like networks,<sup>33</sup> providing theoretical foundations that could inform future topological analysis of fault systems extracted from seismic data.


Figure 6 shows the automatic estimation results of multi-scale fault strike and dip.

Figure 6A superimposes local directional arrows on the fault probability volume. These arrows, derived from gradient field and tensor analysis, indicate the direction of local fault extension and reflect the geometric orientation of

the fault in space. Figure 6B shows the thermal distribution of strike, with different color ranges representing different azimuth angles. The figure clearly shows how the fault's horizontal extension direction varies with position, forming a regular strike zone. Figure 6C presents the corresponding dip distribution. The color variation reveals the vertical tilt of the fault. Some areas with larger dips may represent recently active faults, while more gradual areas may represent relics of earlier tectonic activity. Overall, the multi-scale estimation not only provides a quantitative description of the fault's spatial morphology but also lays the foundation for subsequent tectonic evolution analysis and mechanical modeling.

To quantify the model's ability to interpret fault geometry and activity, we selected eight representative faults in different tectonic regions and extracted their geometric parameters, including length, average dip, displacement, and activity period. Table 4 systematically presents the geometric characteristics and activity parameter calculations for the eight faults, fully demonstrating the excellent performance of the self-supervised feature-driven algorithm in weakly reflective regions.

Fault lengths range from 6.8 to 16.7 km, with F3 being the largest (16.7 km) and F5 the smallest (6.8 km), demonstrating the diversity of faults in the study area. Dip parameters range from 48° to 63°, typical of normal or high-angle reverse faults, with F2 having the highest dip (63°) and F5 the lowest (48°). Displacement varies significantly, with the maximum reaching 150 m (F3) and the minimum only 15 m (F5), reflecting the varying activity intensities of the faults. Activity period estimates indicate that F5 has the longest activity period (4.2 Ma), while F3 has the shortest (1.8 Ma), likely reflecting the regional tectonic evolution. Confidence levels are generally



**Figure 6.** Multi-scale automatic estimation of fault strike and dip attributes from seismic data. (A) Fault probability overlaid with local orientation vectors showing directional analysis based on fault geometry and structural trends; (B) Strike distribution heatmap displaying azimuthal variations with fault-related near-north-south trends and bedding-related east-west orientations; (C) Dip angle distribution showing steep fault planes (80–90°) contrasted with gentle sedimentary layer dips (5–25°), demonstrating structural geological differentiation.


**Table 4.** Comparison of different fault geometry and activity parameters

Fault number	Length (km)	Average inclination (°)	Maximum displacement (m)	Minimum displacement (m)	Activity period (Ma)	Confidence (%)	Toward consistency (%)
F1	12.3	58	120	30	2.1	92	88
F2	8.5	63	85	20	3.5	89	86
F3	16.7	55	150	40	1.8	95	91
F4	10.4	60	90	25	2.9	90	87
F5	6.8	48	70	15	4.2	85	80
F6	14.1	62	110	35	2.4	93	89
F7	9.9	57	95	20	3.0	88	84
F8	11.6	59	105	28	2.6	91	86

high (85%–95%), demonstrating the algorithm's ability to reliably detect faults even in areas of weak reflection, with the highest confidence level (95%) for F3. The strike consistency index ranges from 80% to 91%, demonstrating the algorithm's superiority in extracting fault continuity, particularly in areas with complex geological conditions, where it maintains high spatial consistency.

Figure 7 demonstrates the integration of fault interpretation with sedimentological context.

In Figure 7, different colors are used to distinguish sedimentary facies types (sandstone, shale, limestone, and mixed facies) with an overlay of fault probability volumes (hot colors), creating a coupled visualization of faults and sedimentary environments. Some faults penetrate multiple sedimentary facies units, exhibiting distinct cross-layer cutting characteristics (such as cutting through sandstone, shale, and limestone), which suggests strong regional tectonic control. Other faults, however, are primarily confined to a single sedimentary facies, potentially related



**Figure 7.** Integration of fault probability volumes with sedimentary facies analysis for reservoir characterization. The two color bars on the right define the two datasets: The right color bar corresponds to the four sedimentary facies types distinguished using a categorical color scheme, while the left color bar indicates the fault probability, ranging from low (transparent) to high (dark red).

to local stress field adjustments during sedimentation. This overlay interpretation helps understand the impact of faults on sedimentary facies distribution, reservoir distribution, and fluid migration pathways.

To enhance the reliability of the results, this study also conducted uncertainty quantification and credibility analysis of the fault identification results. Monte Carlo Dropout was used to calculate the variance of the fault probability volume during the inference phase using multiple forward propagations, defining the uncertainty metric:

$$U(x, y, z) = \text{Var}(P_1(x, y, z), P_2(x, y, z), \dots, P_N(x, y, z)) \quad (16)$$

where  $P_N(x, y, z)$  is the fault probability value predicted by the  $N$ -th sampling. A larger  $U$  indicates a lower confidence in the model at that voxel location.

In practical application, the workflow can be adaptive: For new seismic surveys with no prior interpretation, the model first operates in the unsupervised pseudo-labeling mode (as described in Section 4) to generate an initial fault probability volume and geometric attribute estimates. Interpreters can then selectively validate and correct high-uncertainty regions or key structural areas (as indicated by Table 5), creating a small set of high-confidence labels. These labels are then used to switch to the supervised fine-tuning mode, rapidly refining the model specifically for the local geology of the new survey. This hybrid human-in-the-loop workflow maximizes efficiency while minimizing manual annotation effort. Table 5 summarizes the statistical results for different seismic regions. Within tectonic zones with high SNRs, the model predicts higher confidence and lower standard deviations for faults. In contrast, in areas with complex backgrounds, ambiguous fault features, or

high noise, the uncertainty of the model output increases significantly, providing a quantitative reference for manual interpretation.

As shown in Table 5, the mean probability values for areas A1 and A2 are both greater than 0.8, with low standard deviations ( $<0.12$ ), indicating that the model's predictions in these areas are highly consistent and reliable. In contrast, in areas C and D, particularly the unlabeled region of D2, the mean probability is only approximately 0.37, with a high standard deviation of 0.28 and an uncertainty index close to 0.8, indicating that the model lacks sufficient discriminative power in these areas. The results must be verified through manual interpretation

and additional geological data. This approach, based on a combined analysis of mean probability and uncertainty, provides quantifiable metrics for subsequent exploration deployment and risk control, facilitating more robust fault interpretation in complex geological settings.

## 6. Experiment and results analysis

In this study, to systematically evaluate a 3D seismic fault identification method based on SSL, we designed multiple comparative experiments and introduced multi-dimensional performance evaluation metrics for analysis. These comparative experiments included traditional fault detection methods based on seismic attributes (such as

**Table 5. Statistics of mean and standard deviation of uncertainty in fault interpretation in different regions**

Area number	Geological background description	Number of sample points	Average probability means ( $\bar{\mu}$ )	Standard deviation ( $\sigma$ )	Average uncertainty index ( $U = \sigma/\mu$ )
A1	Structural highs, clear reflections	1500	0.87	0.08	0.092
A2	Anticline flanks, good continuity	1300	0.81	0.11	0.136
B1	Fault-intensive areas, high SNR	1200	0.76	0.15	0.197
B2	Regions with strong lithologic variations	1000	0.65	0.18	0.277
C1	High noise interference	900	0.54	0.21	0.389
C2	Thin interbedded sedimentary environments	800	0.49	0.24	0.490
D1	Deep areas with low SNR	700	0.42	0.25	0.595
D2	Unlabeled or blind areas	600	0.37	0.28	0.757

Abbreviation: SNR: Signal-to-noise ratio.

similarity and curvature), fully supervised convolutional neural network methods, a semi-supervised framework, and the proposed self-supervised pre-training + transfer learning method. We also conducted ablation experiments, gradually removing modules such as mask reconstruction, 3D contrastive learning, and multimodal contrast to examine the contribution of each component.

In terms of evaluation metrics, we defined fault segmentation accuracy, recall, F1-score, and IoU to measure segmentation performance. Furthermore, to better meet the needs of seismic interpretation, we introduced the consistency index (CI) to measure the consistency between model predictions and those of different geological interpreters. Let the set of ground truth faults be  $Y$  and the set of predictions be  $\hat{Y}$ , then the IoU is defined as:

$$\text{IoU} = \frac{|Y \cap \hat{Y}|}{|Y \cup \hat{Y}|} \quad (17)$$

The consistency index CI is defined as the weighted ratio of the intersection and union between multiple interpreters:

$$\text{CI} = \frac{1}{M} \sum_{i=1}^M \frac{|Y_i \cap \hat{Y}|}{|Y_i \cup \hat{Y}|} \quad (18)$$

where  $Y_i$  is the fault label of the  $i$ -th interpreter, and  $M$  is the number of interpreters.

Table 6 systematically compares the performance of five different methods in the fault identification task, clearly showing their respective technical characteristics and effect differences.

Beyond in-domain evaluation, we further validated the model's generalization capability on unseen data through a cross-dataset transfer experiment. Specifically, we performed self-supervised pre-training and supervised fine-tuning solely on the F3 dataset from the Dutch North Sea. The trained model was then applied directly for inference on the SEAM Phase I deepwater salt basin synthetic dataset, without any re-training or fine-tuning using SEAM labels. This setup simulates the model's performance when faced with a completely new geological environment (transitioning from a relatively simple deltaic setting to a complex salt-influenced regime) and different acquisition parameters.

Table 7 presents the fault segmentation results of our model on the SEAM dataset, comparing it against a supervised model trained directly on SEAM (Oracle, serving as an upper-performance bound) and a supervised model trained on F3 (without pre-training).

The results demonstrate that when a standard supervised model trained on F3 is directly transferred to SEAM, its performance drops sharply, achieving an IoU of only 0.42 on SEAM compared to 0.72 on its native F3 test set. This significant performance degradation indicates substantial overfitting to the geological context of the

**Table 6. Performance comparison of different methods in fault identification tasks**

Method category	Accuracy	Recall	F1-score	IoU	CI
Similarity attribute + threshold detection	0.74	0.58	0.65	0.49	0.61
Curvature attribute method	0.77	0.61	0.68	0.52	0.64
Supervised convolutional neural network	0.88	0.79	0.83	0.72	0.81
Semi-supervised method	0.85	0.75	0.79	0.68	0.78
Self-supervised + transfer learning in this paper	0.91	0.83	0.87	0.76	0.85

Abbreviations: CI: Consistency index; IoU: Intersection of union.



Table 7. Cross-dataset generalization results (trained on F3 → tested on SEAM)

Method	Precision	Recall	F1-score	IoU
Supervised model (trained on SEAM, Oracle)	0.89	0.81	0.85	0.74
Supervised model (trained on F3, no pre-training)	0.68	0.52	0.59	0.42
GeoSSL (pre-trained & fine-tuned on F3, tested on SEAM)	0.82	0.74	0.78	0.65

Abbreviation: IoU: Intersection of union.


training data. In contrast, our GeoSSL model, leveraging its robust feature representations learned through multi-task self-supervision on F3, exhibits superior generalization, achieving an F1-score of 0.78 and an IoU of 0.65 on SEAM. While there is still a gap compared to the Oracle model trained directly on SEAM (IoU = 0.74), the performance degradation of the GeoSSL model is far less severe than that of the purely supervised model. This strongly suggests that the GeoSSL framework learns more universal and transferable geological features, rather than merely memorizing the texture of a specific dataset. This result directly addresses the challenge of geological variability in practical applications and validates the model's significant potential for real-world exploration scenarios.

Looking at overall performance trends, traditional seismic attribute methods (similarity and curvature attributes) perform relatively poorly across all metrics. The similarity attribute plus threshold detection method exhibits the weakest overall performance, with an accuracy of 0.74, a recall of only 0.58, an F1-score of 0.65, an IoU of 0.49, and a confidence index of 0.61. This indicates that traditional methods suffer from significant underdetection issues in fault identification. The curvature attribute method shows slight improvement, with all metrics showing improvement, but still hovering between 0.6 and 0.7, reflecting the limitations of handcrafted feature-based methods. Supervised convolutional neural network methods demonstrate significant advantages, reaching an accuracy of 0.88, an F1-score of 0.83, an IoU of 0.72, and a confidence index of 0.81, demonstrating the powerful capabilities of deep learning for feature learning. Semi-supervised methods maintain good performance despite

limited labeled data, achieving an accuracy of 0.85 and an F1-score of 0.79, but performing slightly worse than fully supervised methods. The self-supervised combined with transfer learning approach proposed in this paper achieved state-of-the-art results across all evaluation metrics, with precision reaching 0.91, recall reaching 0.83, F1-score of 0.87, IoU of 0.76, and confidence index of 0.85. This significant improvement in recall demonstrates that this approach effectively reduces missed detections and comprehensively outperforms traditional and supervised learning methods, demonstrating the superior performance and practical value of combining self-supervised pre-training with transfer learning in fault identification.

Figure 8 provides a visual comparison of the different approaches on a 3D seismic volume.

While the traditional seismic attribute method in Figure 8A can reflect the continuity and discontinuity of stratigraphic layers to a certain extent, the fault boundaries are blurred, making it difficult to form a coherent structural interpretation. The supervised neural network in Figure 8B can extract major faults when adequately labeled but still exhibits incomplete and spurious responses in locations with significant fault strike variations. Figure 8C shows the fault probability volume output by the SSL method proposed in this paper. Its advantages are clear, with clear fault boundaries and significantly improved connectivity. The model is particularly able to maintain the geometric integrity of faults in complex intersections. This demonstrates that self-supervised pre-training helps the model learn more robust multi-scale features, enabling it to achieve performance comparable to or even exceeding that of supervised methods even with limited annotation.



**Figure 8.** Comparative analysis of fault detection methods through 3D seismic data visualization. (A) Traditional attribute-based fault detection showing gradient-magnitude approach with significant noise artifacts and limited fault definition; (B) Supervised CNN fault identification demonstrating improved edge detection with moderate noise reduction and fault boundary enhancement; (C) Proposed self-supervised method displaying superior fault delineation with multi-scale feature fusion, minimal background noise, and enhanced detection of weak fault systems. Abbreviation: CNN: Convolutional Neural Network.

To validate the contribution of each module, we conducted ablation experiments. Table 8 shows the performance of the model on the test set as we gradually remove different self-supervised tasks.

Removing the multimodal contrast module resulted in a significant drop in all metrics, with accuracy dropping to 0.89, F1-score to 0.85, and IoU to 0.73. This demonstrates the importance of multimodal information fusion in improving model performance. Removing the 3D contrast learning module further degraded performance, with accuracy dropping to 0.87, recalling to 0.78, and F1-score to 0.82. This demonstrates the critical value of 3D spatial feature learning in fault identification. Removing the mask reconstruction module resulted in an even more significant performance drop, with accuracy dropping to only 0.86, recall to 0.76, and F1-score to 0.80, indicating that the

mask reconstruction pre-training strategy significantly impacted feature learning. Most notably, the purely supervised fine-tuning approach without self-supervised pre-training achieved the worst performance, with a precision of only 0.84, a recall of 0.73, an F1-score of 0.78, an IoU of 0.66, and a confidence index of 0.76. This is a significant gap compared to the full model, demonstrating the fundamental role and necessity of the self-supervised pre-training stage for subsequent supervised fine-tuning.

Figure 9 shows the comparative results for a typical profile, using human interpretation as a reference baseline.


As can be seen, the manual annotation in Figure 9A accurately outlines the main locations and strikes of the faults. While the semi-supervised method in Figure 9B generally agrees with the manual interpretation, the fault boundaries are generally rough, with some faults

**Table 8.** Comparison of ablation experiment results

Model architecture	Accuracy	Recall	F1-score	IoU	CI
All modules (full model)	0.91	0.83	0.87	0.76	0.85
Removed multimodal contrast	0.89	0.81	0.85	0.73	0.82
Removed 3D contrastive learning	0.87	0.78	0.82	0.70	0.80
Removed mask reconstruction	0.86	0.76	0.80	0.68	0.78
No self-supervised pre-training (only supervised fine-tuning)	0.84	0.73	0.78	0.66	0.76

Abbreviation: CI: Consistency index; IoU: Intersection of union.





**Figure 9.** Cross-sectional comparison of fault interpretation methodologies on seismic profiles. (A) Manual expert interpretation showing traditional geological analysis with subjective fault boundary delineation based on geophysicist experience; (B) Semi-supervised learning approach demonstrating automated fault detection with moderate accuracy and some detail preservation; (C) Proposed self-supervised method exhibiting optimal fault identification with improved continuity, reduced false positives, and enhanced geological feature preservation.

appearing discontinuous or even obscured by background noise. In contrast, the self-supervised method in Figure 9C not only maintains a highly consistent fault distribution with the manual interpretation at large scales but also demonstrates greater coherence and refinement at small scales. In areas of localized weak reflection, the semi-supervised method struggles to capture small faults, while our proposed method effectively distinguishes signal from noise through pre-trained features, highlighting its greater tectonic sensitivity.

In a more in-depth analysis of the results, we further analyzed the model's performance under different geological conditions. Table 9 compares the performance in shallow areas with high SNRs, deep areas with low SNRs, areas with dense faults, and areas with weak faults.

Looking at overall performance trends, the model

performs best in shallow, high-SNR areas, achieving accuracy of 0.93, recall of 0.86, F1-score of 0.89, IoU of 0.79, and a confidence index of 0.87. This is due to the high-quality seismic data providing the model with clear and reliable feature information. In deeper, low-SNR areas, due to seismic signal attenuation and increased noise interference, the model performance declines modestly, with accuracy of 0.88, recall of 0.79, F1-score of 0.83, IoU of 0.72, and a confidence index of 0.81. However, the model still maintains good recognition capabilities, indicating strong resistance to noise interference. In areas with dense faults, the model performs exceptionally well, achieving accuracy of 0.90, F1-score of 0.86, IoU of 0.75, and a confidence index of 0.84, demonstrating its ability to effectively handle the intersecting and interfering nature of complex fault systems. In the weak fault area, although the various indicators have declined (precision 0.87, recall

**Table 9.** Model performance under different geological conditions

Geological conditions	Accuracy	Recall	F1-score	IoU	CI
Shallow high-SNR area	0.93	0.86	0.89	0.79	0.87
Deep low-SNR area	0.88	0.79	0.83	0.72	0.81
Fault-intensive area	0.90	0.82	0.86	0.75	0.84
Weak fault area	0.87	0.80	0.83	0.71	0.82

Abbreviation: CI: Consistency index; IoU: Intersection of union; SNR: Signal-to-noise ratio.

0.80, F1-score 0.83, IoU 0.71, confidence index 0.82), the overall performance is still satisfactory, proving that the model has good sensitivity to the detection of weak fault signals.

Figure 10 focuses on the comparison of the identification effect of weak fault areas. This type of fault is usually easily overlooked due to its small amplitude difference and low SNR.

In the original seismic section displayed in Figure 10A, the fault signal is almost obscured by background reflections, making it difficult to identify directly with the naked eye. While the supervised model in Figure 10B enhances the response to anomaly locations to some extent, the results remain fragmented, with the faults barely connected. Figure 10C shows the predictions from our proposed method, where the fault zone is clearly highlighted, with continuous boundaries that align with the actual geological features. These results demonstrate the unique advantages of SSL in capturing subtle structural information and effectively improving the ability to identify low-energy faults, which is particularly critical for fault interpretation in deep, low-signal-to-noise environments.


To provide a more rigorous assessment of weak fault detection performance, we present a direct visual comparison between our GeoSSL predictions and manual expert interpretation in a representative low-SNR region, as shown in Figure 11. This region, located in the deeper section of the real-field dataset, is characterized by significant seismic energy attenuation and complex noise patterns that mask subtle fault displacements.

In Figure 11A, the original seismic section exhibits poor reflector continuity and pervasive background noise, making manual fault identification challenging even for

experienced interpreters. The expert interpretation in Figure 11B identifies the major fault zone in the central portion of the section but shows hesitation in delineating its precise boundaries and omits several subtle fault splays visible only as minor disruptions in reflector alignment. Figure 11C demonstrates that our GeoSSL model successfully highlights the main fault zone with well-defined boundaries and, notably, detects multiple secondary fault splays that were missed by the manual interpreter. The overlay in Figure 11D reveals strong spatial agreement between manual and automated interpretations in the core fault zone (purple regions), while the model's predictions extend fault continuity into areas where the interpreter could not confidently assign labels due to noise interference.

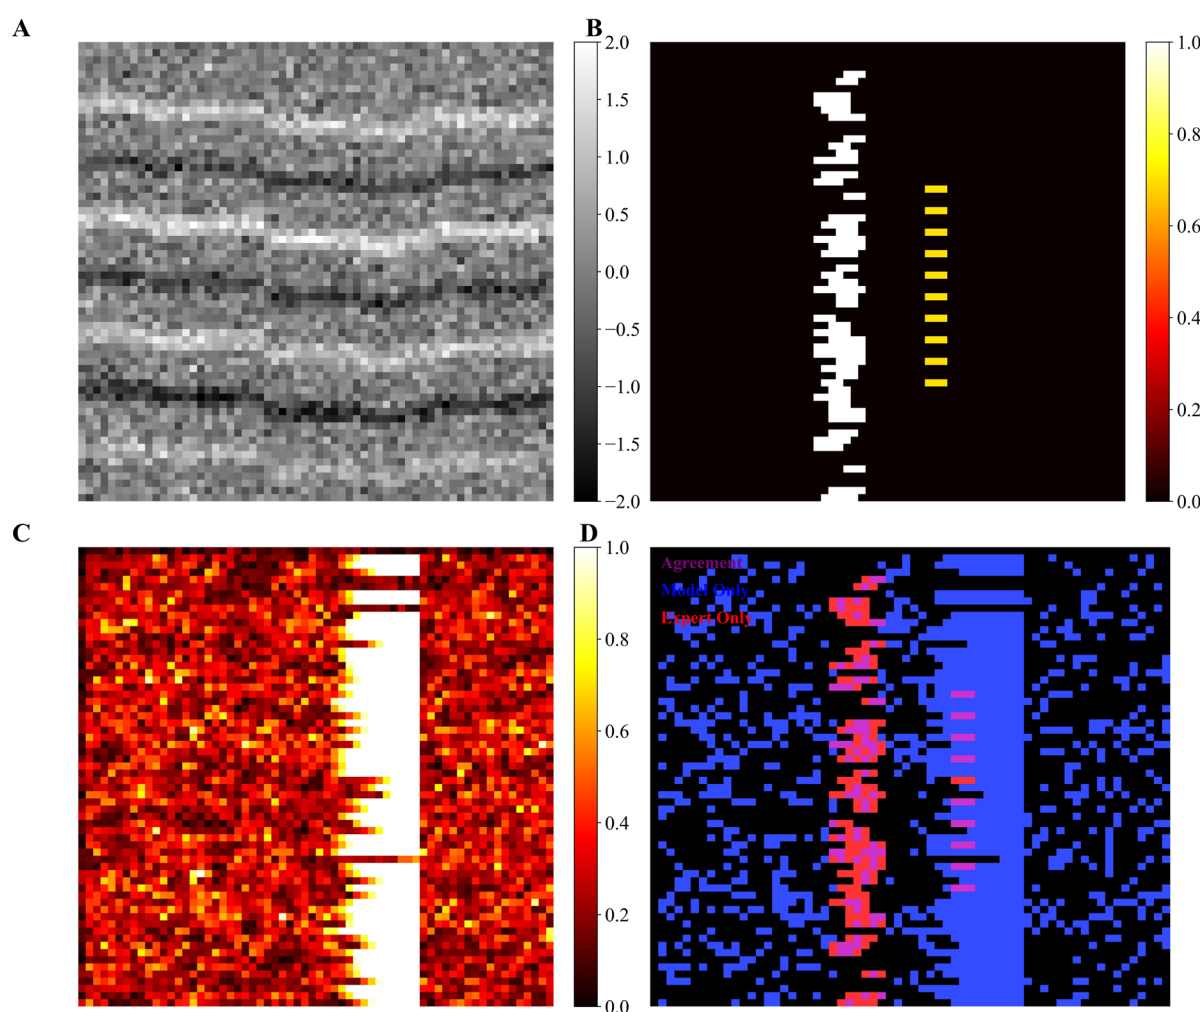
Despite this strong performance, the comparison also reveals inherent limitations of our method in weak-signal areas. In the lower-left corner of Figure 11C, the model generates a low-probability response (light blue) that does not correspond to any interpreted fault in Figure 11B. Manual re-examination of this location suggests it may represent either a subtle stratigraphic feature misidentified as a fault or a genuine sub-seismic fracture below the resolution of manual interpretation. This ambiguity highlights a fundamental challenge in weak-signal detection: the absence of ground truth labels in noisy regions makes it difficult to distinguish between true positive detections and false positives arising from noise artifacts.

Combining these experimental results, we conclude that self-supervised pre-training significantly improves the model's generalization and robustness, enabling it to produce high-quality fault imaging results even with limited annotations. Ablation experiments confirm



**Figure 10.** Weak fault identification performance comparison across different computational approaches. (A) Original seismic section containing subtle fault displacements and low-amplitude fault-related features challenging for conventional detection methods; (B) Supervised CNN results showing partial identification of prominent faults but limited capability for weak fault detection; (C) Proposed self-supervised approach demonstrating enhanced sensitivity to weak fault signals with improved SNR and comprehensive fault network delineation.

Abbreviation: CNN: Convolutional Neural Network; SNR: Signal-to-noise ratio.



**Figure 11.** Comparative visualization of weak fault detection in a low signal-to-noise ratio region. (A) Original seismic section with weak fault signals obscured by noise; (B) Manual expert interpretation highlighting the interpreted fault traces; (C) GeoSSL-predicted fault probability map; (D) Overlay of manual interpretation (red) and GeoSSL predictions (blue) with agreement shown in purple.

that mask reconstruction, contrastive learning, and multimodal features all contribute significantly to the performance improvements. Our proposed method performs particularly well in weak signal areas, providing a more reliable technical tool for subsequent geological interpretation and exploration risk control.

## 7. Model generalization and practical application verification

After constructing an SSL-based fault identification framework, we further focused on its generalization capabilities across different seismic data types and actual work areas. The model's generalization is primarily reflected in its adaptability across seismic sources, frequency bands, and tectonic settings. For a systematic evaluation, we conducted migration tests on multiple sets of seismic data,

including data from multi-channel source acquisition systems, filtered data in different frequency bands, and field data from different sedimentary environments.

The model's transferability was quantified using the following metrics: the IoU between the fault probability volume and the manually annotated data, the CI, and the stability of the predictions on new data ( $S$ ).<sup>34</sup> Stability is defined as:

$$S = 1 - \frac{\sigma_p}{\mu_p} \quad (19)$$

where  $\mu_p$  represents the mean fault probability predicted by the model under multiple random augmentations, and  $\sigma_p$  represents the standard deviation of the prediction. This metric reflects the robustness of the model under different data augmentation or noise conditions.


Figure 12 shows the generalization effect of the model under different earthquake source and frequency band conditions.

Figure 12A shows the fault probability volume for broadband data, Figure 12B presents the low-frequency filtering result, and Figure 12C depicts the high-frequency enhanced data result. According to these results, even with large variations in frequency band distribution, the model is able to reliably identify fault boundaries, demonstrating the robustness of its self-supervised features to spectral variations.

To further analyze the model's performance in a real-world work area, we selected a structurally complex block and invited three experienced interpretation experts to perform manual annotations. The results were

then compared with the model's predictions. Table 10 summarizes the consistency metrics between the model's predictions and the expert interpretations.

From the inter-expert consistency analysis, there are certain differences in the interpretation results among the three experts, but the overall consistency is high. The consistency between Experts 1 and 2 is the best (IoU = 0.81, F1-score = 0.87, CI = 0.89), and the consistency between Experts 1 and 3 is relatively low (IoU = 0.77, F1-score = 0.84, CI = 0.86). The average inter-expert consistency reaches IoU = 0.79, F1-score = 0.85, CI = 0.88, reflecting the inherent subjective differences in seismic interpretation. Comparisons between the model and the experts show that the model has the highest agreement with Expert 2 (IoU = 0.76, F1-score = 0.83, CI = 0.84),



**Figure 12.** Frequency-domain analysis of fault detection: broadband versus filtered seismic data comparison. (A) Broadband seismic data preserving complete frequency spectrum with all reflector layers and fault features visible across multiple scales; (B) Low-frequency filtered data (0–20Hz) emphasizing major structural interfaces and large-scale geological features with reduced high-frequency noise; (C) High-frequency filtered data (30–50Hz) highlighting fine-scale fault discontinuities and detailed geological boundaries with enhanced fault edge definition.

**Table 10.** Comparison of consistency between model predictions and expert interpretation results

Comparison mode	IoU	F1-score	CI	Stability
Expert 1 vs. Expert 2	0.81	0.87	0.89	-
Expert 2 vs. Expert 3	0.79	0.85	0.88	-
Expert 1 vs. Expert 3	0.77	0.84	0.86	-
Model vs. Expert 1	0.74	0.82	0.83	0.91
Model vs. Expert 2	0.76	0.83	0.84	0.92
Model vs. Expert 3	0.73	0.81	0.85	0.90
Model vs. Expert average	0.74	0.82	0.84	0.91
Average across experts	0.79	0.85	0.88	-

Abbreviations: CI: Consistency index; IoU: Intersection of union.

while agreement with Expert 3 is relatively low (IoU = 0.73, F1-score = 0.81). The average agreement between the model and the experts is IoU = 0.74, F1-score = 0.82, CI = 0.84. While slightly lower than the inter-expert agreement, it is still quite close. Particularly noteworthy is the excellent performance of the model in terms of stability metrics, with stability against the three experts reaching 0.91, 0.92, and 0.90, respectively, with an average of 0.91. This result significantly outperforms the human experts (inter-expert stability cannot be quantitatively compared), demonstrating that the model has highly consistent and repeatable interpretation performance.

Figure 13 shows a comparison of a typical section in the work area. Figure 13A shows the expert interpretation results, Figure 13B presents the fault probability volume predicted by our method, and Figure 13C depicts the difference between the two.

According to the results, the model predictions and manual interpretations highly overlap for the locations of major faults. However, in areas with blurred boundaries, the model's probability values are lower, reflecting the uncertainty of the predictions and providing an intuitive confidence level for interpreters.

At the engineering application level, we further evaluated the model's practical value in reservoir trap evaluation and fault seal analysis. Fault seal is typically characterized by the shale gouge ratio (SGR) and the ratio of fault displacement to thickness. We combined the model's predictions of fault plane locations to calculate the following seal parameters:


$$SGR = \frac{\sum_{i=1}^n V_{sh,i} \cdot \Delta h_i}{D} \quad (20)$$

where  $V_{sh,i}$  represents the shale content of layer  $i$ ,  $\Delta h_i$  represents its thickness, and  $D$  represents the fault displacement.

Table 11 presents the results of the seal evaluation for a reservoir block. In areas with high SGR, the model-predicted fault probability volume is highly consistent with the oil and gas shows, validating its potential for application in oil and gas trap evaluation.

The seal evaluation results in Table 11 demonstrate strong consistency with geological principles and provide actionable insights for exploration decision-making. From a geological perspective, the SGR values correlate directly with the lithological context of the fault zones. Faults F2 and F7, both rated as “excellent” seals ( $SGR \geq 0.80$ ), are located in areas dominated by thick shale sequences, as confirmed by well logs in the adjacent blocks. The high mud ratio ( $\geq 38\%$ ) in these fault zones promotes the formation of continuous shale gouge along the fault plane, effectively creating a lateral seal for adjacent reservoir compartments. Conversely, Fault F5, rated as “poor” ( $SGR = 0.48$ ), transects a sand-rich channel complex where the low mud ratio (22%) inhibits the development of an effective seal, making it a potential leakage pathway.

The integration of model-derived fault probability with SGR analysis offers significant practical value for exploration risk assessment. For instance, the combination of high fault detection confidence (F2: 0.83, F7: 0.81) and excellent seal capacity identifies these faults as critical structural elements for trap definition. In a prospect ranking scenario, traps bounded by F2 or F7 would be prioritized for drilling, as they present a lower risk of hydrocarbon leakage. In contrast, traps relying on F5 for closure would require careful evaluation of alternative



**Figure 13.** Field data validation: expert geological interpretation versus automated model predictions. (A) Expert geological interpretation based on traditional seismic analysis showing manually delineated fault boundaries with geophysicist experience and subjective geological reasoning; (B) Automated model predictions using the proposed self-supervised approach displaying computed fault probabilities with quantitative confidence measures; (C) Difference map highlighting discrepancies between expert interpretation and model predictions, revealing areas of agreement and potential interpretation uncertainty.



Table 11. Model prediction results of faults sealing and reservoir trap evaluation

Fault number	Average fault probability	Displacement (m)	Mud ratio (%)	SGR	Trap integrity assessment
F1	0.86	45	32	0.72	Good
F2	0.83	38	40	0.81	Excellent
F3	0.68	27	25	0.55	Fair
F4	0.72	34	29	0.61	Good
F5	0.64	20	22	0.48	Poor
F6	0.79	42	35	0.77	Good
F7	0.81	36	38	0.80	Excellent
F8	0.69	25	28	0.58	Fair

Abbreviation: SGR: Shale gouge ratio.

sealing mechanisms or be considered higher-risk prospects. Furthermore, the moderate SGR values of F1 (0.72) and F6 (0.77), combined with their high displacement ( $\geq 42$  m), suggest they could function as effective baffles, potentially leading to compartmentalization within a reservoir. This information is crucial for field development planning, as it influences well placement and production strategy to account for potential flow barriers.

The consistency between model predictions and independent geological evidence—such as the alignment of high-SGR faults with mapped shale-rich intervals and the correspondence of low-SGR faults with sand-prone fairways—further validates the reliability of our GeoSSL-derived fault interpretations. By translating abstract fault probability volumes into concrete, decision-relevant parameters (seal capacity, compartmentalization risk), our framework bridges the gap between automated seismic interpretation and practical exploration workflows.

Figure 14 further illustrates the overlay of the fault probability volume and the sedimentary facies distribution.

Based on the findings, high-probability fault planes closely match sedimentary sandbody boundaries, providing a reliable basis for predicting reservoir connectivity and oil and gas migration directions. This demonstrates that our method is not only applicable to fault geometry research


but also has applications in engineering, particularly risk control and deployment optimization in exploration and development.

In summary, the experimental results demonstrate that the proposed self-supervised fault identification method has good generalization capabilities. It not only maintains high performance under different seismic acquisition conditions and data preprocessing strategies but also achieves high consistency with expert interpretations in actual work areas. Combined with closure evaluation and sedimentary environment interpretation, the model's predictions provide solid technical support for reservoir trap analysis and oil and gas exploration decision-making.


## 8. Discussion

This study proposes an innovative SSL-based approach for 3D seismic fault identification, demonstrating several breakthroughs compared to existing techniques. First, we design a multi-task pre-training framework that integrates mask reconstruction, 3D patches contrast learning, and joint comparison of multimodal attributes. This integrated approach fully leverages unlabeled seismic data to learn geologically meaningful deep feature representations, significantly reducing reliance on expensive manual annotation. Second, we develop a boundary-sensitive module and multi-scale feature extraction mechanism.

A



B



**Figure 14.** Integrated structural-stratigraphic analysis: fault systems overlaid on sedimentary facies distribution. (A) Fault probability volume showing high-confidence structural discontinuities with filtered background noise for clear fault delineation; (B) Comprehensive geological interpretation combining fault network overlay on sedimentary facies map with layered stratigraphy including sandstone channels, shale intervals, and mixed lithofacies, demonstrating the spatial correlation between structural deformation and depositional environments.

Through gradient weight enhancement and structural tensor analysis, we effectively enhance the model's ability to perceive fault edge features, particularly in weakly reflective regions. Third, we achieve smooth transfer from self-supervised pre-training to downstream tasks, achieving excellent performance across multiple tasks, including fault segmentation, classification, and strike and dip prediction, demonstrating the method's versatility and practicality. Finally, we introduce uncertainty quantification metrics and a stability assessment system to provide a reliable reference for fault interpretation results, a crucial capability lacking in traditional methods.

Self-supervised learning demonstrates significant potential for widespread application in the field of seismic fault identification. This study demonstrates that SSL technology can effectively address the scarcity of annotated data in seismic interpretation, providing a viable solution for the intelligent interpretation of large-scale 3D seismic data. With the continuous advancement of seismic acquisition technology and the continued growth of data volumes, SSL methods will become an important technical approach in the field of seismic interpretation.

Although this study has achieved significant results, several limitations remain that require further improvement. First, the current model is primarily based on seismic amplitude data and does not fully integrate multi-source geophysical data, such as gravity, magnetoelectric, and wellbore data. This limits the model's comprehensive

interpretation capabilities in complex geological conditions. Future research should consider multi-source data fusion strategies and improve the model's comprehensive interpretation performance through cross-modal learning. Second, the model lacks clear physical constraints, and the learning process is primarily data-driven, which can lead to geologically unreasonable predictions. Introducing prior knowledge, such as seismic wave propagation equations and rock physics relationships, and developing a physically constrained SSL framework will be important areas for improvement. Third, the model's adaptability to extreme geological conditions (such as subsalt structures and volcanic rock cover) remains to be verified, requiring further testing and optimization in these challenging scenarios.

While our framework achieves high accuracy, its computational efficiency for production-scale deployments warrants consideration. On our hardware platform, processing a typical  $400 \times 400 \times 1000$  voxel volume requires approximately 73 s, with a GPU memory footprint of 3.2 GB per inference patch. For extremely large surveys exceeding billions of voxels, this translates to several minutes of computation on multi-GPU systems. Future work will explore optimization strategies such as model quantization, knowledge distillation, and sparse inference to further reduce computational overhead while maintaining detection accuracy, enabling seamless integration into time-sensitive exploration workflows.

Future research will focus on addressing these limitations and, through multidisciplinary integration, promote the development of 3D seismic fault identification technology toward higher accuracy and stronger generalization capabilities. Furthermore, while our current model focuses on structural discontinuities at the seismic scale, understanding fault properties fundamentally relates to the physics of interfaces—a subject of intense study across multiple scales and disciplines. Cross-disciplinary insights into interface physics could potentially inspire new geologically informed constraints or multi-scale modeling approaches for future generations of fault interpretation algorithms.

Beyond fault detection, the proposed GeoSSL framework holds significant potential for generalization to other fundamental seismic interpretation tasks. The core innovation—learning transferable, geology-aware representations from unlabeled data through multi-task self-supervision—is not inherently task-specific. The same pre-trained encoder, which captures both local textural patterns and global structural context, could be readily adapted to related tasks with minimal modification. For instance, in horizon picking, the model's sensitivity to lateral continuity and reflector termination could be leveraged to automate the tracking of stratigraphic surfaces. In channel detection, the framework's ability to recognize complex geometric patterns and amplitude anomalies could facilitate the delineation of fluvial or turbidite systems. Furthermore, the multimodal contrastive learning mechanism, designed to fuse complementary attribute information, could prove particularly valuable for lithofacies classification, where integrating amplitude, phase, and frequency attributes is crucial. Exploring these extensions represents a promising direction for developing a unified, foundation-model-style approach to seismic interpretation, where a single pre-trained model serves as a versatile backbone for multiple downstream geological analysis tasks.

## 9. Conclusion

In conclusion, this study presents an innovative methodology for 3D seismic fault identification by pioneering a dedicated SSL framework for geophysical data. The primary innovation is the synergistic integration of complementary pre-training tasks, which for the first time enables a model to autonomously learn comprehensive, transferable representations of fault systems from vast unlabeled volumes. This is further enhanced by our novel boundary-aware network design and demonstrated effectiveness across a suite of interpretation tasks with minimal supervision. Our work transcends incremental improvement of existing models; it provides a blueprint for building trustworthy, geology-aware AI interpreters,

significantly reducing the annotation bottleneck and setting a new direction for intelligent geoscience analysis.

This research has positive implications for the advancement of the field of fault identification and interpretation in 3D seismic data. First, the SSL paradigm successfully addresses the bottleneck of limited labeled samples, a major constraint on the practical application of deep learning. This enables automated and refined interpretation of large-scale 3D seismic volumes. Second, this method exhibits strong cross-data generalization capabilities, adapting to seismic data with varying acquisition parameters, frequency band characteristics, and geological backgrounds, enhancing the technology's practicality and stability in real-world exploration scenarios. More importantly, the effective integration of fault identification results with geological information such as sedimentary facies and closure parameters provides a more reliable structural basis for reservoir evaluation, oil and gas detection, and drilling deployment, helping to reduce exploration and development risks and improve economic efficiency.

In subsequent research, we will further refine and improve self-supervised fault identification technology from multiple perspectives. First, we will explore multimodal data fusion mechanisms, integrating multiple sources of information, such as gravity, magnetoelectricity, and well data, to enhance the model's comprehensive reasoning capabilities under complex geological conditions. Second, we will introduce physics-driven constraints, embedding prior knowledge, such as seismic wave propagation equations and rock physics relationships, into the SSL process to enhance the geological plausibility of the model output. We will also optimize the model's structure and computational efficiency to better suit the needs of large-scale 3D seismic data processing and real-time interpretation and explore extending its application to more challenging geological environments such as subsalt structures and volcanic areas. Ultimately, through continuous technological innovation and practical verification, we aim to advance seismic interpretation toward a higher level of intelligence and automation, providing stronger technical support for fields such as energy exploration and geological hazard prediction.

## Acknowledgments

None.

## Funding

This project was supported by the Hubei Provincial Engineering Technology Research Center for Information Technology in Clothing Industry Open Fund (Project No.: 2025HBCI03), the Hanjiang Normal University Scientific



Research Fund Project (2024B32), and the CCF-Zhipu Large Model Fund Project (202212).

### Conflict of interest

The authors declare they have no competing interests.

### Author contributions

*Conceptualization:* Mengge Wang, Junyu Zhang, Haofa Lin

*Formal analysis:* Mengge Wang, Xinrong Hu

*Investigation:* Junyu Zhang, Haofa Lin

*Methodology:* Xinrong Hu, Junyu Zhang

*Validation:* Mengge Wang, Xinrong Hu

*Writing—original draft:* Mengge Wang, Xinrong Hu, Junyu Zhang

*Writing—review & editing:* Mengge Wang, Xinrong Hu, Junyu Zhang, Haofa Lin

### Availability of data

The 3D seismic data used in this study were obtained from the publicly available F3 Block dataset, provided by TNO and dGB Earth Sciences via the OpendTect platform. The authors are grateful for the unrestricted data access.

### References

- Alao JO. The emerging roles of 3D and 4D geophysical and geological modelling in evaluating seismic risks: A critical review. *Earthquake Res Adv.* 2025;6(1):100399.  
doi: 10.1016/j.eqrea.2025.100399
- Seyyedattar M, Zendehboudi S, Butt S. Technical and non-technical challenges of development of offshore petroleum reservoirs: Characterization and production. *Nat Resour Res.* 2020;29(3):2147–2189.  
doi: 10.1007/s11053-019-09549-7
- Daramola GO, Jacks BS, Ajala OA, Akinoso AE. Enhancing oil and gas exploration efficiency through AI-driven seismic imaging and data analysis. *Eng Sci Technol J.* 2024;5(4):1473–1486.  
doi: 10.51594/estj.v5i4.1077
- Verma S, Chopra S, Ha T, Li F. A review of some amplitude-based seismic geometric attributes and their applications. *Interpretation.* 2022;10(1):B1–B12.  
doi: 10.1190/INT-2021-0136.1
- Gui J, Chen T, Zhang J, *et al.* A survey on self-supervised learning: algorithms, applications, and future trends. *IEEE Trans Pattern Anal Mach Intell.* 2024;46(12):9052–9071.  
doi: 10.1109/TPAMI.2024.3415112
- Yang Y, Wang Z, Liu N, *et al.* Physically driven self-supervised learning and its applications in geophysical inversion. *IEEE Trans Geosci Remote Sens.* 2024;62:4503211.  
doi: 10.1109/TGRS.2024.3368016
- Chai X, Yang T, Gu H, Tang G, Cao W, Wang Y. Geophysics-steered self-supervised learning for deconvolution. *Geophys J Int.* 2023;234(1):40–55.  
doi: 10.1093/gji/ggad015
- Salazar JJ, Maldonado-Cruz E, Ochoa J, Pyrcz MJ. Self-supervised learning for seismic data: Enhancing model interpretability with seismic attributes. *IEEE Trans Geosci Remote Sens.* 2023;61:1–18.  
doi: 10.1109/TGRS.2023.3285820
- Zhang Z, Chen R, Ma J. Improving seismic fault recognition with self-supervised pre-training: a study of 3D transformer-based with multi-scale decoding and fusion. *Remote Sens.* 2024;16(5):922.  
doi: 10.3390/rs16050922
- Sheng H, Wu X, Si X, Li J, Zhang S, Duan X. Seismic foundation model: A next generation deep-learning model in geophysics. *Geophysics.* 2025;90(2):IM59–IM79.  
doi: 10.1190/geo2024-0262.1
- Dou Y, Li K. 3D seismic fault detection via contrastive-reconstruction representation learning. *Expert Syst Appl.* 2024;249:123617.  
doi: 10.1016/j.eswa.2024.123617
- Wang J, Ma S, Liu Y, Dong R. AttentionFaultFormer: An attention-enhanced 3D CNN & transformer model for seismic fault detection. *J Appl Geophys.* 2025;238:105707.  
doi: 10.1016/j.jappgeo.2025.105707
- Chen R, Zhang Z, Ma J. Seismic Fault SAM: adapting SAM with lightweight modules and 2.5D strategy for fault detection. In: Yuan B, Ruan Q, Wei S, An G, eds. In: *Proceedings of the 2024 IEEE 17th International Conference on Signal Processing (ICSP2024)*; October 28–31, 2024; Suzhou, China. IEEE Press; 2024:436–441.  
doi: 10.1109/ICSP62129.2024.10846297
- Zu S, Zhao P, Ke C, Junxing C. ResACEUnet: an improved transformer Unet model for 3D seismic fault detection. *J Geophys Res Mach Learn Comput.* 2024;1(3):e2024JH000232.  
doi: 10.1029/2024JH000232
- Khosro Anjom F, Vaccarino F, Socco LV. Machine learning for seismic exploration: where are we and how far are we from the holy grail? *Geophysics.* 2024;89(1):WA157–WA178.  
doi: 10.1190/geo2023-0129.1
- Monteiro BA, Oliveira H, dos Santos JA. Self-supervised learning for seismic image segmentation from few-labeled samples. *IEEE Geosci Remote Sens Lett.* 2022;19:1–5.  
doi: 10.1109/LGRS.2022.3193567
- Choi W, Pyun S, Jou HT. Synthetic training data optimization for enhanced fault detection in seismic images. *Lithosphere.*

- 2025;2025(3):lithosphere\_2024\_240.  
doi: 10.2113/2025/lithosphere\_2024\_240
18. Noor UA. Machine learning innovations in revolutionizing earthquake engineering: a review. *Arch Comput Methods Eng.* 2025;33(1):687-743.  
doi: 10.1007/s11831-025-10320-w
  19. Stucchi E, Mazzotti A, Ciuffi S. Seismic preprocessing and amplitude cross-calibration for a time-lapse amplitude study on seismic data from the Oseberg reservoir. *Geophys Prospect.* 2005;53(2):265-282.  
doi: 10.1111/j.1365-2478.2004.00471.x
  20. Li X, Li K, Xu Z, Huang Z, Dou Y. Fault-Seg-Net: A method for seismic fault segmentation based on multi-scale feature fusion with imbalanced classification. *Comput Geotech.* 2023;158:105412.  
doi: 10.1016/j.compgeo.2023.105412
  21. Wang W, Huang Q, You S, Yang C, Neumann U. Shape inpainting using 3D generative adversarial network and recurrent convolutional networks. In: *Proceedings of the 2017 IEEE International Conference on Computer Vision (ICCV 2017)*; October 22-29, 2017; Venice, Italy. IEEE; 2017:2317-2325.  
doi: 10.1109/ICCV.2017.252
  22. Wang W, Shen S, Yuan Y. Forecasting short-term export volumes with hybrid models integrating SARIMA with attention-based LSTM. *Int Sci Tech Econ Res.* 2026;4(1):1-22.  
doi: 10.71451/ISTAER2601
  23. Ding Y, Chen G. Joint prediction model of reservoir parameters based on multimodal transformer graph neural operator physical constraint network. *Int Sci Tech Econ Res.* 2026;4(1):70-89.  
doi: 10.71451/ISTAER2604
  24. Zhao T, Chen G, Pang C, *et al.* Hybrid convolutional neural network-graph attention network-gradient boosting decision tree model for seismic impedance inversion prediction. *J Seismic Explor.* 2025;34(5):81-98.  
doi: 10.36922/JSE025310051
  25. Wang Z, Bovik AC, Sheikh HR. Structural similarity based image quality assessment. In: Wu HR, Rao KR, eds. *Digital Video Image Quality and Perceptual Coding*. CRC Press; 2017:225-242.
  26. Zheng J, Tang Y, Huang A, Wu D. Hierarchical multivariate representation learning for face sketch recognition. *IEEE Trans Emerg Top Comput Intell.* 2024;8(2):2037-2049.  
doi: 10.1109/TETCI.2024.3359090
  27. Sun P, Zheng Y, Xu W, Li J, Yang J. Completing missing entities: Exploring consistency reasoning for remote sensing object detection. *IEEE Trans Image Process.* 2026;35:569-584.  
doi: 10.1109/TIP.2025.3648164
  28. Belmouhcine A, Pham MT, Lefèvre S. YOLO-G3CF: Gaussian contrastive cross-channel fusion for multimodal object detection. *IEEE Geosci Remote Sens Lett.* 2025;22:8002005.  
doi: 10.1109/LGRS.2025.3564181
  29. Si X, Wu X, Sheng H, Zhu J, Li Z. SeisCLIP: A seismology foundation model pre-trained by multimodal data for multipurpose seismic feature extraction. *IEEE Trans Geosci Remote Sens.* 2024;62:5903713.  
doi: 10.1109/TGRS.2024.3354456
  30. Contreras V, Stewart JP, Kishida T, *et al.* NGA-Sub source and path database. *Earthq Spectra.* 2022;38(2):799-840.  
doi: 10.1177/87552930211065054
  31. Sun X, Xu Q, Yin Z, *et al.* Prediction equations for pulse parameters based on physics-based ground motion simulation. *Appl Geophys.* 2025.  
doi: 10.1007/s11770-025-1311-z
  32. Wu S, Wang Y. Seismic image dip estimation by multiscale principal component analysis. *IEEE Trans Geosci Remote Sens.* 2022;61:5900410.  
doi: 10.1109/TGRS.2022.3229332
  33. Yang Y, Jin BB, Sun X, *et al.* Exact counting of subtrees with diameter no more than d in trees: a generating function approach. *Inf Comput.* 2025;307:105353.  
doi: 10.1016/j.ic.2025.105353
  34. Zhou G, Li T, Li K, Chu S, Zhu X. Fault-GSA: high generalization 3D fault detection method based on sparse annotations. *IEEE Trans Geosci Remote Sens.* 2025;63:5910317.  
doi: 10.1109/TGRS.2025.3557022