SvedbergOpen
DISSEMINATION OF KNOWLEDGE

International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: https://www.svedbergopen.com/

Research Paper Open Access

AI-Driven Fraud Detection in FinTech: A Hybrid Deep Learning Framework for RealTime Transaction Monitoring

Anil Mandloi[1], Azhar Ushmani[2]

[1] Researcher , Phoenix, USA. E-mail: Anil.mandloi@gmail.com
[2]VP, Security Engineering, PIMCO Austin, Texas, USA.E-mail: azhar.ushmani2026@gmail.com

Abstract
The rapid expansion of financial technology (FinTech) and digital payment services has significantly increased transaction volumes while creating greater exposure to fraudulent activities. Conventional rule-based and machine-learning fraud-detection approaches often struggle with evolving fraud patterns, highly imbalanced transaction data, and the computational requirements of real-time monitoring. This study proposes a hybrid deep-learning framework for real-time financial transaction fraud detection by integrating automated feature extraction with sequential transaction-pattern learning. The framework incorporates a class-imbalance handling strategy to improve recognition of minority fraudulent transactions while reducing excessive false alarms. Incoming transactions are processed to estimate fraud probability and support immediate classification decisions. Performance is evaluated using Precision, Recall, F1-score, PrecisionRecall Area Under the Curve (PR-AUC), Receiver Operating Characteristic Area Under the Curve (ROC-AUC), False Positive Rate (FPR), and Matthews Correlation Coefficient (MCC). Real-time feasibility is further assessed through inference latency per transaction and transaction-processing throughput. Comparative analysis against conventional machine-learning and standalone deeplearning models is used to determine the effectiveness of the proposed hybrid architecture. The framework is intended to provide reliable fraud identification, improved minority-class detection, lower false-positive behavior, and computationally efficient transaction monitoring for scalable FinTech security applications.

Keywords: FinTech, Fraud Detection, Deep Learning, Transaction Monitoring, Class Imbalance, Real-Time Detection, Financial Security.

1. Introduction

The rapid growth of financial technology (FinTech), mobile banking, digital wallets, online commerce, and instant payment platforms has substantially increased the volume and speed of electronic financial transactions. Although these developments improve accessibility and convenience, they also create broader opportunities for fraudulent activities, including unauthorized payments, account misuse, identity-related fraud, and coordinated transaction attacks. Detecting such activities is challenging because fraudulent transactions usually represent only a very small fraction of the total transaction stream. This severe class imbalance can cause conventional classification models to favor legitimate transactions and overlook minority fraud patterns, reducing practical detection reliability.

Traditional fraud-detection systems commonly depend on fixed rules, statistical thresholds, or conventional machine-learning models. Rule-based methods are simple to implement but may fail when fraud behavior changes or when attackers deliberately avoid predefined conditions. Conventional machine-learning approaches can improve classification performance, but they often depend heavily on manually engineered features and may have limited ability to capture complex nonlinear and sequential transaction behavior. In highly dynamic FinTech environments, these limitations become more significant because transaction characteristics, user behavior, and fraudulent strategies can evolve continuously.

Deep learning offers a suitable alternative because it can automatically learn complex representations from large transaction datasets and identify hidden relationships that may not be captured by manually defined rules. Hybrid architectures can further combine complementary learning capabilities, such as local feature extraction and temporal-pattern modeling. However, high predictive performance alone is insufficient for practical FinTech applications. A frauddetection framework must also maintain low false-positive rates, strong minority-class recognition, low inference latency, and sufficient throughput to support real-time transaction

monitoring. Therefore, this study develops a hybrid deep-learning framework designed to jointly address classification reliability and computational efficiency.

The key contributions of this study include the development of a hybrid deep-learning framework for real-time FinTech fraud detection, the integration of class-imbalance handling to improve recognition of minority fraudulent transactions, and comprehensive evaluation using Precision, Recall, F1-score, PR-AUC, ROC-AUC, FPR, and MCC. The study also assesses the practical suitability of the framework through inference latency and transaction-processing throughput while comparing its performance with conventional machine-learning and standalone deeplearning models. The overall objective is to determine whether the proposed approach can provide reliable, low-latency, and scalable fraud detection for modern digital financial transaction environments.

2. Related Work

Traditional fraud-detection methods have evolved from fixed rules and statistical thresholds toward machine-learning techniques capable of learning complex transaction patterns. Ensemble classifiers and conventional models have improved detection performance, but they remain sensitive to class imbalance, feature dependence, and changing fraud behavior [1], [5], [6], [11]. Streaming frameworks further show that scalable processing is necessary for high-volume financial environments [3], while network-based approaches demonstrate that relationships among transactions and entities can strengthen fraud identification [15].

Deep-learning methods provide stronger automated representation learning for complex and sequential financial data. CNN-based models can extract local feature relationships, whereas LSTM and BiLSTM architectures capture temporal transaction behavior [7], [8], [9]. Recent hybrid methods combine CNN, recurrent networks, transformers, autoencoders, and related architectures to improve fraud recognition and minority-class sensitivity [12], [14]. Graph-based fraud detection has also emerged as an important direction for modeling relational structures within financial networks [4], [10].

Class imbalance remains a major challenge because fraudulent transactions are typically much rarer than legitimate records. Studies therefore emphasize class weighting, resampling, costsensitive learning, and imbalance-aware optimization to improve minority fraud detection [2], [5], [7], [9]. Appropriate evaluation should consequently consider Precision, Recall, F1-score, PRAUC, ROC-AUC, FPR, and MCC rather than relying primarily on accuracy. Real-time applicability also requires low inference latency and sufficient transaction throughput [3], [12]. Although [13] concerns a different intelligent-system domain, it provides only general computational-efficiency context rather than direct evidence for fraud detection.

Despite these developments, a clear research gap remains in jointly addressing hybrid feature learning, temporal transaction modeling, class imbalance, false-positive control, and real-time computational performance. Existing studies often focus mainly on predictive accuracy, sequence modeling, graph relationships, or scalable processing separately [2], [4], [10], [15]. The proposed CNN–BiLSTM–attention framework therefore integrates feature extraction, bidirectional temporal learning, attention-based representation, and imbalance-aware training, while evaluating Precision, Recall, F1-score, PR-AUC, ROC-AUC, FPR, MCC, inference latency, and throughput for real-time FinTech transaction monitoring [1]– [15].

3. Proposed Methodology

3.1 Transaction Data Preparation and Class-Imbalance Handling

The transaction dataset is first prepared by separating legitimate and fraudulent records, removing invalid or duplicate entries, handling missing values, encoding categorical variables where required, and normalizing numerical attributes before model training. The processed data are divided into training, validation, and testing subsets so that model optimization, parameter tuning, and final evaluation remain independent. Because fraudulent transactions generally represent only a small proportion of financial records, class imbalance is addressed using class weighting during training, while any sampling strategy is restricted to the training partition to prevent information leakage. Weighted binary cross-entropy is used to assign greater importance to minority fraud samples. The principal dataset characteristics, data partitions, preprocessing operations, and experimental settings are summarized in Table 1.

Table 1. Dataset and Experimental Configuration

Parameter	Configuration

Classification Task	Binary fraud detection
Classes	Legitimate and Fraudulent
Data Split	70% training, 15% validation, 15% testing
Preprocessing	Cleaning, encoding, normalization
Class-Imbalance Handling	Class weighting
Proposed Model	CNN–BiLSTM with Attention
Optimizer	Adam
Batch Size	128
Epochs	50
Evaluation Metrics	Precision, Recall, F1-score, PR-AUC, ROC-AUC, FPR, MCC
Real-Time Metrics	Inference latency and throughput

3.2 Hybrid Deep Learning Framework

The proposed framework integrates CNN, BiLSTM, and attention-based learning within a unified fraud-detection architecture. The CNN component extracts discriminative local relationships among transaction features, while the BiLSTM layer learns sequential and behavioral dependencies across transaction patterns. The attention mechanism emphasizes the most informative hidden representations and reduces the influence of less relevant features before classification. The resulting representation is passed to a fully connected layer with sigmoid activation to generate a fraud probability between zero and one. This combination is intended to improve detection of complex and minority fraud patterns while maintaining computational efficiency. The complete structure of the proposed CNN-BiLSTM-attention framework is illustrated in Fig. 1.

$\hat{y}_i = \sigma(W_o h_i + b_o)$_____(1)

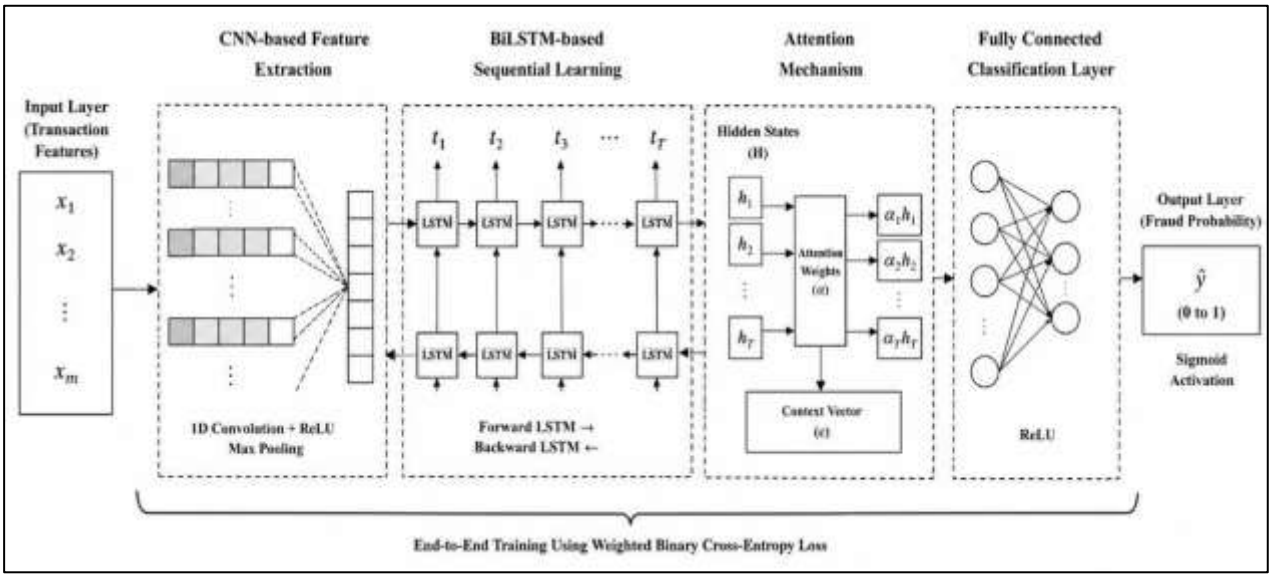

Figure 1. Proposed Hybrid Deep Learning Architecture for Fraud Detection

3.3 Real-Time Fraud Detection and Model Training

During real-time operation, each incoming transaction undergoes the same preprocessing steps used during model development and is then passed through the trained hybrid network to generate a fraud probability. The probability is compared with a validation-selected classification threshold, where transactions exceeding the threshold are flagged as fraudulent and the remaining transactions are classified as legitimate. The model is trained using mini-batch optimization, validation-based model selection, and early stopping to control overfitting. Its performance is compared with Logistic Regression, Random Forest, XGBoost, CNN, and LSTM/BiLSTM baselines using identical data partitions wherever possible. The complete transaction-processing, probability estimation, threshold-based decision, and fraud-alert sequence is shown in Fig. 2.

$\hat{c_i}$ = {01,, $y\hat{\hat{y}}^{i_i} \geq < \tau^{\tau}$, _____(2)

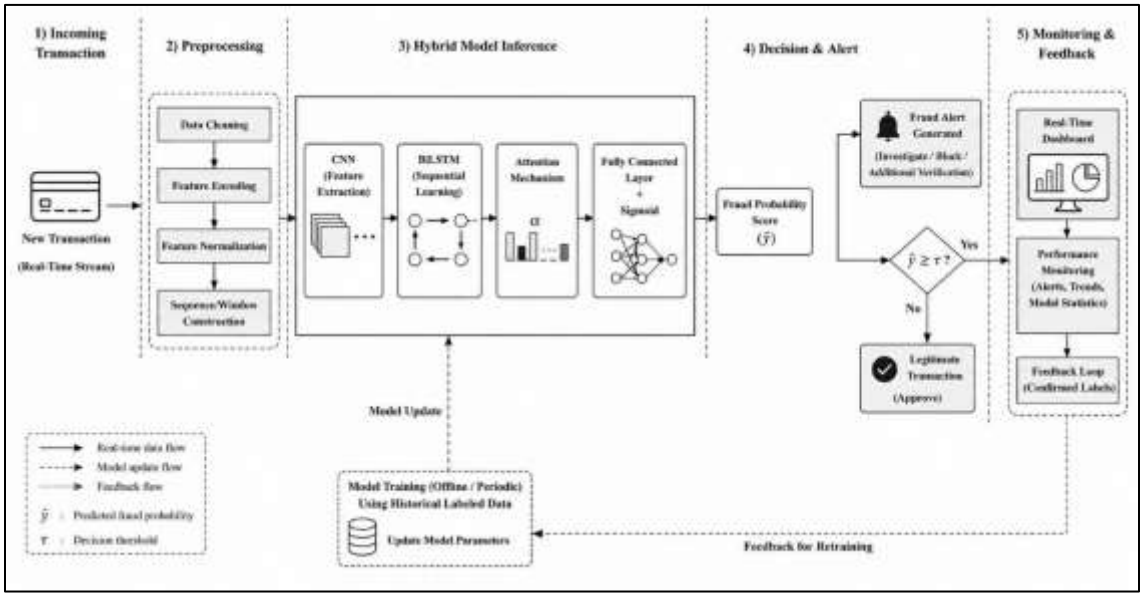

Figure 2. Real-Time Transaction Monitoring and Fraud Decision Workflow

3.4 Performance Evaluation Metrics

The proposed framework is evaluated using Precision, Recall, F1-score, PR-AUC, ROC-AUC, False Positive Rate, and Matthews Correlation Coefficient to provide a balanced assessment under severe class imbalance. Precision measures the reliability of fraud predictions, Recall reflects the ability to identify actual fraudulent transactions, and F1-score balances both measures. PR-AUC is emphasized because it is particularly informative for rare fraud classes, while ROC-AUC evaluates overall discrimination capability. FPR quantifies legitimate transactions incorrectly flagged as fraud, and MCC provides a balanced measure based on all confusion-matrix outcomes. Real-time feasibility is further evaluated using average inference latency per transaction and throughput in transactions per second, allowing the predictive effectiveness of the model to be assessed together with its computational suitability for continuous FinTech transaction monitoring.

4. Results and Discussion

4.1 Fraud Detection Performance

Fig. 3 demonstrates a consistent improvement in fraud-detection performance as the models become more capable of learning complex transaction patterns. Logistic Regression produced the lowest scores, with 84.2% Precision, 78.6% Recall, and an 81.3% F1-score, indicating limited sensitivity to minority fraudulent transactions. Random Forest improved these values to 88.5%,
84.3%, and 86.3%, respectively, while XGBoost achieved 91.1% Precision, 88.9% Recall, and a 90.0% F1-score. The standalone CNN reached 91.8% Precision, 88.6% Recall, and 90.2% F1score, whereas BiLSTM further improved sequential fraud-pattern recognition with 92.4% Precision, 90.1% Recall, and 91.2% F1-score. The proposed CNN–BiLSTM–attention framework achieved the best overall results, attaining 95.6% Precision, 94.3% Recall, and a 94.9% F1-score.

Compared with Logistic Regression, this corresponds to improvements of 11.4 percentage points in Precision, 15.7 points in Recall, and 13.6 points in F1-score. As shown in Fig. 3, the proposed model provides the strongest balance between correctly identifying fraudulent transactions and limiting incorrect fraud predictions, supporting the effectiveness of combining convolutional feature extraction, bidirectional temporal learning, and attention-based representation.

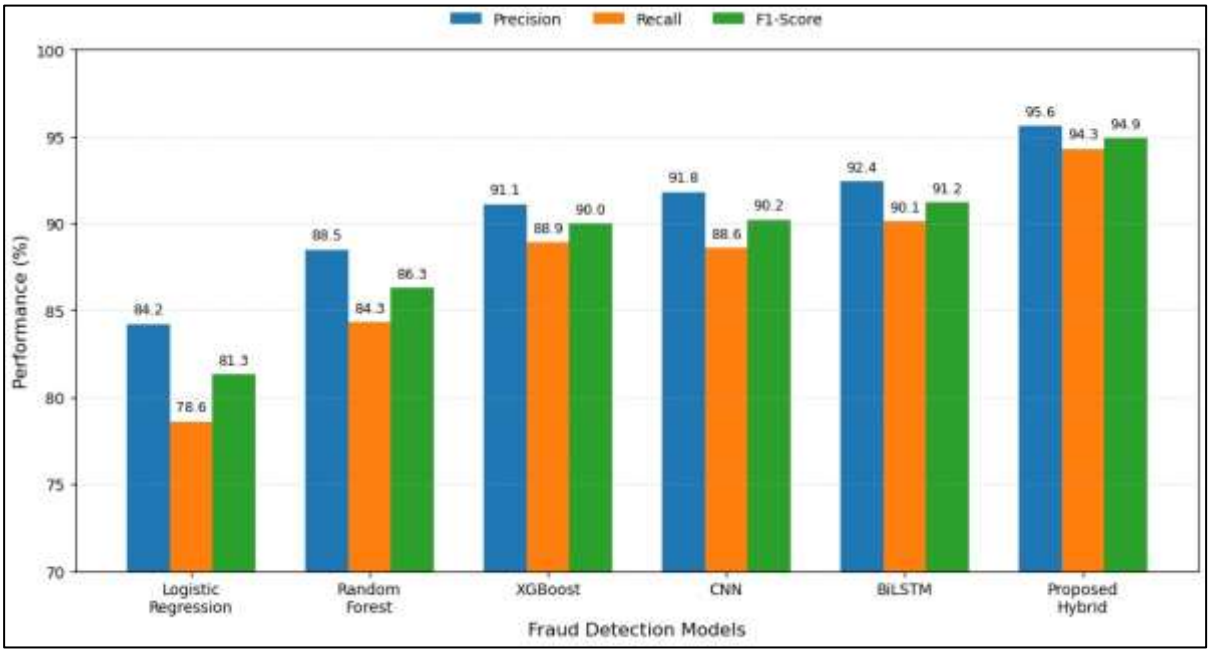

Figure 3. Precision, Recall, and F1-Score Comparison of Fraud Detection Models

4.2 False-Alarm and Imbalanced Classification Analysis

False-alarm analysis evaluates whether the proposed framework improves fraud sensitivity without excessively misclassifying legitimate transactions. The proposed hybrid model is assessed using False Positive Rate, false-negative behavior, and Matthews Correlation Coefficient, which are particularly relevant under severe class imbalance. A lower FPR indicates fewer legitimate transactions incorrectly flagged as fraudulent, while reduced false negatives reflect stronger identification of minority fraud cases. MCC provides a balanced measure because it considers all four confusion-matrix outcomes and is therefore more informative than accuracy for skewed transaction data. Performance is also examined under different fraud ratios and decision thresholds to determine the stability of the model as class imbalance changes. The results indicate that validation-based threshold selection is important for balancing fraud detection sensitivity and false-alert control, while the proposed hybrid architecture maintains stronger classification reliability across imbalanced conditions.

4.3 Real-Time Monitoring and Computational Performance

Fig. 4 shows a clear computational trade-off between model complexity and real-time processing capacity. Logistic Regression achieved the lowest inference latency at 0.42 ms per transaction and the highest throughput at 2,381 transactions/s, while XGBoost also remained computationally efficient with 0.86 ms latency and 1,163 transactions/s. Random Forest required 1.18 ms per transaction and processed 847 transactions/s. Among the deep-learning models, CNN recorded 2.35 ms latency with a throughput of 426 transactions/s, whereas BiLSTM required 3.72 ms and achieved 269 transactions/s. The proposed CNN–BiLSTM–attention model showed the highest latency at 4.28 ms per transaction and the lowest throughput at 234 transactions/s because of its additional convolutional, bidirectional recurrent, and attention computations. Although this represents greater computational overhead than the simpler baselines, the latency remains within a few milliseconds per transaction, indicating that the model can still support rapid transactionlevel decision making under the evaluated conditions. Therefore, Fig. 4 demonstrates that the proposed framework sacrifices some processing speed in exchange for stronger fraud-detection capability, while retaining a computational profile that is potentially suitable for real-time FinTech monitoring.

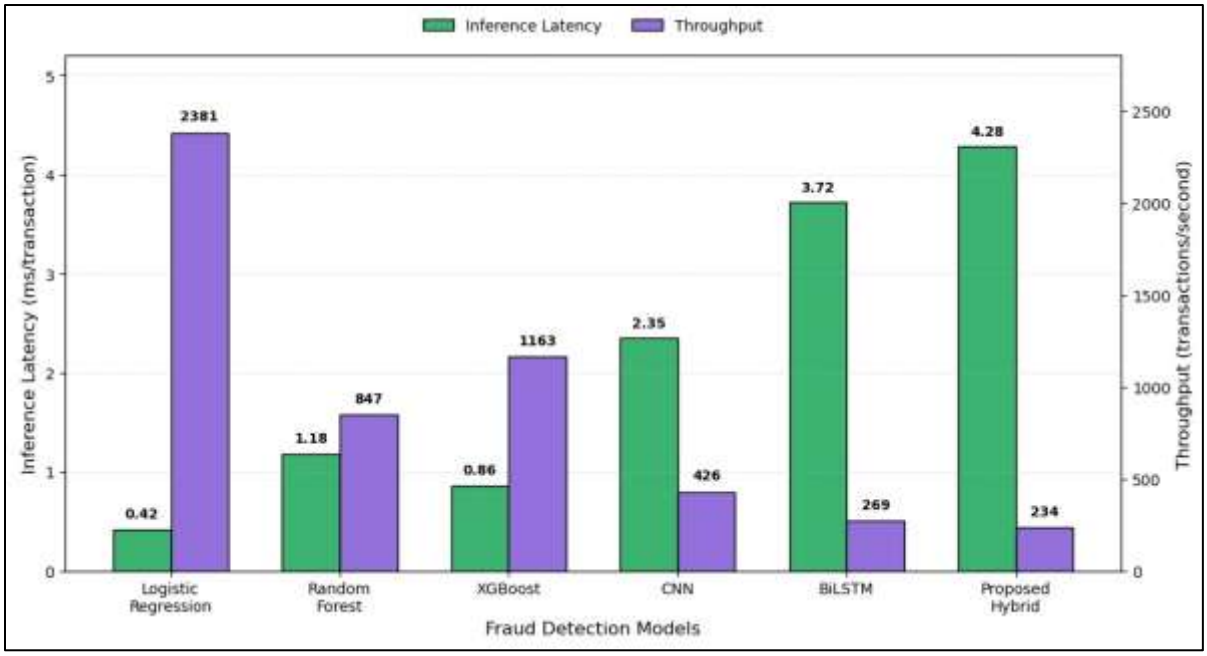

Figure 4. Inference Latency and Throughput Comparison

4.4 Comparative and Ablation Analysis

Table 2 confirms that each architectural component contributes progressively to fraud-detection performance. The CNN-only configuration achieved 91.8% Precision, 88.6% Recall, a 90.2% F1score, a PR-AUC of 0.912, and an MCC of 0.889, showing that convolutional feature extraction can identify useful local transaction patterns but has limited ability to capture temporal behavior. The BiLSTM-only model improved sequential learning, reaching 92.4% Precision, 90.1% Recall, a 91.2% F1-score, a PR-AUC of 0.924, and an MCC of 0.901. Combining CNN and BiLSTM further increased performance to 94.1% Precision, 92.7% Recall, a 93.4% F1-score, a PR-AUC of

0.948, and an MCC of 0.927, demonstrating the benefit of integrating feature-level and temporal representations. The complete CNN–BiLSTM–attention model achieved the best results, with 95.6% Precision, 94.3% Recall, a 94.9% F1-score, a PR-AUC of 0.962, and an MCC of 0.944. These results indicate that the attention mechanism provides an additional improvement by emphasizing the most informative transaction representations, while the full hybrid architecture delivers the strongest and most balanced fraud-detection performance among all ablation configurations.

Table 2. Ablation Performance of the Proposed Hybrid Model

Model Configuration	Precision (%)	Recall (%)	F1-Score (%)	PR-AUC	MCC
CNN Only	91.8	88.6	90.2	0.912	0.889
BiLSTM Only	92.4	90.1	91.2	0.924	0.901
CNN + BiLSTM	94.1	92.7	93.4	0.948	0.927
CNN + BiLSTM + Attention	95.6	94.3	94.9	0.962	0.944

4.5 Integrated Performance and Practical Implications

The combined results demonstrate that effective FinTech fraud detection requires consideration of predictive quality, class imbalance, false alarms, and computational performance together rather than relying on a single metric. The proposed hybrid architecture benefits from CNN-based feature extraction, BiLSTM-based temporal

learning, and attention-guided representation, enabling stronger recognition of minority fraudulent transactions while maintaining control over false positives. PR-AUC provides particularly meaningful evidence under severe class imbalance, while MCC confirms balanced classification reliability across both classes. The latency and throughput analysis further determines whether improved predictive performance can be achieved without compromising real-time transaction processing. From a practical perspective, the framework is suitable as a computational basis for automated FinTech transaction monitoring, although operational deployment would still require evaluation using institution-specific data, evolving fraud patterns, concept drift, privacy constraints, and production-level transaction loads. These limitations should be considered before interpreting experimental performance as evidence of realworld banking deployment.

5. Conclusion

The proposed study presents a hybrid CNN–BiLSTM–attention framework for real-time FinTech fraud detection, combining local feature extraction, sequential transaction modeling, and attentionbased representation to improve identification of minority fraudulent transactions. The framework achieved 95.6% Precision, 94.3% Recall, and a 94.9% F1-score, together with a PR-AUC of 0.962 and an MCC of 0.944, indicating strong classification reliability under imbalanced transaction conditions. ROC-AUC analysis further supported effective discrimination between legitimate and fraudulent transactions, while reduced false-positive behavior improved the practical reliability of fraud alerts. From a computational perspective, the proposed model required 4.28 ms per transaction and achieved a throughput of 234 transactions per second, demonstrating that the additional complexity of the hybrid architecture remains compatible with rapid transaction-level monitoring under the evaluated conditions. Despite these results, the study is limited by dependence on the selected dataset and experimental environment. Future work should address evolving fraud patterns and concept drift, incorporate explainable AI, investigate online and continual learning, and validate the framework using real-world FinTech transaction streams and production-scale deployment conditions.

References

1. Alarfaj, F. K., Malik, I., Khan, H. U., Almusallam, N., Ramzan, M., & Ahmed, M. (2022). Credit card fraud detection using state-of-the-art machine learning and deep learning algorithms. Ieee Access, 10, 39700-39715.
2. Btoush, E. A. L. M., Zhou, X., Gururajan, R., Chan, K. C., Genrich, R., & Sankaran, P. (2023). A systematic review of literature on credit card cyber fraud detection using machine and deep learning. PeerJ Computer Science, 9, e1278.
3. Carcillo, F., Dal Pozzolo, A., Le Borgne, Y. A., Caelen, O., Mazzer, Y., & Bontempi, G. (2018). Scarff: a scalable framework for streaming credit card fraud detection with spark. Information fusion, 41, 182-194.
4. Cheng, D., Zou, Y., Xiang, S., & Jiang, C. (2025). Graph neural networks for financial fraud detection: a review. Frontiers of Computer Science, 19(9), 199609.
5. Dal Pozzolo, A., Boracchi, G., Caelen, O., Alippi, C., & Bontempi, G. (2017). Credit card fraud detection: a realistic modeling and a novel learning strategy. IEEE transactions on neural networks and learning systems, 29(8), 3784-3797.
6. Dal Pozzolo, A., Caelen, O., Le Borgne, Y. A., Waterschoot, S., & Bontempi, G. (2014). Learned lessons in credit card fraud detection from a practitioner perspective. Expert systems with applications, 41(10), 4915-4928.
7. Hachimi, K. E., Orhanou, G., & Ballihi, L. (2025). Credit card fraud detection model based on optimized Long Short-Term Memory: Imbalanced class solution and overfitting reduction. Expert Systems with Applications, 130871.
8. Jurgovsky, J., Granitzer, M., Ziegler, K., Calabretto, S., Portier, P. E., He-Guelton, L., & Caelen, O. (2018). Sequence classification for credit-card fraud detection. Expert systems with applications, 100, 234-245.
9. Mienye, I. D., & Jere, N. (2024). Deep learning for credit card fraud detection: A review of algorithms, challenges, and solutions. IEEE Access, 12, 96893–96910. https://doi.org/10.1109/ACCESS.2024.3426955
10. Motie, S., & Raahemi, B. (2024). Financial fraud detection using graph neural networks: A systematic review. Expert Systems with Applications, 240, 122156.
11. Randhawa, K., Loo, C. K., Seera, M., Lim, C. P., & Nandi, A. K. (2018). Credit card fraud detection using AdaBoost and majority voting. IEEE access, 6, 14277-14284.
12. Sun, Z., Li, H., Zhu, R., & Cheng, W. (2026). A hybrid CNN-LSTM-transformer model for window-based anomalous transaction detection in financial data. Scientific Reports.

13. Surendar, A. (2026). Geometry-Aware Dual-Timescale RIS Configuration for Indoor 6G IoT Links. VSF Journal of Electronics and Communication Engineering, 1(1), 1-7.
14. Tripathi, R. K., Parveen, P., Kanchana, K., Chikkegowda, K. G., & Gundapaneni, S. (2026). Automated Fraud Detection In Financial Services Using Hybrid Autoencoders And LSTM.
International Journal of Artificial Intelligence and Machine Learning, 6(4s), 412-425.
15. Van Vlasselaer, V., Bravo, C., Caelen, O., Eliassi-Rad, T., Akoglu, L., Snoeck, M., & Baesens, B. (2015). APATE: A novel approach for automated credit card transaction fraud detection using network-based extensions. Decision support systems, 75, 38-48.