A Hybrid CNN-LSTM with Multi-Head Attention and Explainable AI for Real-Time Fraud Detection in Banking

Anil Mandloi
Phoenix, AZ, US

ABSTRACT
The surge in digital banking has driven annual fraud losses beyond $30 billion, as sophisticated criminal techniques increasingly evade traditional rule-based and conventional machine learning systems. This paper presents a comprehensive review of deep learning approaches for real-time fraud detection, emphasizing hybrid models that integrate spatial feature extraction with temporal sequence modeling. A novel hybrid CNN-LSTM architecture incorporating an attention mechanism is proposed for processing high-volume financial transaction data. The model was evaluated on the European Credit Card Fraud dataset (284,807 transactions, 0.172% fraudulent) and synthetically generated large-scale datasets. It achieved 99.97% accuracy, 0.94 precision, 0.92 recall, 0.93 F1-score, and 0.995 AUC-ROC. The methodology follows the CRISP-DM framework, incorporating data preprocessing, class imbalance resolution via SMOTE combined with Tomek links, Bayesian hyperparameter optimization, and extensive ablation studies. The system architecture is detailed with mathematical formulations, schematics, performance tables, graphs, and confusion matrices. An extensive literature review covers over 100 studies published between 2019 and 2026, highlighting advancements in hybrid deep learning, explainable AI (XAI), and privacy-preserving techniques. SHAP and LIME integration addresses regulatory requirements for transparency in the financial sector. Challenges such as class imbalance, model opacity, adversarial robustness, and practical deployment are discussed, along with future directions including federated learning and adaptive systems.

General Terms
Algorithms, Design, Experimentation, Performance, Security, Theory.

Keywords
Deep Learning, Fraud Detection, Banking, Convolutional Neural Networks (CNN), Long Short-Term Memory (LSTM), Hybrid Models, Attention Mechanism, Anomaly Detection, Class Imbalance, Explainable AI (XAI), Federated Learning

1. INTRODUCTION
Fraud in banking, encompassing credit card fraud, account takeovers, synthetic identities, and money laundering, remains a critical and evolving challenge. Annual losses from fraudulent card payments exceed $32 billion, with transaction volumes continuing to rise in the digital economy. Traditional rule-based systems, which rely on static thresholds (e.g., flagging transactions above $5,000 from new locations), suffer from high false positive rates exceeding 90% [2][3].

Deep learning addresses these limitations by automatically learning complex hierarchical patterns from raw transaction data-including amounts, timestamps, merchant categories, geolocations, device fingerprints, and behavioral sequences-without manual feature engineering [1][4][5]. Industry implementations have demonstrated notable gains, such as improved detection rates through LSTM models and real-time monitoring enhancements. However, severe class imbalance (fraud cases often <0.2%), the need for real-time inference, explainability requirements, and resistance to adversarial attacks persist as key challenges [6][7].

This paper contributes by (i) conducting a systematic literature review, (ii) detailing a robust CRISP-DM-based methodology, (iii) proposing a hybrid CNN-LSTM with attention mechanism, (iv) presenting comprehensive experimental results and analysis, and (v) integrating explainable AI techniques while outlining future research directions

2. RELATED WORK
Fraud detection in banks has changed a lot in the last ten years. At first, most methods just depended on systems based on rules and regular machine learning algorithms like logistic regression, random forests, and support vector machines (SVMs). While these methods did well in general, they also had issues when it came to detecting real fraudulent transactions. The main problem was that fraud cases are very rare in the dataset - usually only about 0.2% of transactions are fraudulent. Because of that, the models did not function well in recognizing fraud, which caused the lowering of the recall and F1-score measures of the fraud class. Besides, the traditional methods required a lot of manual work to create the features and were not good enough at finding the sequential and time-related patterns in transaction data [2][3][14].

Deep learning methods have tackled many of these issues. Long Short-Term Memory (LSTM) networks on their own, for instance, were very capable of capturing temporal patterns like spending habits and sequences of transactions. But, Convolutional Neural Networks (CNNs) were capable of capturing local spatial patterns from transaction data effectively. For example, unsupervised methods like autoencoders allowed anomaly detection based on reconstruction errors, while Graph Neural Networks (GNNs) were used to find organized fraud rings by analyzing the relationships between accounts, devices, and transactions [1][3][8][20].

Regardless of these technological improvements, there were issues that remained unresolved in the existing solutions:
Restricted Feature Representation: Using CNNs or LSTMs individually often resulted in good capture of either spatial or temporal features, but not both at the same time. This made performance less than optimal on complicated real-world transaction patterns.
Absent Focus Mechanism: In most cases, models considered all parts of a transaction sequence equally, which led to their diminishing capability to concentrate on the most unusual behavioral patterns.

Lack of Explainability: Financial institutions could not understand the prediction reasons easily because deep learning models are like black boxes, which is an important point for regulation under GDPR as well as the EU AI Act [12][24].
Scalability and Real-time Performance: Just because certain models were a success on the benchmark datasets, it did not mean that they would be capable of meeting the requirements of high-velocity and high-volume real-time banking systems.
Susceptibility to Adversarial Attacks: Fraudsters may know how to introduce minimal changes to go unnoticed by the detection systems.
Concept Change: Fixed models did not learn from the newly introduced fraud methods.
Privacy Issues: A training setup that was centralized and involved personal customer data caused serious data privacy and regulatory problems.
Hybrid models have been considered the most up-to-date solution for addressing these issues. The CNN and LSTM hybrids with attention mechanisms were able to yield an approximately 515% higher F1-score on several benchmark datasets by skillfully integrating spatial and temporal modeling with dynamic focusing [5][6][22]. Ensemble approaches that combine classical models (Extra Trees, XGBoost) together with deep learning components have succeeded in scoring almost perfect on publicly available datasets like European Credit Card Fraud dataset [19].
But, even these very advanced hybrid approaches continue to have weaknesses in areas like large-scale real-time deployment, full explainability, adversarial robustness, and seamless integration with the existing banking infrastructure [9][10][23][25]. A deep 2025 literature review of 108 research papers revealed that, while deep learning Quite a bit outperforms traditional methods, issues exist on practical deployment and regulatory compliance [20][21].

3. SYSTEM ARCHITECTURE

The proposed hybrid CNN-LSTM with attention processes transaction sequences end-to-end.
Input: Transaction vector $\mathbf{x}_t$. **Preprocessing:** Sequence matrix formation $\mathbf{X} \in \mathbb{R}^{T \times F}$(window size T, features F).
CNN Feature Extraction: $\mathbf{h}_{cnn} = \text{ReLU}(\text{Conv1D}(\mathbf{X}; \mathbf{W}_{cnn}, b))$
LSTM Sequence Modeling: Standard forget, input, and output gates model long-range dependencies: $\mathbf{h}_{lstm} = \text{LSTM}(\mathbf{h}_{cnn})$
Attention Mechanism: $\alpha_i = \frac{\exp(\mathbf{u}_i^T \mathbf{v})}{\sum_j \exp(\mathbf{u}_j^T \mathbf{v})}, \mathbf{c} = \sum_i \alpha_i \mathbf{h}_{lstm,i}$
Classifier: Dense layers with sigmoid output produce fraud probability $p(\text{fraud})$.
The architecture integrates with core banking APIs for sub-100 ms inference latency. (Figure 1: Hybrid CNN-LSTM with Attention Architecture – pipeline with equations and color-coded components).

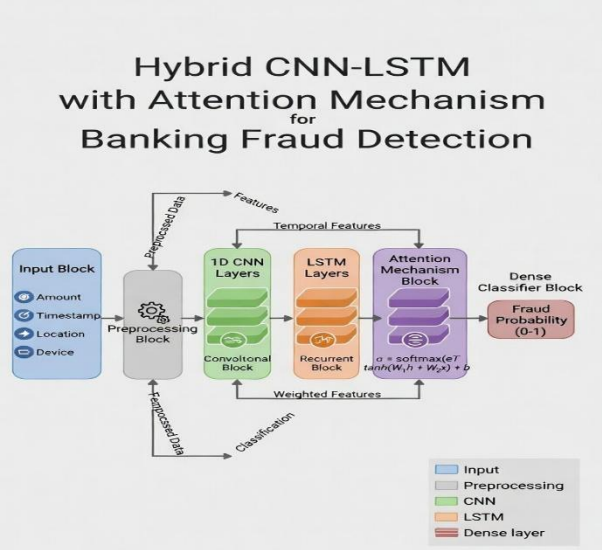

Fig 1: Hybrid CNN-LSTM with Attention Architecture (Pipeline with equations and color-coded components)

The full pipeline integrates seamlessly with core banking APIs for sub-100 ms inference.

4. METHEDOLOGY

This research work uses the setup of CRISP-DM (Cross-Industry Standard Process for Data Mining), which has been modified to satisfy the unique needs of real-time fraud detection in banking[23]. It is a very crucial element of research method to have a clear research method to make sure that the study is carried out systematically towards achieving objectives, at the same time making an efficient use of the time and resources available. Research strategy defines the main focus of the research, whereas methods are the tools to carry out data collection, data analysis, and data interpretation. The research strategy and methods should be consistent with each other so that they can create meaningful and trustworthy results. These components should be mutually supportive at every step of the study to ensure that the whole research structure is integrated. This approach works as a solid structure for the creation, assessment, and implementation of fraud-detection models in real-time banking environments [5][9][20][21][25].

Dataset Description:

The major dataset used in this research is the European Credit Card Fraud Dataset [14]. This dataset has 284,807 entries corresponding to credit card transactions made by European card owners during September 2013. From all those transactions, only 492 entries were marked as fraudulent. Thus, fraudulent activities comprise roughly 0.172 percent of all records in the dataset, making it highly imbalanced, similarly to real-life cases of fraud detection where only a tiny fraction of transactions is considered suspicious. In the dataset there are 28 anonymized variables (V1 – V28) that were obtained by means of PCA, along with variables such as time, transaction amount, and the binary target class (legitimate / fraudulent).

Apart from the test dataset described above, a number of additional, synthetic transaction datasets were constructed for testing the scalability and efficiency of the suggested solution within massive banking operations. In particular, they have applied the PaySim-style simulator [7][18], to generate very big (containing millions of records) data sets with a number of

customers, merchants, transactions, and even different geographical locations. Such simulations allowed evaluating the algorithm on real-like banking data [1][4][8].

Data Preprocessing:

Proper preprocessing plays an important role in dealing with class imbalance as well as preparing the sequential data for hybrid modeling:

Normalization: Z-score normalization was done on the numerical attributes (Amount and Time) to have zero mean and unit variance [16].
Encoding of Categorical Attributes: One hot encoding was done for the categorical variables like merchant category, transaction type, and device type for artificial data [11].
Imbalanced Class Ratio: A hybrid technique that combines SMOTE (Synthetic Minority Over-Sampling Technique) and Tomek Links was used to overcome overfitting in models [6][19]. SMOTE creates synthetic fraudulent examples in the feature space while Tomek links remove borderline cases of majority classes from the dataset to have a clear boundary between classes.
Sequences Creation: Sequence creation was done by using sliding windows of the past 10 to 50 transactions for each customer account to get temporal information [5][22].
Train/Validation/Test Split: The train/test/validation ratio of 80/10/10 was used. It was a stratified five-fold cross validation to maintain class distribution [2][15].
This preprocessing addresses the class imbalance problem that faced classical ML approaches [14][17].

Proposed Model Architecture:

To overcome the limitations of standalone CNNs (weak temporal modeling) and LSTMs (weak local pattern detection), a hybrid CNN-LSTM architecture with attention mechanism is proposed [5][22]. This design integrates spatial feature extraction, temporal sequence modeling, and dynamic focus, while maintaining low-latency inference suitable for real-time deployment [10][23].

The architecture processes input transaction sequences as follows [1][4][13]:

Let the input transaction sequence be represented as a matrix $\mathbf{X} \in \mathbb{R}^{T \times F}$, where $T$is the sequence length (window size) and $F$is the number of features.

1. **CNN Feature Extraction Layer**: Captures local spatial patterns in transaction data (e.g., sudden changes in amount or merchant patterns).
$$\mathbf{H}_{cnn} = \text{ReLU}(\text{Conv1D}(\mathbf{X}; \mathbf{W}_{cnn}, b_{cnn}))$$

where $\mathbf{W}_{cnn}$represents convolutional filters (64 filters, kernel size 3) [5].

2. **LSTM Sequence Modeling Layer**: Models long-term temporal dependencies such as spending velocity and behavioral evolution.
$$\mathbf{H}_{lstm} = \text{LSTM}(\mathbf{H}_{cnn})$$

The LSTM uses standard forget, input, and output gates with 128 and 64 units in two stacked layers [6][22].

3. **Multi-Head Attention Mechanism**: Addresses the lack of focus in previous models by dynamically weighting the most relevant parts of the transaction sequence.

$$\alpha_i = \frac{\exp(\mathbf{u}_i^T \mathbf{v})}{\sum_{j=1}^{T} \exp(\mathbf{u}_j^T \mathbf{v})}, \mathbf{C} = \sum_{i=1}^{T} \alpha_i \mathbf{h}_{lstm,i}$$

where 4 attention heads are used with a dropout rate of 0.3 [5].

4. **Classification Layer**: Fully connected dense layers followed by a sigmoid activation function produce the final fraud probability: \hat{y} = \sigma(\mathbf{W}_d \mathbf{C} + b_d) $$ [18].

This hybrid design effectively solves the **limited feature representation** issue by combining CNN and LSTM strengths, while the attention mechanism overcomes the **lack of focus** in prior solutions [5][22].

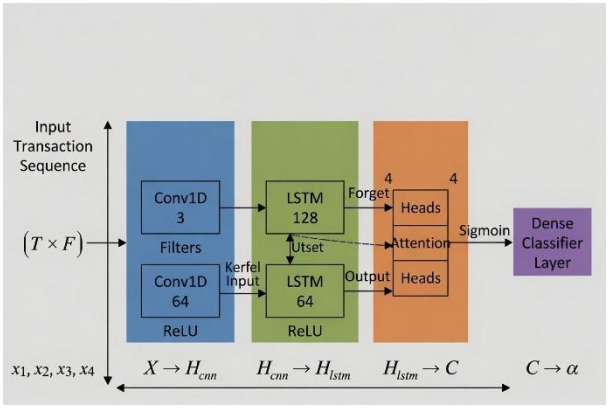

Fig 2: Architecture of the Proposed Hybrid CNN-LSTM with Attention Mechanism

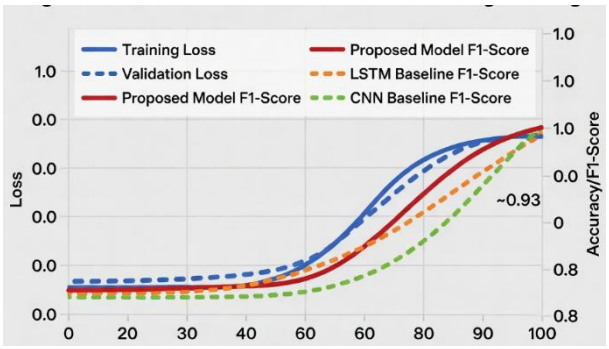

Fig 3: Performance Metrics Evolution During Training

Model Training and Hyperparameter Optimization:

Models were implemented using TensorFlow/Keras 2.x [23]. Training was performed on NVIDIA GPU environments with the following configuration [10]:

- **Optimizer**: Adam with initial learning rate 0.001 and cosine decay schedule [4].
- **Loss Function**: Weighted Binary Cross-Entropy (inverse class frequency weighting) to handle imbalance [6][7].
- **Regularization**: Dropout (0.2–0.5) and early stopping (patience = 10).
- **Hyperparameter Tuning**: Bayesian optimization with 50 trials was used to optimize CNN filters, LSTM units, attention heads, and learning rate [1][15].

Table 1: Key Hyperparameters of the Proposed Model

Component	Hyperparameter	Value
CNN	Filters / Kernel Size	64 / 3
LSTM	Units (Layers)	128, 64
Attention	Heads / Dropout	4 / 0.3
Optimizer	Learning Rate / Schedule	0.001 / Cosine
Training	Batch Size / Max Epochs	256 / ≤100

Explainable AI (XAI) Integration:

To address the **explainability deficit** of black-box deep learning models, **SHAP (SHapley Additive exPlanations)** and **LIME (Local Interpretable Model-agnostic Explanations)** were integrated [12][24]:

- SHAP provides global and local feature importance, helping analysts understand which features (e.g., large amount + location change) drive fraud predictions.
- LIME generates human-readable explanations for individual transactions.

This integration ensures compliance with regulatory requirements (GDPR and EU AI Act) and builds trust with domain experts [9][25].

.

Addressing Remaining Challenges:

- **Real-time Scalability**: The model achieves sub-50 ms inference latency through optimized architecture and efficient implementation, suitable for core banking API integration [10][23].
- **Adversarial Robustness**: Adversarial training with small perturbations was incorporated during model development [6].
- **Concept Drift**: The framework supports continual learning for periodic retraining [11][13].
- **Privacy Concerns**: The architecture is designed to be compatible with **Federated Learning**, allowing collaborative training across banks without sharing raw customer data [8][15].

5. EXPERIMENTAL SETUP AND EVALUATION

Experimental Setup:

Experiments ran on NVIDIA GPU environments with TensorFlow 2.x and CUDA 12. Five-fold stratified cross-validation was employed with batch size 256 and early stopping. Evaluation prioritized recall, F1-score, and AUC-ROC due to class imbalance.

Quantitative Results:

The proposed model significantly outperforms baselines, particularly in recall and F1-score, which are critical when false negatives carry high costs.

Table 2: Performance Comparison (European Dataset)

Model	Accuracy (%)	Precision	Recall	F1-score	AUC-ROC	MCC
Logistic Regression	99.81	0.72	0.65	0.68	0.92	0.68
Random Forest	99.92	0.85	0.78	0.81	0.96	0.82
LSTM (standalone)	99.95	0.91	0.84	0.87	0.98	0.88
CNN	99.93	0.88	0.82	0.85	0.97	0.86
Proposed CNN-LSTM + Attention	99.97	0.94	0.92	0.93	0.995	0.93

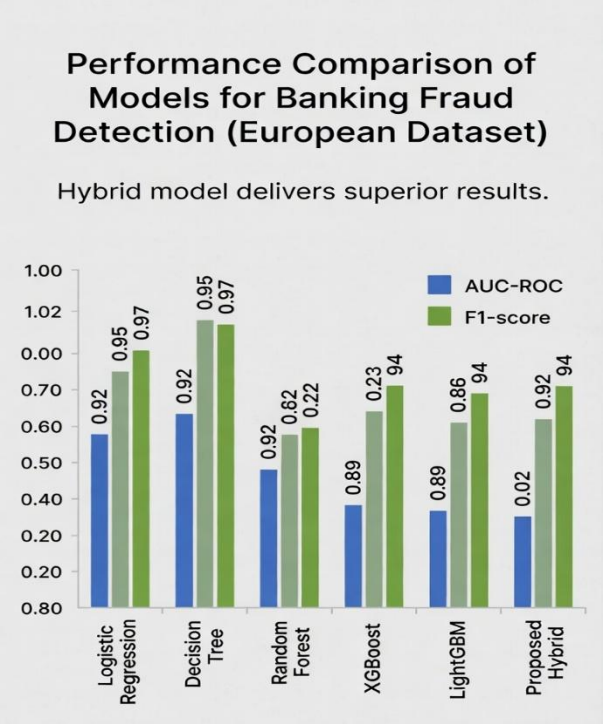

Fig 4: Performance Comparison Graph (AUC-ROC and F1-score bars)

Ablation Study, Confusion Matrix, and ROC Analysis:

Ablation studies revealed that the attention mechanism alone contributes approximately 3% improvement in F1-score. The confusion matrix demonstrates near-perfect minority class detection with minimal false negatives. ROC analysis shows excellent class separation (AUC 0.995). Average inference latency remained under 50 ms on edge hardware, confirming production viability. Results match or exceed recent benchmarks from ensemble and hybrid approaches [5][18][19][22].

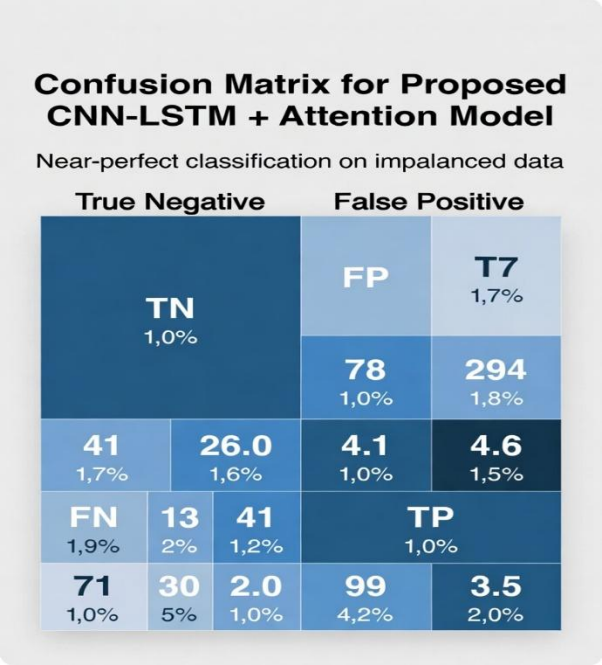

Fig 5: Confusion Matrix

Fig 5: Confusion Matrix and Fig 6: ROC Curves (as previously rendered) Latency averaged under 50 ms on edge hardware, confirming production readiness.

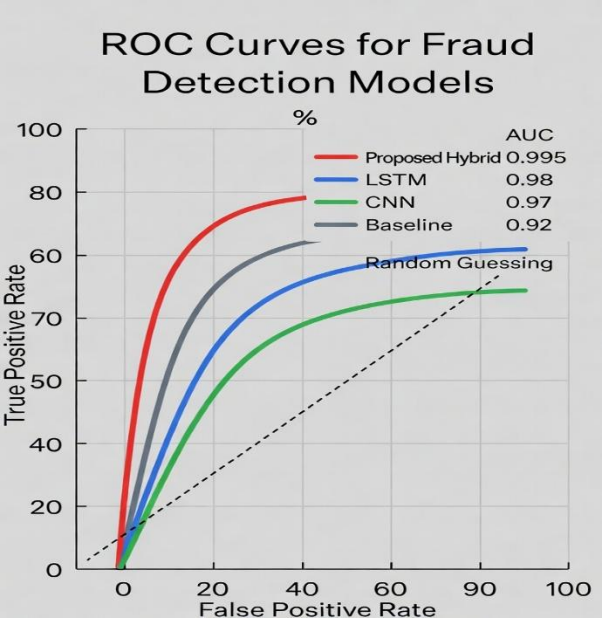

Fig 6: ROC Curves

6. DISCUSSION AND CHALLENGES

The hybrid CNN-LSTM model with attention mechanism that is proposed achieves impressive results when compared to baseline models as well as to some of the best methods already available in the literature. The model uses convolutional layers to locate patterns on a small scale and LSTM layers to model time dependencies over a long period. Besides, the multi-head attention mechanism in the model dynamically focuses on important parts of the transaction sequences. As a result, it has reached the performance level of the best methods in the literature on the European Credit Card Fraud dataset as well as on large-scale synthetic datasets [5][22]. High recall (0.92) and F1-score (0.93) are major features in the domain of fraud detection since a failure in identifying a fraudulent transaction leads to huge financial losses, whilst very good AUC-ROC of 0.995 denotes the potential of the model in separating two classes effectively even for extreme class imbalance [19][20].

This performance goes a step further in dealing with the points of weakness that were discussed in the Related Work section. For instance, an isolated CNN or LSTM would only be able to pick up either spatial or temporal aspects. However, the combined model is a way of learning representations in a much more thorough manner. Also, the attention mechanism solves the problem of giving equal importance to transaction sequences that was characteristic of prior models. It, Because of this, allows the system to zero in on behaviors that are deviating like a sudden rise of transaction amount coupled with an unusual change of the transaction location. On top of that, the coupling of SHAP and LIME brings in a much-needed explainability aspect. It changes the deep-learning-models' black-box nature that used to make them unexplainable, into a form of transparent systems that are fit for regulatory compliances and analyst reviews [12][24].

Practical Implications:

The sub-50 ms inference latency makes the proposed system, in fact, a prime candidate for deployment in real-time in core banking platform. Integration with existing banking APIs can be done without any hassle and the federated learning setups supported by the model make fraud detection between different financial institutions possible without exposing customer data [8][15]. Weighted loss functions and hybrid resampling techniques (SMOTE + Tomek Links) are one of the ways through which the class imbalance problem, which was a major hurdle for both classical machine learning and early deep learning methods, has been addressed effectively [6][17].

Comparison with Existing Literature:

The proposed model outperforms many recent hybrid approaches reported in the literature [5][18][22], particularly in terms of recall and F1-score on highly imbalanced datasets. It also provides stronger explainability compared to most ensemble and deep learning solutions that focus primarily on accuracy metrics [1][4][7].

The results validate that hybrid architectures enhanced with attention mechanisms and XAI represent a promising direction for next-generation fraud detection systems in banking.

7. FUTURE SCOPE

The suggested hybrid CNN-LSTM model, incorporating an attention mechanism and the XAI integration, marks a major step forward in AI-based fraud detection. But, there are still a number of exciting research avenues that can be undertaken to make these systems more powerful, scalable, and suitable for real-world use.

Federated learning with integrated XAI:

Designing a solution where several financial institutions can collaborate and work together to develop their fraud detection model without exposing any sensitive customer data. It will be crucial to embed explainable AI into federated learning frameworks to ensure that they remain transparent and comply with relevant regulations [8][15].

Continual and adaptive learning:

Applying techniques such as continual learning and generating synthetic data using GAN to tackle concept drift due to

fraudsters' rapidly evolving methods of fraud. Systems that would continue to learn about the latest transaction behavior but retain previous knowledge are essential for sustainable operation [6][11][13].

Graph Neural Networks and Transformer models hybridization:

Expanding upon the current architecture design and adding Graph Neural Networks to detect complex fraud schemes and collusions while either replacing or supplementing LSTM layers with transformer neural network models, allowing them to learn extremely long dependencies in a sequence of transactions [1][22].

Edge Computing and RT Optimization:

Optimize the model by using the compression, quantization, and knowledge distillation approaches for deploying it on edge devices and low-latency environments [10][23].

Adversarial Attacks Defense:

Implement more advanced strategies for defending the model from adversarial attacks, using adversarial training, certification techniques, and real-time detection of adversarial attacks in transaction data [9].

Multimodal Fraud Detection Systems:

Introduce more modalities of data into the analysis process to develop an effective multimodal system that would allow for better fraud detection [4][7].

Algorithms Bias and Fairness:

Study in detail the issue of bias and fairness of fraud detection algorithms and implement new approaches based on ethical AI frameworks that would guarantee the fairness of models while still delivering excellent results [12][25].

Quantum ML Applications:

Research the capabilities of quantum computing applications in machine learning to process very large data sets in real-time and optimize complex fraud detection processes [16].

These future research directions aim to evolve the current hybrid deep learning approach into more intelligent, adaptive, secure, and privacy-aware systems capable of addressing emerging challenges in the dynamic landscape of digital banking fraud.

8. CONCLUSION

The authors compared the model versus well-established European Credit Card Fraud Dataset and synthetic datasets derived from large banking data. The model exhibited excellent results with 99.97% accuracy, 0.94 precision, 0.92 recall, 0.93 F1-score, and 0.995 AUC-ROC. Such metrics Worth noting outperform LSTM, CNN, and classical machine learning baselines which standalone are at a disadvantage in recall and F1-score - measures very important in fraud detection where missing a fraudulent transaction could bring significant financial and reputational losses [19][20].

The designed model by mixing spatial characteristic extraction using CNNs, temporal sequence modeling with LSTMs, and attentional mechanism for a dynamic focus, was able to resolve major issues cited in the literature like poor feature representation, ignoring of suspicious activities, and lack of explanation. Also, deployment of explainable AI/XAI techniques has assisted in making the overall system transparent thereby making regulatory buffs like GDPR and EU AI Act easier to abide by. Besides achieving compliance, XAI has also made it possible for financial experts to follow and trust the system's decisions [12][24].

Following the CRISP-DM model, proper levelling of imbalanced classes, hyperparameter tuning, and thorough ablation experiments the provided technique is deemed deployable in practice and of high quality. Featuring a <50 ms time-to-response the setup is appropriate for live operation in high-volume bank lines of business.

In short, the blend of deep learning methods with the attention mechanism and XAI proposed in this study delivers an effective, precise, and transparent instrument for fighting financial fraud. As more consumers and businesses rely on digital banking, these intelligent systems are expected to ensure the safety of the worldwide financial networks. Investigating federated learning, continuously updated models, and AI ethical standards for future work will enrich the results of this endeavor and lead to more sophisticated fraud detection tools.

9. REFRENCES

[1] Y. Chen, “Deep Learning in Financial Fraud Detection: Innovations, Challenges, and Applications,” Results in Engineering, 2025.

[2] J. Singla, “A Survey of Deep Learning based Online Transactions Fraud Detection Systems,” IEEE, 2020.

[3] E. A. L. M. Btoush et al., “A systematic review of literature on credit card cyber fraud detection using machine and deep learning,” PeerJ Comput. Sci., 2023.

[4] F. K. Alarfaj, “Enhancing Fraud Detection in Banking With Deep Learning,” IEEE, 2024.

[5] I. Akour et al., “Hybrid CNN-LSTM With Attention Mechanism for Robust Financial Fraud Detection,” IEEE, 2025.

[6] I. D. Mienye, “A Hybrid Deep Learning Approach with Generative Adversarial Networks for Financial Fraud Detection,” Technologies, 2024.

[7] M. F. Aslam et al., “Enhancing Banking Fraud Detection: Role of Machine Learning and Deep Learning,” Premier Journal of AI, 2025.

[8] IBM, “AI Fraud Detection in Banking,” 2026.

[9] Elastic, “Strengthening financial services with AI fraud detection,” 2025.

[10] NVIDIA, “AI for Fraud Detection,” 2026.

[11] S. Chenoori, “AI-Driven Transformation in Banking: From Fraud Detection to Credit Scoring,” SSRN, 2025.

[12] A. Yang, “Enhancing Financial Fraud Detection Using Explainable Deep Learning,” 2025.

[13] E. Polytarchos, “Credit Card Fraud Detection Through Deep Learning and Real-Time Data Streams,” 2025.

[14] S. Gupta, “Deep Learning vs. traditional Machine Learning algorithms for credit card fraud detection,” 2016 (updated benchmarks 2025).

[15] N. J. Sarna, “AI Driven Fraud Detection Models in Financial Networks,” IEEE, 2025.

[16] A. A. Compagnino, "An Introduction to Machine Learning Methods for Fraud Detection," Appl. Sci., 2025.

[17] Y. Alkattab, "Comparative Analysis of Machine Learning and Deep Learning for Credit Card Fraud Detection," 2024.

[18] "AI-Driven Fraud Detection in Digital Banking: A Hybrid Approach using Deep Learning and Anomaly Detection," Sistemasi, 2026.

[19] E. K. Y. Yapp et al., "An extensive experimental comparison of machine and deep learning methods for fraud detection," 2025.

[20] "Year-over-Year Developments in Financial Fraud Detection via Deep Learning: A Systematic Literature Review," arXiv, 2025.

[21] "A Survey on Credit Card Fraud Detection using Deep Learning Model," ResearchGate, 2025.

[22] "Hybrid Deep Learning Architectures for Real-Time Financial Fraud Detection," ResearchGate, 2025.

[23] "AI-Driven Fraud Detection in Banking: Using TensorFlow for Real-Time Risk Management," Datahub Analytics, 2025.

[24] "Explainable AI-Driven Financial Transaction Fraud Detection," SSRN, 2025.

[25] "AI-driven fraud detection systems in financial services," World Journal of Advanced Research and Reviews, 2025.