

ARTICLE

Deep learning for three-dimensional seismic fault prediction using convolutional neural networks

Yasir Bashir^{1*}, Dilek Yüksel¹, Begüm Akin¹, Derin Gezer¹, Muhsan Ehsan², Muhammad Khan³, and Syed Haroon Ali⁴¹Department of Geophysical Engineering, Faculty of Mines, İstanbul Technical University, İstanbul, Turkey²Department of Earth and Environmental Sciences, Bahria School of Engineering and Applied Sciences, Bahria University, Islamabad, Pakistan³Saudi Aramco, Dhahran, Saudi Arabia⁴Department of Earth Sciences, Faculty of Science, University of Sargodha, Sargodha, Punjab, Pakistan

Abstract

Understanding seismic faults is essential for generating prospects, modeling reservoirs, and assessing carbon dioxide storage. Identifying faults in complex tectonic regimes presents significant challenges, especially in areas that have undergone multiple phases of tectonic activity. Even with advances in structural seismic attributes and machine learning, interpreters often rely on manual methods to examine complex fault systems. This work introduces a method for predicting three-dimensional seismic faults using convolutional neural networks (CNNs), effectively overcoming the constraints of conventional interpretation techniques. The research uses CNNs to demonstrate the effectiveness of seismic attributes in training models that identify faults with high accuracy and consistency. This approach, unlike manual interpretation, reduces time, costs, and subjective errors by leveraging automated learning techniques, thereby improving reproducibility and efficiency while reducing interpreter bias. The research highlights the growing importance of solid computational tools in geophysics, especially as seismic datasets become more complex and wider. The approach significantly enhances confidence in artificial intelligence-assisted geological analysis by validating its performance with real-world data. The validation accuracy rises from 0.936 to 0.9436 across configurations, while the validation loss increases from 0.5635 to 0.9941 across diverse patches of the trained model.

Keywords: Deep learning; Convolutional neural networks; Fault likelihood; Fault cube; Fault prediction

***Corresponding author:**Yasir Bashir
(ybashir@itu.edu.tr)

Citation: Bashir Y, Yüksel D, Akin B, *et al.* Deep learning for three-dimensional seismic fault prediction using convolutional neural networks. *J Seismic Explor.* 2026;35(3):026180079.
doi: 10.36922/JSE026180079

Received: April 30, 2026**Revised:** May 21, 2026**Accepted:** May 28, 2026**Published online:** June 18, 2026

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

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

1. Introduction

In geophysics, the entire process of seismic structural interpretation depends significantly on the expertise and specialized knowledge of the interpreters involved. This dependence increases the chances of human error and variability, particularly in

complex geological settings.¹ Recently, artificial intelligence and machine learning (ML) have become significant tools in geophysical exploration, capable of replicating human expertise and enhancing their performance through training on a broader range of datasets.^{2–6} Geophysical explorations require seismic fault detection. The problem is that this method is complicated and time-consuming. Conventional interpretation methods rely on human interpreters' knowledge and judgment, which increases variability and error risk, especially when handling large datasets or complex geological features. Manual methods are becoming insufficient as seismic data volumes and complexity increase.⁶ The issue at present limits scalability and the detection of shortcomings, emphasizing the necessity for more effective, objective, and automated solutions, particularly those based on ML.

Existing geophysical methods lack tuning and automation to handle the variety and complexity. Without these answers, human error and fault identification across geological contexts remain issues. The scenarios demonstrate the importance of detailed fault mapping for hydrocarbon exploration and reservoir characterization to reduce risk and boost production. Recently, Wu *et al.*⁷ proposed an automatic defect interpretation method based on a convolutional neural network (CNN). A 7-layer CNN was used to estimate fault orientations, notably dips and strikes, from small patches of entire seismic pictures. The estimated orientations yielded anisotropic Gaussian functions that extend in the projected fault directions.⁷ By arranging locally oriented Gaussian functions, a fault probability image is generated that enhances the clarity and continuity of fault features. Despite being trained on synthetic seismic data, the CNN shows the capability to accurately estimate fault orientations from actual seismic images. This approach offers greater accuracy and clarity than conventional fault-attribute techniques. Guo *et al.*⁸ used CNN to interpret faults and horizons in a coalfield case study. Structurally modeled datasets achieved high accuracy, eliminating the need for manual labeling.⁸ Both Wu *et al.*⁷ and Guo *et al.*⁸ highlighted the potential of CNNs for seismic structural interpretation; however, there are several methodological gaps that this study explicitly addresses. Initially, Wu *et al.*⁷ focused on estimating fault orientation through patch classification. However, their method did not yield spatially coherent fault-probability volumes that could be directly interpreted geologically, nor did it address the simultaneous prediction of faults and horizons. Guo *et al.*⁸ achieved joint fault-horizon prediction; however, their approach to label generation via structural modeling, although innovative, was tested only on relatively simple geological configurations. Furthermore, it did not explicitly address the significant class imbalance in

fault detection tasks, where fault pixels typically constitute only 5–10% of the total labeled volume. Additionally, neither study conducted a systematic validation of label quality nor reported metrics for inter-annotator agreement regarding the training data. The present study advances the current understanding in several important areas: (i) A patch-based training strategy that includes targeted oversampling of fault-bearing patches to address class imbalance; (ii) the implementation of a weighted focal loss function to further tackle the frequency imbalance between fault and non-fault classes during training; (iii) a thorough binary mask quality assurance protocol that involves independent verification and spatial consistency checks; and (iv) validation on both synthetic and real coalfield seismic data, utilizing quantitative performance metrics such as precision, recall, and the F1-score, which offers a more detailed evaluation of model generalization than previously established.

Research indicates that seismic facies analysis methodologies have evolved, highlighting a transition from manual interpretation to ML-based techniques. The study describes the difficulties associated with identifying seismic features, particularly for new interpreters, and introduces a method that integrates UVQ and beam forming seismic.⁹ This approach employs computational weights to pinpoint the key seismic features, thereby enhancing objectivity and reproducibility. The article recommends using spectral decomposition in conjunction with ML for interpretation and initial validation. Spectral decomposition helps interpreters identify lithological and stratigraphic changes by providing frequency-based information from seismic volumes instead of well log data. A detailed Gulf of Guinea case study shows how attribute combinations affect facies classifications. It shows how the selection of input data affects unsupervised techniques. Zhao *et al.*¹⁰ perform a comparative analysis of six commonly used classification techniques for seismic facies recognition, tackling the escalating challenges posed by increasingly large three-dimensional (3D) seismic datasets and the variety of available seismic attributes.¹⁰ Since manually reading seismic lines and time slices is difficult, this research evaluates supervised and unsupervised ML systems. This study integrates a semblance-based fault-likelihood assessment with the widely used ant-tracking method to reduce the inefficiencies and subjectivity of manual fault interpretation. This integration creates high-resolution discontinuity volumes that enable interpreters to automatically extract defects with less input.^{11,12} The study's use of directed semblance, based on Dave Hale's fault-oriented semblance algorithm, improves alignment with fault geometry over coherency-based methods. The approach was tested on a fractured reservoir in a Malaysian

basin, demonstrating its reliability in noisy, gas-related signal-attenuated settings. The method improves fault mapping accuracy and efficiency and shows potential for detecting fractures, stratigraphic discontinuities, and gas chimneys.

Recent studies have shown the increasing efficacy of CNNs and sophisticated deep learning frameworks in complex medical imaging applications.^{13,14} Methods including privacy-preserving federated learning paired with extended reality for analyzing pediatric cerebral palsy, lightweight deep learning alongside virtual reality visualization for offline tumor segmentation, spiking CNNs employing spike-timing-dependent plasticity for glioma segmentation, and advanced U-Net and ConvNeXt architectures incorporating attention mechanisms and data augmentation for breast tumor segmentation have demonstrated notable advancements in feature extraction, segmentation accuracy, computational efficiency, and the intelligent interpretation of high-dimensional datasets.^{15,16} These advancements underscore the extensive potential of CNN-based models for automated pattern recognition and semantic interpretation, which can also be tailored for seismic data analysis, subsurface characterization, and geophysical feature detection.

A hybrid method is developed that employs variable-mode decomposition and support vector machines for the detection of small-scale faults in coal mining, where vertical displacements range from 2 to 5 m. Their approach effectively minimized seismic noise, enhancing fault detection under challenging conditions.^{17,18} We also discussed the growing interest in semi-supervised and unsupervised methods, which reduce labeling while maintaining model efficacy. Recent advancements in seismic interpretation leverage sophisticated hybrid architectures, such as combining CEEMDAN multiscale fuzzy entropy with NSGA-III-optimized CNNs for event recognition, or integrating graph attention networks with gradient-boosted decision trees for enhanced impedance inversion prediction.^{19,20} A deep convolutional neural network model for automatic fault detection was introduced by An *et al.*,²¹ framed as an image segmentation problem.²¹ The model achieved performance comparable to that of humans across multiple datasets while significantly decreasing processing time.

1.1. The study area

Structurally complex, the F3 Block in the northern Dutch offshore is impacted by Devonian to Paleogene tectonic periods. Ter Borgh *et al.*²² and Terranubis' 2020 high-resolution 3D seismic dataset show that the F3 area is important for regional fault reactivation, salt tectonics, and

basin development. The study area is between the Elbow Spit Platform and the Step Graben. Terranubis seismic data confirms pre-Zechstein deformation, fault reactivation, and salt as a mechanical decoupling horizon. The traits align with regional structural evolution and help explain trap formation, fault sealing, and compartmentalization. The F3 Block captures Ter Borgh *et al.*'s²² regional tectonic evolution, while the Terranubis 2020 seismic dataset provides substantial image support for local and regional fault system correlation.

A unique extensional regime formed N360° (N–S)-trending faults during the Triassic and Early Cretaceous periods, which helped construct the Step and Dutch Central Grabens. These faults have large throws in the F3 Block and play a major role in the distribution of Zechstein salt and in structural deformation. Finally, N070° (WSW–ENE)-trending dextral strike-slip faults, active from the Jurassic to the Paleogene, locally overprint earlier structures and present as subtle flower structures within coherence and attribute displays of the F3 volume, especially in salt-free areas²² (Figure 1).

1.2. Convolutional neural networks

Research indicates a significant transition towards supervised ML methods for seismic interpretation, with CNNs recognized as the most prevalent and efficient technique for fault detection. CNNs exhibit greater precision and resilience in pattern recognition, particularly compared with conventional manual or attribute-driven interpretation techniques. The ability to autonomously extract spatial characteristics from seismic data makes them especially effective at detecting geological discontinuities, including faults. Research consistently shows that CNNs outperform traditional algorithms in terms of precision and generalizability, especially in complex or noisy subsurface settings. Hybrid models that combine CNNs with signal processing techniques or temporal frameworks, such as long short-term memory networks, show greater potential to improve detection sensitivity and maintain fault continuity.

The main goal of this research is to develop a robust, automated system for detecting seismic faults using supervised ML methods. The study focuses on training CNN models that can accurately identify and highlight fault features within seismic volumes. CNNs serve as the primary algorithm because of their established success in intricate pattern recognition challenges and their flexibility in handling geophysical data formats. The integration of these models into seismic interpretation workflows aims to improve interpreters' productivity and accuracy. This is achieved by generating suggestive fault frameworks and




Figure 1. Location of the dataset used for research. (A) Google Earth map of the study area, and (B) Focusing on the structural framework of the northern Dutch offshore. Faults are displayed at the Base Permian level or, for older structures, at the highest affected stratigraphic horizon. The map also shows the depth to the base of the Zechstein Group, or to the next overlying unit where Zechstein is absent, adapted from Ref.²² Abbreviation: 3D: Three-dimensional.

associated confidence levels, which, in turn, reduce manual effort and enhance consistency in subsurface structural analysis.

2. Materials and methods

The method for identifying seismic faults uses ML techniques, with a particular focus on CNNs. The work begins with an overview of the core principles of CNNs, explaining their importance and advantages in seismic data analysis, particularly for segmentation tasks. CNNs are chosen for their ability to independently extract features and recognize complex spatial patterns in seismic data, making them well-suited for detecting fault zones.

A comprehensive overview of the theoretical foundations of CNNs is followed by an in-depth explanation of the specific workflow in this research. The next phase of the workflow involves preprocessing real field seismic data. The process of selecting seismic attributes for fault detection is carefully carried out, ensuring that only the most relevant features are retained for the model. Following this, the model architecture is defined, with the U-Net selected for its demonstrated efficacy in segmentation tasks. The process involves training the CNN model using labeled seismic data patches, followed by evaluation and performance testing.

The workflow is thoughtfully explained, incorporating

patch-based training, model evaluation using accuracy and loss metrics, and visualization techniques for assessing fault detection. This methodology is designed to develop an automated solution for detecting seismic faults that is both effective and efficient. The method relies on ML to improve both the precision and efficiency of seismic interpretation.

2.1. Machine learning approaches for fault detection

Machine learning algorithms can be classified into two main categories: supervised learning, which functions with labeled datasets, and unsupervised learning, which operates on unlabeled datasets. As indicated in Table 1,⁹ supervised learning techniques are especially advantageous when the dataset includes predefined output values or labels. This method enables the model to identify patterns by analyzing the connection between the input data and the associated output labels. Conversely, unsupervised learning is used when the output labels are unavailable, requiring the algorithm to identify hidden patterns or clusters in the data without predefined classifications.

For fault detection in seismic data, we chose supervised ML because it can be applied to datasets with labeled values. Supervised learning provides significant benefits by enabling the model to learn directly from labeled data, resulting in enhanced accuracy and reliability. By leveraging the established geophysical values, supervised

learning enhances the model's ability to accurately identify and predict fault zones, making it a more appropriate choice for this application.

Table 1. Primary categories of machine learning methodologies, along with their corresponding algorithms

Supervised machine learning	Unsupervised machine learning
- Convolutional neural networks - Support vector machines - Random forest - Convolutional recurrent neural networks - Probabilistic neural networks	- Self-organizing maps - K-means clustering - Principal component analysis - Independent component analysis - Generative topographic maps - Unsupervised vector quantizer - Convolutional autoencoder

Supervised ML techniques have demonstrated effectiveness in seismic interpretation tasks by reducing the need for manual intervention and improving the accuracy of structural identification. CNNs have proven significant potential due to their ability to autonomously extract

features and identify spatial patterns within seismic data.²³ CNNs are particularly effective for binary segmentation tasks, such as distinguishing between faulted and non-faulted regions, making them an ideal option for this work, as shown in Figure 2.

A typical CNN architecture comprises three essential components: convolutional layers, pooling layers, and fully connected layers. The convolutional layers play a crucial role in identifying low- and high-level features in the input seismic slices, including edges, discontinuities, and textures, by using learnable filters. Pooling layers subsequently reduce the dimensionality of the extracted features, thereby helping the model generalize and mitigate overfitting (Figure 2A). The fully connected layers at the conclusion of the network aggregate these features into a definitive classification output, signifying the presence or absence of faults.²⁴

2.2. Convolutional neural networks

Convolutional neural networks represent an innovative deep learning approach for detecting seismic faults, offering notable benefits over conventional techniques. The main factor contributing to their success is the computational power of contemporary graphics processing units. These




Figure 2. The working of a convolutional neural network (CNN) and internal process. (A) Generalized scheme of CNN including convolutional, pooling, and fully connected layer for output, and (B) specific flow of the machine learning model explaining the CNN architecture used to estimate fault orientation from an input seismic patch.⁷

Abbreviations: COV: Convolutional layer; FC: Fully connected.

enable CNNs to concentrate on local feature connections while streamlining the overall architecture of neural networks. This feature allows CNNs to handle large seismic datasets with exceptional efficiency, eliminating the requirement for extensive libraries of waveform templates that were once crucial in conventional methods. Avoiding this dependency allows CNNs to reduce computational complexity and achieve robust generalization, enabling the detection of seismic features in waveforms that were not part of the training set.²⁵ Additionally, CNNs excel at extracting significant features directly from raw seismic data (uninterpreted), enabling them to model intricate, non-linear relationships among variables. This feature proves particularly advantageous for analyzing the complex, often subtle patterns encountered in seismic fault-detection tasks, establishing CNNs as a formidable asset in the domain, as shown in Figure 2A.

Additionally, CNNs extend beyond the scope of

seismic fault detection and demonstrate significant adaptability across a range of seismic applications. Their primary strength lies in the capacity to generalize beyond the specific training data, enabling them to manage the inherent variability present in geological structures and seismic signals. This capability holds particular significance in areas with limited labeled data, as CNNs can leverage patterns acquired from one dataset to enhance others, thereby markedly increasing the accuracy of predictions (Figure 3). The capacity to generalize across datasets, manage complex and noisy data, and optimize the computational process establishes them as a transformative technology in the field, serving as a fundamental element for progressing both research and practical applications in seismic interpretation.

2.3. Workflow of research

The workflow for this research is a systematic, structured




Figure 3. The machine learning algorithms have been used for the most common interpretation applications. The three gray circles illustrate the relationship between their dimensions and the fractions they represent in the applications.¹

Abbreviations: 4D: Four-dimensional; CNN: Convolutional neural network; DBSCAN: Density-based spatial clustering of applications with noise; GAN: Generative adversarial network; GTM: Generative topographic mapping; KNN: k-nearest neighbor; MLP: Multilayer perceptron; PCA: Principal component analysis; PNN: Probabilistic neural network; PSVM: Proximal support vector machine; RNN: Recurrent neural network; SOM: Self-organizing map; SVM: Support vector machine; SVR: Support vector regression.

approach that ensures each component of the ML model's development and application is comprehensively examined. The main goal is to use CNNs to detect seismic faults. Each phase is structured to ensure optimal model performance while addressing the complexities of interpreting seismic data, as illustrated in [Figure 4](#).

The process begins with data collection and preparation. The study used real-world open-access seismic datasets. The Dutch F3 offshore block (Terranubis F3-Demo-2020) has a high-resolution 3D seismic volume and a detailed fault system description. The dataset is model-ready. Seismic properties are chosen to find faults, preserving geological features such as reflection terminations for study. The next step is model selection and configuration. The CNN using the U-Net architecture is chosen for its segmentation performance. U-Net's balanced encoder-decoder architecture and skip links retain spatial precision needed for geophysical data interpretation, making it effective for seismic fault detection. To improve model performance and accuracy, the number of epochs and patch sizes were tested. This technique identified the best training configuration, improving fault detection accuracy and reliability.

The results highlight the strengths and drawbacks of the seismic fault identification model. Future proposals include improving model accuracy, integrating more seismic datasets, and modifying model hyperparameters. The following suggestions will guide seismic fault detection and improve the model's ability to handle more complex geological formations. The approach to developing the seismic fault detection ML framework relies on supervised learning and CNNs. These networks can automatically extract characteristics and identify spatial patterns, making them ideal for segmentation and seismic fault detection. The methodology begins with seismic data collection and preprocessing, followed by fault detection and seismic attribute identification. After data preparation, the CNN model is patch-trained. This method solves memory restrictions and improves model generalization. Different epochs and patch sizes were tested to increase training performance and accuracy.

3. Results

The application of the ML model to seismic data is followed by a detailed study of its performance and implications for fault detection. The main goal is to evaluate the CNN model's performance in identifying fault zones in seismic sections using North Sea seismic data. The model is trained and tested on a publicly available 3D seismic dataset using a patch-based approach to optimize memory usage and increase the diversity of training samples. The extensive

training process, model configuration, and assessment measures indicate that CNNs can accurately detect errors in complex, noisy datasets. This section examines how TensorFlow, Keras, and Matplotlib are used to build, train, and visualize the model. Probability maps of the model's fault predictions allow comparisons between projected fault zones and seismic structures.

3.1. Machine learning training and testing

The architecture is a U-Net-based CNN. Due to its balanced encoder-decoder architecture and skip connections that preserve spatial resolution, the U-Net is well-suited for segmentation tasks such as seismic volume analysis. Consequently, this architecture influenced the data management strategy during experimentation. The raw seismic data were evaluated and grouped into a 3D array with inline, crossline, and temporal (or depth) dimensions before training. After extracting and organizing each trace, the amplitude values were normalized to (0, 1). This preprocessing step ensured numerical stability and consistency during model training. Six manually labeled binary masks aided supervised learning. Seismic inline slices 100, 200, 300, 400, 500, and 600 correspond to masks. The original tutorial includes tested and labeled inlines 100–500. This project used inline 600 to improve model learning and evaluation ([Figures 5 and 6](#)).

Each mask was a PNG image with binary pixel values indicating defects. However, these label masks had uneven beginning size and did not match the seismic slice resolution. Using masks directly in training and evaluation is difficult due to dimensional irregularity, which prevents identified defects from aligning with their counterparts in the raw data. This issue was fixed by resizing all masks to 462 pixels high and 951 pixels wide. By preventing interpolation artifacts, nearest-neighbor interpolation preserved mask binary integrity during resizing. After scaling, masks were converted to binary to ensure pixel values represented fault or non-fault circumstances. Standardizing mask dimensions ensured spatial alignment between seismic inlines and labels. If this phase is skipped, the array shapes will differ, preventing the model from accurately learning correlations between seismic features and fault structure. Applying this modification consistently to all six labeled inlines created a consistent and compatible training and validation dataset.

The high dimensionality and complexity of seismic data make direct training of a CNN on complete seismic slices computationally intensive and inefficient. A patch-based approach was employed to enhance the model's training efficiency. This approach entails extracting smaller, square subimages, referred to as patches, from the seismic inline




Figure 4. Workflow of the research: A convolutional neural network (CNN)-based fault detection method using a multi-epoch and training patches approach

Abbreviation: 3D: Three-dimensional.




Figure 5. Seismic inline slices from 100 to 600 at 100-interval steps. Panels (A) through (F) correspond to inline 100, 200, 300, 400, 500, and 600, respectively.




Figure 6. Seismic inline slices with labeled fault masks from 100 to 600 at 100-interval steps. Panels (A) through (F) correspond to inline 100, 200, 300, 400, 500, and 600, respectively.

slices, along with their associated binary label masks. A specialized function for patch extraction was created to accomplish this task. It randomly selects square regions of size 64×64 pixels from both the seismic data and the label mask. Each patch from the seismic image is spatially aligned with its corresponding label patch to maintain

pixel-level correspondence, a crucial requirement for supervised learning (Figure 7).

A threshold condition was used to ensure that extracted patches contained significant features. Only patches containing a sufficient number of fault-indicating pixels




Figure 7. Patch samples along with their associated faults (yellow lines indicate faults)

were retained. Filtering removes irrelevant or empty regions of seismic volume from the training data. Patch extraction continues until a sufficient number of high-quality samples are collected. Array orientation discrepancies are corrected using the function. If the seismic slice and label mask differ in shape due to axis ordering, the function automatically transposes the data to align their dimensions.

For each training session, accuracy and loss metrics for the training and validation datasets were recorded. Table 2 and Figure 8 show the performance evolution plotted and examined. Figure 8A–D shows the model's learning behavior over time by plotting accuracy and loss values for each training period. A confusion matrix was created to evaluate classification performance using predictions from the validation dataset. The confusion matrix lists the true positives, false positives, true negatives, and false negatives for all patches. Normalizing and visualizing the data as a heatmap showed the model's ability to distinguish fault and non-fault classes. Even with the typical class imbalance in seismic fault datasets, background classification was accurate, and fault pixels were detected well (Figure 9). This method reliably assessed the model's generalization to full-scale data and demonstrated its capacity to recognize major and complex geological features (Figures 10 and 11).

Model performance was evaluated in the third dimension by selecting a depth level within the seismic volume. A horizontal time slice at depth index 420 showed a consistent time sample across all inlines and crosslines (Figure 12A). The two-dimensional slice showed subsurface structures at the selected depth as amplitude values. The patch-based prediction algorithm for vertical inlines was modified to use a sliding window in this horizontal plane. During the training phase, a 64×64 -pixel window was used to traverse the slice and provide predictions at the center of each patch. The model systematically scanned the time slice to create a fault probability map (Figure 12B and 12C).

3.2. Machine learning output

The patchification procedure expanded the training dataset and augmented the data, allowing the model to encounter more fault patterns and improve its fault-detection generalization (Figure 7). Epoch counts were tested and compared during training for the U-Net-structured CNN model. By analyzing accuracy and loss measurements across multiple epochs, it is possible to determine when additional training continues to improve performance and to reduce the risk of overfitting. Despite limited processing resources and a short dataset, the model learned from the data. A validated design, adequate filter depths, and regulated training helped the model to recognize and pinpoint geological fault structures with sufficient precision.

Table 2 summarizes U-Net model training and validation performance metrics for epochs 3–15. Training accuracy, validation accuracy, and validation loss are used. The objective was to determine an appropriate number of training epochs to improve generalization while avoiding overfitting. Training accuracy rises from 96.81% in epoch 3 to 99.47% in epoch 15. Similarly, training loss drops from 0.0821 to 0.0124. This shows that the model learned and fit the training data better with each epoch. Validation accuracy is steady and high throughout all epochs, with minimal changes between 93.12% and 94.2%, peaking at 94.2% in epoch 14. The validation loss trend is different. Starts at 0.335 in epoch 3, rises to 0.9141 in epoch 11, and varies between 0.75 and 0.76 thereafter.

Overfitting begins when both training accuracy (A) and validation loss (D) increase. The model starts to remember training data patterns at epoch 8 or 9, rather than generic characteristics. The validation accuracy remains good, but the rising validation loss suggests that the model may be overconfident in its predictions, especially when applied to unseen data, reducing reliability. Epochs 8 to 10 show a good balance, with epoch 10 achieving 99.19% training accuracy, 93.62% validation accuracy, and a validation loss of 0.6878, demonstrating control before overfitting.

Table 2. A comparison of accuracy and loss is presented as a function of epoch number, with validation accuracy and loss also calculated

Epoch no	Accuracy	Loss	Validation accuracy	Validation loss
3	0.9681	0.0821	0.9359	0.3350
4	0.9746	0.0630	0.9327	0.3953
5	0.9861	0.0336	0.9342	0.5120
6	0.9888	0.0269	0.9353	0.5371
7	0.9859	0.0334	0.9326	0.5443
8	0.9917	0.0196	0.9387	0.6434
9	0.9922	0.0184	0.9390	0.6069
10	0.9919	0.0192	0.9362	0.6878
11	0.9928	0.0171	0.9312	0.9141
12	0.9932	0.0161	0.9361	0.5953
13	0.9949	0.0120	0.9397	0.7290S
14	0.9943	0.0136	0.9420	0.7584
15	0.9947	0.0124	0.9347	0.7607

In the early training phases (Figure 8A, epochs 3–5), the model improves training accuracy while maintaining validation accuracy, with a slight divergence after epoch 4. Overfitting begins when validation loss rises while training loss decreases. Although the model has modest generalization, it fails to understand the complexity of fault structures. In epochs 10–12 (Figure 8A), training accuracy saturates above 99%, and loss approaches zero. However, validation accuracy plateaus while validation loss increases, showing a performance disparity. The model improves on the training set, but its generalization performance declines. Due to its high accuracy and good validation performance, epoch 12 (Figure 8A) was chosen for final implementation. In the final phases of training (Figure 8A, epochs 13–15), training accuracy remains high, whereas validation accuracy plateaus or declines. Meanwhile, validation loss is rising, indicating that the model is overconfident in its predictions and that these improvements are not matched by generalization gains. The latter epochs exhibit overfitting, indicating that additional training yields diminishing returns and degrades the model's performance on new seismic data.

The confusion matrix is a tool for assessing a model's classification performance by comparing predicted labels

to actual labels. This provides a summary of the correct and incorrect predictions for each class, offering insights into the model's ability to differentiate between fault and non-fault pixels. The matrix allows for the derivation of key metrics, including true positives, false positives, true negatives, and false negatives, which are essential for evaluating accuracy, precision, and recall, as demonstrated in Figure 9.

The confusion matrices in Figure 9 show that the model achieves 98–99% accuracy in detecting no-fault pixels across all epochs, due to their abundance in the sample. Despite advances, fault identification remains limited, especially in the earlier epochs (3–5), where true-positive rates drop to 19% (Figure 9A–C). In epochs 10–12, defect detection reaches 25%, suggesting model sensitivity improvement after training (Figure 9D–F). In epochs 13–15, fault prediction accuracy stabilizes or decreases, suggesting overfitting. This suggests that while the model improves the dominant class, it does not generalize to fault locations (Figure 9G–I).

To visually evaluate the model, predicted fault probability maps were overlaid on seismic inlines. The results for epochs 5 (Figure 10) and 10 (Figure 11) were reviewed to evaluate fault identification progress during training.

The findings show that epoch 10 forecasts are more advanced and more consistent across all inline portions than epoch 5 forecasts. Both sets of overlays show fault detection in red highlights on the seismic amplitude background, but there are numerous key differences. Epoch 5 shows many expected fault patterns, some broken and others partially aligned with geological discontinuities. The model identifies major fault regions, though it still misses parts, especially on faint or low-contrast seismic lines. At epoch 10, predictions are more consistent and aligned with seismic structures. In complex sections such as lines 300 and 500 (Figure 5C and 5E), faults are more accurately identified. Epoch 10's red overlays align more closely with fault planes and exhibit less noise, improving the model's ability to recognize both significant and subtle fault characteristics. This visual upgrade confirms the quantitative data that epoch 10 showed improved training accuracy and validation performance. It appears the model has increased its capacity to generalize from labeled data, leading to more geologically relevant predictions.

Fault prediction was performed on a time-slice segment at depth index 420, removed from training and without labeled data, to evaluate model generalization (Figure 12A). To assess the model's performance on unseen data, the results from epochs 5 and 10 were compared (Figure 12B and 12C).




Figure 8. Performance metrics for training and validation are illustrated across epochs. (A) Accuracy during training, (B) loss during training, (C) accuracy during validation, and (D) loss during validation.

The prediction overlays show that model performance varied from epochs 5 to 10, especially in a segment of the data that was not used during training and lacked labeled input. At epoch 5, fault forecasts are scattered and less continuous, with multiple discontinuities and isolated false positives (Figure 12B). Many fault structures are underrepresented or partially recognized, suggesting the model has not fully generalized from the training data. The predictions at epoch 10 are clearer, more consistent, and aligned with the seismic imaging fault patterns (Figure 12C). The model covers more fault network locations, especially curved or branched ones, while limiting noise and isolated detections. This improvement shows enhanced generalization and fault sensitivity after lengthy training. In epoch 10, the model recognized intricate fault structures even in areas without explicit supervision, indicating a considerable improvement in prediction quality.

After selecting an appropriate epoch number (epoch = 12), research was conducted to test model performance by varying the number of training patches while maintaining a constant number of epochs. Epoch 12 was chosen for its accuracy and low validation loss. Patch extraction settings aid this process. Each patch was 64×64 pixels, with a threshold of 0, suggesting that no minimum defect-pixel

count was required. The lesson set the training patch count to 8,000. However, this count was raised to 12,000 to test the effect of data volume on model learning in this study. By adding 24,000 and 48,000 training patches, the data volume impact was better analyzed. The validation set was always 2,000 patches from one labeled inline. This ensured that validation results were unaffected by changes to the data source, enabling trustworthy comparisons of training configurations.

Table 3. A comparison of accuracy and loss for 12K, 24K, and 48K training patches at epoch 12 is shown in the following list

No. of training	Accuracy	Loss	Validation accuracy	Validation loss
12,000	0.9929	0.0168	0.936	0.5636
24,000	0.9966	0.0082	0.9403	0.8801
45,000	0.9979	0.0053	0.9436	0.9941

The data show that adding training patches gradually improves performance (Table 3). Training accuracy rises




Figure 9. Normalized confusion matrices for (A) epoch 3, (B) epoch 4, (C) epoch 5, (D) epoch 10, (E) epoch 11, (F) epoch 12, (G) epoch 13, (H) epoch 14, and (I) epoch 15

from 0.9929 to 0.9979 with 48,000 patches. Meanwhile, training loss drops from 0.0168 to 0.0053. Trends show that a larger training set helps the model grasp data better. Validation metrics show no comparable improvement. Validation accuracy increases from 0.936 to 0.9436 across the three configurations, whereas validation loss increases from 0.5635 with 12,000 patches to 0.9941 with 48,000 patches. Validation accuracy peaks at 24,000 patches (0.9403), whereas validation loss rises to 0.8801. The model's inconsistency implies that its confidence in its predictions is rising, but this does not ensure increased accuracy, suggesting overfitting as training data grows. Additional training patches may enhance performance on the training dataset, but they do not reliably improve

generalization to unseen data. The model, trained with 12,000 patches, balances accuracy with validation loss. This indicates enough learning capacity without overfitting (Figure 13). The accuracy and loss graphs for 12,000, 24,000, and 48,000 patch models show consistent trends.

In all three circumstances, training accuracy rapidly increases and stabilizes at over 99%. Training loss decreases steadily, reaching zero in the final epochs. The observed trends show that the model learns from training data regardless of dataset size. Validation accuracy is steady and much lower than training accuracy, demonstrating little improvement as patches grow. The validation loss increases throughout training, peaking for the model trained with 48,000 patches (Figure 13A).




Figure 10. Predicted fault overlays for epoch 5 on inline slices (A) 100, (B) 200, (C) 300, (D) 400, (E) 500, and (F) 600




Figure 11. Predicted fault overlays for epoch 10 on inline slices (A) 100, (B) 200, (C) 300, (D) 400, (E) 500, and (F) 600

The confusion matrices show that the model accurately detects the no-fault class across all three patch volumes (Figure 14). High performance is expected because of the dataset's no-fault pixels. True positive rates for fault classification are consistent, with the model recognizing 25% of problem pixels with 12,000 patches and 24% with 24,000 and 48,000 patches. Despite adding training data, false negatives—fault pixels misclassified as no-fault—remain at 75–76%. The results show that adding patches does not improve the model's fault structure sensitivity.

Despite the robustness of no-fault detection, the fault detection rate has stabilized, supporting the idea that higher patch volumes do not help model generalization.

To enhance comprehension, the outcomes of the training with a progressively larger number of training patches are analyzed (Figures 15–17).

Distinct structural discontinuities suggest that core fault detection is unaffected by patch volume. Figure 15 shows that the 12,000 patch model captures key fault zones




Figure 12. Results of fault prediction at time slice 420: (A) Input data (uninterpreted), (B) epoch 5, and (C) epoch 10

with sparse forecasts. The model focuses on critical traits, reducing false positives and scattered noise. The 24,000 patch model has slightly more expected flaws (Figure 16). Additional minor characteristics have been found, but continuity and alignment are no better than with the 12,000-patch scenario. The 48,000 patch model predicts a wider spread, especially in low-contrast locations (Figure 17). However, this increases visual clutter by creating more disconnected fault segments that do not correspond with seismic structures. According to earlier quantitative data, training accuracy improved, but validation loss rose, suggesting overfitting.

Figure 18A shows that the 12,000-patch model predicts fault lines that are confined and focused along major reflector terminations and structural boundaries. Highlighting only the most significant and geologically relevant discontinuities, the model shows restraint. This generates a low-noise, high-precision output that matches the seismic backdrop. The 24,000 patch model (Figure 18B) predicts more fault segments in more nuanced locations.

Some forecasts are accurate, but others look inconsistent with seismic texture, indicating overprediction. The model becomes more sensitive but loses selectivity. The density of anticipated fault segments increases in the 48,000-patch model (Figure 18C). Many marked regions no longer match geological discontinuities. Essential faults are still found, but the findings are noisier, resulting in more isolated predictions and misinterpretations of non-fault reflectors. This suggests that the model is over-adapting to tiny training data properties that do not apply to new situations. Validation loss increased with performance, suggesting overfitting.

4. Discussion

This research validates the effectiveness of CNNs, particularly the U-Net architecture, in identifying 3D seismic faults. Labelled seismic inlines assisted the model's identification of fault structures in space, proving its ability to recognize geologic discontinuities. The tests involved a comprehensive assessment of model training




Figure 13. Graphs depicting training and validation accuracy and loss for 12K, 24K, and 48K training patches at epoch 12: (A) Training accuracy, (B) training Loss, (C) validation precision, and (D) validation error




Figure 14. Normalized confusion matrices for (A) 12K, (B) 24K, and (C) 48K training patches at epoch 12

and performance. The implementation of multiple training epochs resulted in a significant impact. The curves for training and validation accuracy showed initial improvements in model learning but reached an average and began to diverge around the tenth epoch. The training accuracy continued to grow up to this threshold; however, the validation accuracy reached a peak while the validation loss exhibited an increasing trend. This pattern indicates overfitting, where the model becomes excessively developed for the training data and is unable to generalize. Therefore, epoch 12 represented the optimal balance

between training performance and generalization.

We also considered training patch volume and observed that increased patches from 12,000 to 48,000 improved training data fit but not validation metrics. More patches increased validation loss, and fault prediction overlays had more noise and false positives at higher volumes. In imbalanced datasets with non-fault pixels, adding data may enhance noise patterns or non-essential features above a threshold. The algorithm trained with 12,000 patches consistently made clearer, geologically interpretable predictions. This shows that moderately well-




Figure 15. Fault prediction was conducted for 12,000 training patches at epoch 12 on the inline seismic section for line numbers (A) 100, (B) 200, (C) 300, (D) 400, (E) 500, and (F) 600




Figure 16. Fault prediction results for 24,000 training patches at epoch 12 on the Inline seismic section line numbers (A) 100, (B) 200, (C) 300, (D) 400, (E) 500, and (F) 600

prepared data is more beneficial than excessive volume. Confusion matrices supported the findings. All models had high classification performance for the dominating non-fault class, but fault pixel identification did not improve with more training data. Fault true positive rates leveled at 24–25%, indicating that minority characteristics in imbalanced datasets are challenging to categorize. The

findings suggest that future research should consider class-weighting or loss function adjustments.

The model's performance was evaluated by predicting a time-slice part that was not used during training of the model. The model trained with 12,000 patches produced the most structurally consistent and interpretable outputs.




Figure 17. Fault prediction results for 48,000 training patches at epoch 12 on inline seismic section line numbers (A) 100, (B) 200, (C) 300, (D) 400, (E) 500, and (F) 600




Figure 18. Prediction overlays at time slice 420 for (A) 12K, (B) 24K, and (C) 48K training patches at epoch 12

Models with 24,000 and 48,000 patches produced larger, noisier fault networks, indicating less geological alignment. This supports the earlier finding that increased patch volume sensitivity may weaken precision without regularization or validation-guided training. The findings show that deep learning's seismic interpretation performance is not solely dependent on model complexity or dataset quantity, but also on the dependable, interpretable, and generalizable results; parameter selection, data preparation, and validation must be done carefully. The findings indicate that deep learning has the capability to identify faults; however, there is a necessity for methodological enhancements.

5. Conclusion

U-Net, a CNN developed and implemented for seismic fault detection. A carefully constructed training procedure using 3D seismic data assesses the model's effectiveness in identifying fault structures within 3D seismic data. The

research's final results are listed below:

- A patch-based CNN approach for detecting seismic faults using 64×64 picture segments was demonstrated, enabling efficient training with limited resources. 12,000 patches of training data yield optimal generalization, resulting in more accurate fault predictions than bigger datasets.
- Excessive training data volume leads to overfitting and decreased prediction quality, emphasizing the importance of data balance over quantity. Training accuracy improved, but class imbalance remained, with true positive fault detection stabilizing at 24–25%.
- The conventional binary cross-entropy loss proves insufficient for highly unbalanced seismic datasets, requiring a larger loss function or data sampling methods. The model effectively picked faults on unlabeled seismic time-slices, proving the developed model's ability.

The research demonstrates that deep learning techniques serve as an effective tool for automated seismic interpretation. The results validate the use of a U-Net architecture for fault detection and define a systematic approach for integrating CNN-based models into geophysical analysis. Conclusions emphasize essential components for upcoming applications, including the impact of data volume, class imbalance, and overfitting, which must be considered to ensure accuracy and reliability in structural interpretation.

Acknowledgments

The authors would like to thank Geophysical Engineering at Istanbul Technical University for providing the facilities used in this research. This research is associated with the ITU DAP project (Code: TDA-2023-44996).

Funding

None.

Conflict of interest

The authors declare that they have no competing interests, financial or non-financial, that could have appeared to influence the work reported in this paper. The authors also declare that there are no personal relationships, affiliations, funding-related conflicts, or institutional interests that could be perceived as influencing the publication of this study.

Author contributions

Conceptualization: Yasir Bashir, Dilek Yüksel, Begüm Akin, Derin Gezer

Formal analysis: Yasir Bashir, Dilek Yüksel, Begüm Akin, Derin Gezer, Muhammad Khan

Investigation: Yasir Bashir, Dilek Yüksel, Begüm Akin, Derin Gezer

Methodology: Yasir Bashir, Dilek Yüksel, Begüm Akin, Derin Gezer

Supervision: Yasir Bashir, Muhsan Ehsan, Muhammad Khan, Syed Haroon Ali

Validation: Yasir Bashir, Dilek Yüksel, Begüm Akin, Derin Gezer

Writing—original draft: Yasir Bashir, Dilek Yüksel, Begüm Akin, Derin Gezer

Writing—review & editing: Yasir Bashir, Muhsan Ehsan, Muhammad Khan, Syed Haroon Ali

Availability of data

Data are available from the corresponding author upon reasonable request, and the codes are available online at the GitHub repository.

Code availability section

Name of the code/library: CNN-fault-prediction

Contact: (Email) dryasir.bashir@live.com, ybashir@itu.edu.tr; (Phone) +905366224503

Hardware requirements: Intel(R) Xeon(R) CPU E5-2678 v3 @ 2.50 GHz and an NVIDIA GeForce GTX 1080 Ti GPU.

Program language: Python

Software required: Anaconda, Jupyter Notebook

Program size: 5 Gigabytes

The source codes are available for download at: <https://github.com/dryasir/CNN-Fault-Prediction>

References

1. 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
2. Xiong W, Ji X, Ma Y, *et al.* Seismic fault detection with convolutional neural network. *Geophysics*. 2018;83(5):O97-O103. doi: 10.1190/geo2017-0666.1
3. Babikir I, Elsaadany M. Machine learning-based seismic characterization of deepwater turbidites in the Dangerous Grounds area, Northwest Sabah, offshore Malaysia. *Acta Geophys*. 2024;73(1):379-391. doi: 10.1007/s11600-024-01396-2
4. Bashir Y, Khan M, Mahgoub M, *et al.* Machine Learning Application on Seismic Diffraction Detection and Preservation for High Resolution Imaging. In: *International Petroleum Technology Conference*. IPTC; 2024:IPTC-23668. doi: 10.2523/IPTC-23668-EA
5. Ismail A, Radwan AA, Leila M, Abdelmaksoud A, Ali M. Unsupervised machine learning and multi-seismic attributes for fault and fracture network interpretation in the Kerry Field, Taranaki Basin, New Zealand. *Geomech Geophys Geo-Energy Geo-Resour*. 2023;9(1). doi: 10.1007/s40948-023-00646-9
6. Khan M, Bery AA, Bashir Y, Sharoni SMH, Ali SS. Automated Fault Network Extraction in Complex Tectonic Regimes: A Hybrid Machine Learning and Structural Attributes Approach. *Appl Comput Geosci*. 2025;27:100264. doi: 10.1016/j.acags.2025.100264
7. Wu X, Shi Y, Fomel S, Liang L. Convolutional neural networks for fault interpretation in seismic images. In: *SEG Technical Program Expanded Abstracts 2018*. Tulsa, OK: Society of Exploration Geophysicists; 2018:1946-1950. doi: 10.1190/segam2018-2995341.1
8. Guo Y, Peng S, Du W, Li D. Fault and horizon automatic

- interpretation by CNN: A case study of coalfield. *J Geophys Eng.* 2020;17(6):1016-1025.
doi: 10.1093/jge/gxaa060
9. Owusu BA, Boateng CD, Asare VDS, Danuor SK, Adenutsi CD, Quaye JA. Seismic facies analysis using machine learning techniques: a review and case study. *Earth Sci Inform.* 2024;17(5):3899-3924.
doi: 10.1007/s12145-024-01395-3
 10. Zhao T, Jayaram V, Roy A, Marfurt KJ. A comparison of classification techniques for seismic facies recognition. *Interpretation.* 2015;3(4):SAE29-SAE58.
doi: 10.1190/INT-2015-0044.1
 11. Imran QS, Siddiqui NA, Latiff AHA, *et al.* Automated fault detection and extraction under gas chimneys using hybrid discontinuity attributes. *Appl Sci.* 2021;11(16):7218.
doi: 10.3390/app11167218
 12. Khan M, Bery AA, Bashir Y, Sharoni SMH, Gnappragasan J, Imran QS. Optimizing petrophysical property prediction in fluvial-deltaic reservoirs: a multi-seismic attribute transformation and probabilistic neural network approach. *J Pet Explor Prod Technol.* 2025;15(2):29.
doi: 10.1007/s13202-024-01912-6
 13. Kamsari M, Sadeghi S, de Oliveira GG, *et al.* The role of data augmentation and attention mechanisms in UNet and ConvNeXt architectures for optimizing breast tumor segmentation. *Sci Rep.* 2025;15:45268.
doi: 10.1038/s41598-025-29841-2
 14. Anari S, Ranjbarzadeh R, Cunneen M, Bendechache M. Privacy-Preserving Federated Learning for Human Intention Modeling in Pediatric Cerebral Palsy Using Extended Reality. In: *2025 IEEE 49th Annual Computers, Software, and Applications Conference (COMPSAC)*. IEEE; 2025:1-6.
doi: 10.1109/COMPSAC65507.2025.00209
 15. Ranjbarzadeh R, Anari S, Cunneen M, Bendechache M. Lightweight Deep Learning with Virtual Reality Visualization for Offline Tumor Segmentation in Rural Environments. In: *2025 IEEE 49th Annual Computers, Software, and Applications Conference (COMPSAC)*. IEEE; 2025:1577-1582.
doi: 10.1109/COMPSAC65507.2025.00211
 16. Bagherian Kasgari A, Sadeghi S, Zarbakhsh P, Jafarzadeh Ghouschi S, Ranjbarzadeh R. A spiking convolutional neural network for glioma brain tumor segmentation using a spike-timing-dependent plasticity method. *Neurocomputing.* 2026;660:131903.
doi: 10.1016/j.neucom.2025.131903
 17. Zeng A, Yan L, Huang Y, Ren E, Liu T, Zhang H. Intelligent detection of small faults using a support vector machine. *Energies.* 2021;14(19):6242.
doi: 10.3390/en14196242
 18. Bashir Y, Babasafari AA, Alashloo SYM, Muztaza NM, Ali SH, Imran QS. Seismic wave propagation characteristics using conventional and advance modelling algorithm for d-data imaging. *J Seism Explor.* 2021;30:21-44.
 19. 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 Seism Explor.* 2025;34(5):81-98.
doi: 10.36922/JSE025310051
 20. Pang C, Zhao T, Chen G, *et al.* Earthquake and blast recognition based on CEEMDAN multiscale fuzzy entropy and NSGAIII optimized 1D-CNN neural networks. *J Seism Explor.* 2025;34(1):22-42.
doi: 10.36922/JSE025260029
 21. An Y, Guo J, Ye Q, Childs C, Walsh J, Dong R. Deep convolutional neural network for automatic fault recognition from 3D seismic datasets. *Comput Geosci.* 2021;153:104776.
doi: 10.1016/j.cageo.2021.104776
 22. Ter Borgh MM, Jaarsma B, Rosendaal EA. Structural development of the northern Dutch offshore: Paleozoic to present. In: *Geological Society, London, Special Publications*. Vol 471. London, UK: Geological Society; 2019:115-131.
doi: 10.1144/SP471.4
 23. Lima G, Zeiser FA, Da Silveira A, Rigo S, de Oliveira Ramos G. An encoder-decoder deep neural network for binary segmentation of seismic facies. *Comput Geosci.* 2024;183:105507.
doi: 10.1016/j.cageo.2023.105507
 24. Waldeland AU, Jensen AC, Gelius LJ, Solberg AHS. Convolutional neural networks for automated seismic interpretation. *Leading Edge.* 2018;37(7):529-537.
doi: 10.1190/tle37070529.1
 25. Wei XL, Zhang CX, Kim SW, *et al.* Seismic fault detection using convolutional neural networks with focal loss. *Comput Geosci.* 2022;158:104968.
doi: 10.1016/j.cageo.2021.104968
 26. Jozinović D, Lomax A, Štajduhar I, Michelini A. Rapid prediction of earthquake ground shaking intensity using raw waveform data and a convolutional neural network. *Geophys J Int.* 2020;222(2):1379-1389.
doi: 10.1093/gji/ggaa233