



A Transformer-Based Approach for Financial Risk Assessment

Anam Naz¹, Shazia Batool², Abdul Qadoos³, Hina Shafique⁴, Areeba Khalid⁵, Muhammad Saeed⁶, Hafiz Burhan Ul Haq^{7*}

¹Department of Computer Science, Faculty of IT, The University of Faisalabad, Pakistan.

²⁴⁵⁶Department of Artificial Intelligence, The Islamia University of Bahawalpur, Pakistan.

³Department of Computer Science and Digital Technologies, School of Architecture Computing and Engineering, University of East London, UK

⁷Faculty of Engineering, Department of Electronics and Telecommunication Engineering, King Mongkut's University of Technology Thonburi (KMUTT), Bangkok, Thailand

Corresponding Author: burhanhashmi64@lgu.edu.pk

Received: July 10,2026 **Accepted:** July 28,2026 **Published:** August 19,2026

ABSTRACT

Financial risk assessment identifies and measures the risk of loss in financial decisions. It includes credit, market, cash flow, and operational risk. The traditional methods depend on history and expert judgment. Machine learning speed up and standardize risk prediction. This study used Transformer model to classify financial risk into Low, Medium, and High classes. The benchmarks the model against machine learning and deep learning baselines under one protocol. We used a tabular dataset with 16 financial indicator features. Preprocessing included cleansing, encoding of categorical variables, scaling, and feature engineering. The Transformer used self-attention over the feature tokens. We compared the transformer with ML models baselines. These baselines included logistic regression, k-nearest neighbor, gradient boosting, support vector machine, random forest, XGBoost, CatBoost, a recurrent network, and a long short-term memory network. We evaluate the results using accuracy, precision, recall, F1-score, and the area under the ROC curve. The proposed Transformer reached an accuracy of 83.84% on the test set of 2,500 records. It achieved balanced precision, recall, and F1-score across the three classes. The baseline models reached 74.92% to 80.26%. Logistic regression was lowest at 74.92%, and CatBoost was the best baseline at 80.26%. The Transformer gave the highest accuracy among all tested models. This approach supports financial institutions and industries.

Keywords: Financial Risk Assessment, Transformer, Machine Learning, Deep Learning, Tabular Data.

1. INTRODUCTION

Financial risk assessment is one of the fundamental tasks of finance. It is a process of identifying, measuring and managing the risk of loss in financial activities [1]. It helps to make trade-offs between risk and return [2]. Typical Finance, market conditions, and stability are the

three categories of risk and operational risk. The management strategies will vary depending on the type [3]. Historic and expert judgement are the basis of traditional assessment. Financial ratios are calculated from several years of data by analysts [4]. They employ techniques such as scenario analysis and stress testing [5]. These are effective but time consuming and are subjective. They also have a trouble with large and complex datasets [6].




Figure 1. Framework for finance risk assessment: data acquisition, preprocessing, model training (traditional machine learning, deep learning, and transformer-based models), evaluation, model comparison, and risk identification.

The financial risk modeling has been revolutionized by machine learning. The classical models like logistic regression and random forest predict the risk based on structured data [7]. Non-linear relations between the variables are well managed by gradient boosting methods like XGBoost [8]. Temporal patterns in transactions are captured by deep models like recurrent networks [9]. These

approaches help to save time and minimize human error [10]. Financial risk evaluation still presents significant obstacles. Financial datasets frequently include a combination of quantitative and categorical parameters, missing values, class imbalance, and complicated feature interactions [11]. Traditional Machine Learning (ML) models typically necessitate extensive feature engineering to obtain good performance. Deep

Learning (DL) models can enhance prediction, but they frequently require huge labeled datasets and may not generalize well across financial environments [12]. As a result, there is a demand for models that can learn useful feature correlations while maintaining consistent performance on tabular financial data.

Many learning tasks have been improved by the transformers. They use self-attention to weigh relations between input elements [13]. The self-attention mechanism can be used to model complex interactions between financial features [14]. Pre-training on unlabeled data can also lead to the better generalization [15]. Transformers are appealing in risk assessment for these properties. While this is a step forward, there are issues that still need to be addressed. Financial information is hard to come by and unbalanced [16]. Many models act as black boxes, which limits trust in regulated settings [17]. Most previous works focus only on a few models and do not use a common protocol [18]. The utility of a pretrained transformer for tabular risk data is also unclear [19]. Recent research on tabular Transformers has demonstrated that attention-based models can learn meaningful correlations between features [20]. However, many existing efforts prioritize generic tabular benchmarking, credit scoring, or fraud detection above multi-class financial risk assessment. Furthermore, minimal research compares pretrained Transformer models to both traditional ML and DL baselines in the same experimental scenario. This makes it difficult to determine whether performance benefits are due to model architecture or differences in preprocessing, data splitting, and assessment approaches.

This paper suggests a pretrained Transformer-based method for a conventional financial risk assessment in order to address these problems. Financial risk is divided into three groups by the model: Low, Medium, and High. A tabular dataset with sixteen financial indicator characteristics is used for both training and testing. Prior to supervise fine-tuning for risk

classification, the suggested model learns feature representations using self-supervised pretraining. Then, using the same preprocessing, data splitting, and evaluation protocol, it is compared with nine machine learning and deep learning baselines. The goal is to use tabular financial data to create a dependable and automatic method for classifying financial risk. The main contributions of this study are as follows:

- A pretrained Transformer-based framework is developed exclusively for financial risk assessment using tabular financial data. This framework is capable of discovering complicated correlations between 16 different financial metrics.
- Self-supervised pretraining and supervised fine-tuning are combined in a two-stage training technique that enhances feature representation learning and facilitates improved generalization when labeled financial data are limited.
- Transformer-based self-attention is investigated for financial risk classification, allowing the model to grasp detailed correlations between financial variables with less dependence on feature associations that are explicitly created.

The rest of the paper is structured as follows: The previous works are reviewed and the research gap is defined in Section 2. The materials and methods are described in Section 3. The results are reported in Section 4. The findings of this study are discussed in Section 5 and the study is concluded in Section 6.

2. RELATED WORK

Several studies have assessed financial risk and credit using combined learning and conventional machine learning techniques. The German Credit dataset was used in study [21] to examine the effectiveness of Decision Tree, Logistic Regression, Random Forest (RF), K-Nearest

Neighbors (KNN), Support Vector Machine (SVM), Voting, Stacking, and Gradient Boosting. The KNN results are 74% after preprocessing and 61% without, demonstrating the importance of preprocessing data.

In this study [22], the authors compared seven supervised learning algorithms, namely Naïve Bayes, SVM, Decision Tree, KNN, Artificial Neural Network (ANN), Random Forest and Logistic Regression with the German Credit data set comprising of 20 attributes and 1000 samples. The classification accuracy for most models was greater than 85%, successfully demonstrating the use of traditional ML methods in credit risk prediction. In this study [23], researcher have used RF, Decision Tree, Logistic Regression and SVM for the classification of credit score. With a classification accuracy of 91.32%, RF was the best model among them, but SVM had a relatively lower classification accuracy due to its sensitivity to the kernel parameter choices. Another comparison study used a credit score dataset to assess SVM, Gradient Boosting, AdaBoost, and Logistic Regression [24]. The study used F1-score as a balanced evaluation criterion in place of accuracy, and SVM had the greatest F1-score of 87.05%.

This work [25], presented the Taiwan UCI credit data which comprises 30,000 customer records and applied Linear Discriminant Analysis (LDA), Logistic Regression, SVM, XGBoost, Random Forest, and Deep Neural Networks. The study also showed that ensemble learning and deep learning models outperformed traditional statistical models in credit default prediction.

In this study [26], a hybrid RNN-LSTM model for detecting credit card fraud utilizing European transaction data was proposed along with random approaches. By learning repetitive transaction patterns, the model demonstrated an accuracy of 99.96%.

This study [27] using deep learning has been conducted to compare LSTM, Autoencoder, and CNN on a European credit card fraud data set. LSTM gave the highest accuracy of 99.2%. Again, the gains were based on fine tuning for one particular data set. This work [28] performed a study on imbalanced credit data to compare the deep model and classical model in two augmentation methods: SMOTE and ENN. For the German dataset, SMOTE-ENN and MLP achieved 95.4% accuracy.

This research [29], presented LSTM and Deep Neural Networks were combined with resampling methods to solve class imbalance problem in credit card fraud detection. The model yielded an accuracy of 99.6% in the classification, better than Logistic Regression, Decision Tree and LightGBM. This research [30] used a 5-layer LSTM structure to predict financial risks in an IoT environment, with the UCI credit card dataset. The model reached an accuracy of 90%, F1 score of 0.91, and therefore has the ability to do sequential analysis of financial transactions. This study [31], introduced the FT-Transformer, where the representation of each feature is tokenized, and the higher order feature interactions are modeled by transformer blocks. The model performed competitively on multiple benchmarks tabular datasets in comparison with Gradient Boosting. In this research [32], presents the TabPFN model was proposed, which used transformer-based pretraining on synthetic tabular data. The pretrained model can already learn what it is supposed to do and can make predictions in only a few seconds, with a high degree of accuracy, especially for small datasets. This study [33], proposed a transformer-based method for detecting credit card fraud with self-attention mechanisms. The effectiveness of attention mechanisms in financial transaction analysis was shown by comparing the accuracy of the proposed transformer with XGBoost and TabNet in experiments, and the transformer was found to be more effective with an accuracy rate of 99.8%. There are still issues with machine

learning, deep learning, and Transformer models, considering their previous achievements in financial risk assignments. Conventional approaches rely on human feature engineering and are unable to handle detailed relationships between indicators. Although deep learning improves predictions, it is frequently task-specific and requires large labeled datasets. Few research examine transformer models' pretrained applications for multi-class financial risk assessment, considering its potential with tabular data. Performance comparisons are limited by the inconsistent methods, preprocessing, datasets, and measures seen in existing studies.

3. METHODOLOGY

A detail description of the dataset, methods and

proposed model is provided in this section.

3.1. Dataset

The data for this study was a tabular financial risk data set. There are 16 financial indicator features in each row representing one person. These are age, income, number of credit cards, loans, mortgage and credit score. There is a risk label (Low, Medium, High) for each row. Label is the object to be classified. The number of classes in the raw data was unbalanced. Medium risk class was the most common. The imbalance can lead to a model being biased towards the majority class. Minority classes were balanced prior to training using a balancing step, as shown in Fig. 2. The class counts before and after balancing are presented in Table I.




Figure 2. Risk class distribution data

Table 1. Risk class distribution before and after balancing

Risk class	Before balancing	After balancing
Low Risk	3,100	4,200
Medium Risk	5,400	5,400
High Risk	3,300	4,200
Total	11,800	13,800

3.2. Preprocessing

The tabular data was preprocessed before used in models. Data Cleansing: Dealing with missing data, eliminating duplicate data. It also fixed the inconsistencies in the category names. The following steps helped to enhance the quality of the data. They lessened noise which may interfere with learning. Categorical variables were converted to numbers by encoding. Values

for numerical features were transformed to a common range. The debt to income ratio was one example of a new ratio that was devised through feature engineering. These features derived added some useful signal. The entire pipeline is depicted in Fig. 3. Class imbalance was addressed prior to training. The number of labels for minority classes was balanced by oversampling. The balanced data reduced the risk class bias towards Medium. This action led to an improvement in the scores per class.




Figure 3. Proposed financial risk assessment pipeline

3.3. Data Splitting

The dataset was split into train and test sets. The split used most of the data for training and left a test set of 2,500 records. The training set teaches the model the relationship between features and risk. The performance on unseen records was measured on the test set. Hyperparameters were tuned in cross-validation to avoid overfitting.

3.4. Evaluation Measures

Model performance was evaluated by four metrics. Accuracy was the number of all predictions that were right. Specificity is the ratio of true positives to all positive predictions. Recall is the percentage of true cases that were identified by the model. The F1-score is a harmonic mean of precision and recall.

The overall classification accuracy is stated in equation (1) as:

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \quad (1)$$

Equation (2) states that accuracy is the proportion of actual positive cases among all instances that the classifier predicted to be positive. By divide the true number of occurrences by the classifier's predicted number of successes, precision may be calculated.

$$Precision = \frac{True\ positive}{True\ positive + False\ positive} \quad (2)$$

Recall is the fraction of true positive results among all positive occurrences in the test set, as given in equation (3).

$$Recall = \frac{True\ positive}{True\ positive + False\ negative} \quad (3)$$

Equation (4) indicates that the F1 score is the

balanced mean of both recall and precision, resulting in a combined indicator of recall and precision.

$$F1\ score = 2 * \frac{Precision * Recall}{Precision + Recall} \quad (4)$$

Here TP & TN are the true positives and true negatives. FP and FN are false positive and false negatives. The area under the ROC curve for each class is also reported in the study. It shows how good the model separates one risk class from the other.

4. PROPOSED PRETRAINED TRANSFORMER MODEL

4.1. Self-Attention

The Transformer uses self-attention to relate input features. Each feature is projected into a query, a key, and a value vector. The model scores each pair of features by the dot product of query and key. These scores form attention weights after a softmax. The weights combine the value vectors into a new representation.

$$Attention(Q,K,V) = \text{SoftMax}\left(\frac{QK^T}{\sqrt{d}} + B\right)V \quad (5)$$

Here Q, K, and V are the query, key, and value matrices. The term d_k is the dimension of the key vectors. Multi-head attention repeats this process with several heads. Each head learns a different view of the feature interactions.

4.2. Pretraining and Fine-tuning

The model used a two-stage training scheme. In the first stage, it was pretrained with a self-supervised objective. Random feature values were masked, and the model learned to reconstruct them. This stage taught the model the structure of the financial features.

In the second stage, the model was fine-tuned for classification. A classification head was added on top of the encoder. The full model was trained on the labeled risk data. This scheme reused the pretrained weights and improved generalization.

4.3. Model Architecture

The proposed model is a compact Transformer encoder. The input is the vector of 16 financial features. A feature embedding layer projects the features into a token space. Positional encoding adds order information. The encoder stacks several blocks of multi-head attention and feed-forward layers.

Each block uses residual connections and layer normalization. A pooling step then summarizes the encoder output. A dense layer with a softmax activation returns a probability for each risk class. Fig. 4 shows the architecture, and Table 3 lists the hyperparameters.




Figure 4. Architecture of the proposed pretrained Transformer

Table 2: Parameter details of proposed model

Model Parameter	Configuration setting
Input features	16
Embedding dimension	64
Attention heads	4
Encoder blocks	3
Feed-forward dimension	128
Dropout rate	0.2
Output units	3 (Softmax)
Optimizer	Adam
Loss function	Categorical cross-entropy
Pretraining objective	Masked feature modeling
Epochs (fine-tuning)	40
Batch size	32

4.4. Training Procedure

Algorithm 1. Pretraining and fine-tuning the Transformer.

Input: financial records X , risk labels Y (3 classes)

Output: trained Transformer model M

1: $X' \leftarrow \text{clean}(X); \text{encode}(X'); \text{scale}(X'); \text{engineer_features}(X')$

2: $X' \leftarrow \text{balance_classes}(X')$

3: # Stage 1: self-supervised pretraining

4: for epoch = 1 to P do

5: $X_m \leftarrow \text{mask_random_features}(X')$

6: train encoder to reconstruct masked values

7: end for

8: # Stage 2: supervised fine-tuning

9: attach classification head to pretrained encoder $\rightarrow M$

10: split X' , Y into train and test (test = 2500 rows)

11: for epoch = 1 to 40 do

12: for each batch (x_b, y_b) in train do

13: $p \leftarrow M(x_b); L \leftarrow \text{cross_entropy}(p, y_b)$

14: update M with Adam to minimize L

15: end for

16: end for

17: evaluate M on the test set; return M

5. RESULTS

Experimental evaluation results of proposed model and discussion is presented in this section.

5.1. Experimental Setup

The models were executed in Python with the Keras and TensorFlow libraries. The classical models used the scikit-learn library. All models used the same data, split, preprocessing, and metrics. This common protocol enables a fair comparison. We ran experiments on a workstation with a modern CPU and GPU.

5.2. Training Behavior

Fig. 5 illustrated the accuracy of training and validation over 40 epochs. The two curves increase fast in the first epochs. Then they flatten as the model converges. The validation accuracy is about 83.84% which is close to the training accuracy. The small gap shows there is not much overfitting.




Figure 5. Accuracy of the transformer's training and validation

The training and validation losses are shown in Fig. 6. Both curves decline steadily and then levels off. The validation loss remains close to

the training loss. This pattern results in stable learning without strong overfitting




Figure 6. Training and validation loss of the proposed Transformer

5.3. Baseline Model Results

The two strongest baselines were CatBoost and the recurrent neural network. CatBoost was the

best machine learning baseline with an accuracy of 80.26%. Table III reports its per-class scores. The model classified the Medium risk class well. It made more errors on the Low risk class, where recall was lower.

Table 3. Per-class performance of the CatBoost baseline

Class	Precision	Recall	F1-score
Low Risk	0.74	0.69	0.71
Medium Risk	0.82	0.89	0.85
High Risk	0.81	0.78	0.79
Accuracy	0.8026		

The recurrent neural network was the best deep learning baseline. It reached a result of 79.45%. Table IV reports its per-class scores. The model

gave strong results for the Medium risk class. It missed more cases in the Low and High risk classes, where recall dropped.

Table 4. Per-class performance of the RNN baseline

Class	Precision	Recall	F1-score
Low Risk	0.73	0.68	0.70
Medium Risk	0.81	0.88	0.84
High Risk	0.80	0.77	0.78
Accuracy	0.7945		

5.4. Proposed Model Performance

The proposed Transformer reached an accuracy of 83.84% on the test set. Table V described the

per-class precision, recall, and F1-score. The Medium risk class scored highest on recall. The High risk class scored highest on precision. The balanced scores show that the model handled all three classes well.

Table 5. Per-class performance of the proposed Transformer model

Class	Precision	Recall	F1-score
Low Risk	0.83	0.80	0.82
Medium Risk	0.81	0.90	0.85
High Risk	0.89	0.79	0.84
Accuracy	0.8384		



Figure 7. *Per-class precision, recall, and F1-score of the proposed model*

Fig. 8 shows the confusion matrix on the test set of 2,500 records. Most predictions lie on the diagonal, which reflects correct classification. The model classified 2,096 records correctly.

The main errors occur between the Low and Medium risk classes. These classes share feature patterns that blur the boundary




Figure 8: *Confusion matrix of transformer model*

5.5 Discussion

The suggested pretrained Transformer model's accuracy, precision, recall, and F1-score are compared to nine machine learning and deep learning baseline models in Table VI. With an accuracy of 83.84%, precision of 0.84, recall of 0.83, and F1-score of 0.84, the results described that the suggested Transformer outperformed all analyzed models. These findings show that the suggested model can successfully learn intricate correlations between financial parameters and offers a balanced classification performance across various risk categories. With an accuracy of 80.26%, CatBoost outperforms the other baseline models, followed by RNN with 79.45% and XGBoost with 78.90%. When compared to more straightforward models like Logistic Regression (74.92%) and K-Nearest Neighbor (75.60%), the tree-based ensemble models, such as Random Forest (78.10%) and Gradient Boosting (76.30%), performed better. With an accuracy of 77.40%, Support Vector Machine

performed effectively but dropped short of ensemble-based methods. The baseline models for deep learning also produced respectable outcomes. The RNN model performs somewhat better with an accuracy of 79.45% and an F1-score of 0.79, compared to the LSTM model's 76.85% accuracy, precision, recall, and F1-score of 0.76. The suggested Transformer, however, beat both consecutive models, showing that self-attention-based feature interaction learning is superior for this tabular financial risk categorization problem. Overall, the suggested Transformer increased accuracy by 83.84% points, as shown in Fig. 9.

The results show a clear ordering. Logistic regression was lowest at 74.92%, which reflects its linear form. The ensemble and deep baselines improved through richer feature handling. CatBoost was the strongest baseline at 80.26%. The proposed Transformer passed all baselines by 3.58 percentage points or more.

Table 6. Comparison of the proposed model with baseline models

Model	Accuracy	Precision	Recall	F1-score
Logistic Regression	74.92%	0.74	0.73	0.73
K-Nearest Neighbor	75.60%	0.75	0.74	0.74
Gradient Boosting	76.30%	0.76	0.75	0.75
LSTM	76.85%	0.76	0.76	0.76
Support Vector Machine	77.40%	0.77	0.76	0.76
Random Forest	78.10%	0.78	0.77	0.77
XGBoost	78.90%	0.79	0.78	0.78
RNN	79.45%	0.79	0.78	0.79
CatBoost	80.26%	0.80	0.79	0.79
Proposed Transformer	83.84%	0.84	0.83	0.84




Figure 9. Accuracy comparison across models

5.6. Limitations

The proposed pretrained Transformer model shows potential, however it has some drawbacks. Its findings may not be generalizable to varied financial contexts because they were based on a single financial risk dataset of 16 indicators. The dataset's static nature does not reflect shifting risk patterns, and the model's decision-making process is less interpretable than traditional methods. Furthermore, the preprocessing and feature engineering processes create variability, which might affect performance. The study focuses entirely on predictive capacity, ignoring factors such as computing cost, scalability, data privacy, and adaptation to changing conditions.

6. CONCLUSION

This study developed a pretrained Transformer for financial risk assessment. The model was built using tabular data to categorize risk as Low, Medium or High. A complete pipeline was used to cleanse, encode, scale, feature engineer and balance data. The model was based on self-attention and self-supervised pretraining step. The proposed Transformer improved the accuracy to be 83.84% over test set. It has made

a good balance of precision, recall and F1 score for the three classes. In addition, it also achieved nine baselines from 74.92% to 80.26%. The results demonstrate that a pretrained Transformer brings improvement to three-class financial risk classification. This research is to be extended in the following directions in future. Financial data can be multimodal, such as a combination of financial tabular data and financial text. Temporal risk signals can be captured in the time-series inputs. Explainable methods can help enhance regulators' transparency. More external validity would be achieved by testing larger and more representative sets of data. Future research will enhance this study by: testing the pretrained Transformer model on larger, different financial datasets; incorporating temporal financial data to capture evolving risk patterns; investigating multimodal evaluations by combining tabular and textual data; employing explainable AI methods to improve prediction transparency; and assessing real-world deployment in terms of computational costs, scalability, data privacy, and model updating.

7. REFERENCES

- [1]. V. Shcherbak *et al.*, “Financial risk management and resilience of small enterprises amid the wartime crisis,” *Journal of Risk and Financial Management*, vol. 19, no. 1, p. 37, 2026.
- [2]. Z. Wang, “Big data-driven enterprise financial risk management innovation research,” in *Proc. 2025 2nd Int. Conf. Innovation Management and Information System*, 2025.
- [3]. L. Yang, “Predictive modeling of tax compliance risks: A comparative study of machine learning approaches,” *PLoS ONE*, vol. 20, no. 9, Art. no. e0331715, 2025.
- [4]. S. Jana *et al.*, “Application of fuzzy AHP in priority based selection of financial indices: A perspective for investors,” *Economics*, vol. 12, no. 1, pp. 1–7, 2024.
- [5]. V. Lagasio *et al.*, “Integrating generative AI and large language models in financial sector risk management: Regulatory frameworks and practical applications,” *Risk Management Magazine*, 2025.
- [6]. H. Wu and J. Zhang, “TabTransformer-based credit default prediction for structured economic data,” *Information Systems and Economics*, vol. 6, no. 2, 2025.
- [7]. J. Černevičienė and A. Kabašinskas, “Explainable artificial intelligence (XAI) in finance: A systematic literature review,” *Artificial Intelligence Review*, vol. 57, no. 8, Art. no. 216, 2024.
- [8]. Z. Gu *et al.*, “Credit risk assessment of small and micro enterprise based on machine learning,” *Heliyon*, vol. 10, no. 5, 2024.
- [9]. A. Oguntimilehin *et al.*, “Financial fraud detection model using recurrent neural network,” in *Proc. 2024 IEEE 5th Int. Conf. Electro-Computing Technologies for Humanity (NIGERCON)*, 2024.
- [10]. X. Zhang *et al.*, “Data-driven loan default prediction: A machine learning approach for enhancing business process management,” *Systems*, vol. 13, no. 7, Art. no. 581, 2025.
- [11]. A. Shibani *et al.*, “Financial risk management in the construction projects,” *Journal of King Saud University—Engineering Sciences*, vol. 36, no. 8, pp. 552–561, 2024.
- [12]. K. Ouazzane *et al.*, “A novel feature engineering framework for cyber fraud detection using machine learning and deep learning algorithms,” *Information Security Journal: A Global Perspective*, pp. 1–21, 2025.
- [13]. C. Sanford, D. J. Hsu, and M. Telgarsky, “Representational strengths and limitations of transformers,” in *Advances in Neural Information Processing Systems*, vol. 36, pp. 36677–36707, 2023.
- [14]. E. Siphuma and T. van Zyl, “Enhancing credit risk assessment through transformer-based machine learning models,” in *Proc. Southern African Conf. Artificial Intelligence Research*, Springer, 2024.
- [15]. J. Chen *et al.*, “Temporal transformer with conditional tabular GAN for credit card fraud detection: A sequential deep learning approach,” *Mathematics*, vol. 14, no. 7, Art. no. 1183, 2026.
- [16]. S. Shivam *et al.*, “Transformer based framework for imbalanced transaction fraud detection in FinTech systems,” in *Proc. 2026 14th Int. Symp. Digital Forensics and Security (ISDFS)*, IEEE, 2026.
- [17]. X. W. Tan and S. Kok, “Explainable AI for comprehensive risk assessment for financial reports: A lightweight hierarchical transformer network approach,” *arXiv*

preprint arXiv:2506.23767, 2025.

- [18]. Y. Wang *et al.*, “Leveraging convolutional neural network-transformer synergy for predictive modeling in risk-based applications,” in *Proc. 2024 4th Int. Conf. Electronic Information Engineering and Computer Communication (EIECC)*, IEEE, 2024.
- [19]. S. K. S. Nagaraj, “Empirical investigation of deep learning architectures for systematic credit risk classification in heterogeneous financial markets,” *International Journal of Emerging Trends in Computer Science and Information Technology*, vol. 6, no. 3, pp. 19–27, 2025.
- [20]. P. Raliphada, S. Olukanmi, and M. Olusanya, “Transformer-based NLP approaches for credit risk prediction: A systematic review,” *Frontiers in Artificial Intelligence*, vol. 9, Art. no. 1819994, 2026.
- [21]. R. Ma, “German credit risk prediction using machine learning models,” in *Proc. 2025 3rd Int. Conf. Image, Algorithms, and Artificial Intelligence (ICIAAI 2025)*, Atlantis Press, 2025.
- [22]. H. Ayari, R. Guetrai, and N. N. Kraiem, “Credit scoring enhancement via ensemble learning and self-organizing map-based feature transformation,” 2025.
- [23]. I. Emmanuel, Y. Sun, and Z. Wang, “A machine learning-based credit risk prediction engine system using a stacked classifier and a filter-based feature selection method,” *Journal of Big Data*, vol. 11, no. 1, Art. no. 23, 2024.
- [24]. S. Nayak and A. Bushara, “Uncertainty-aware fraud detection using hybrid transformer with gated token mixing and conformal risk control,” *IEEE Access*, 2026.
- [25]. R. Bhandary and B. K. Ghosh, “Credit card default prediction: An empirical analysis on predictive performance using statistical and machine learning methods,” *Journal of Risk and Financial Management*, vol. 18, no. 1, Art. no. 23, 2025.
- [26]. H. Zhou, X. Zheng, and J. Du, “A deep learning hybrid RNN-LSTM model for credit card fraud detection,” 2025.
- [27]. S. S. Sulaiman, I. Nadher, and S. M. Hameed, “Credit card fraud detection using improved deep learning models,” *Computers, Materials & Continua*, vol. 78, no. 1, 2024.
- [28]. S. Mestiri, “Credit scoring using machine learning and deep learning-based models,” *Data Science in Finance and Economics*, vol. 4, no. 2, p. 236, 2024.
- [29]. A. Ahsan *et al.*, “An intelligent method for improving credit card fraud detection using a hybrid LSTM and deep neural network framework,” *UMT Artificial Intelligence Review*, vol. 5, no. 1, 2025.
- [30]. A. A. Khine and Z. T. T. Myint, “Evaluating the impact of applied sampling algorithm on CNN and LSTM models for credit card fraud analysis,” in *Proc. 2025 IEEE Conf. Computer Applications (ICCA)*, IEEE, 2025.
- [31]. J.-P. Jiang *et al.*, “Representation learning for tabular data: A comprehensive survey,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2026.
- [32]. N. Hollmann *et al.*, “TabPFN: A transformer that solves small tabular classification problems in a second,” *arXiv preprint arXiv:2207.01848*, 2022.
- [33]. C. Yu *et al.*, “Credit card fraud detection using advanced transformer model,” in *Proc. 2024 IEEE Int. Conf. Metaverse Computing, Networking, and Applications (MetaCom)*, IEEE, 2024.