IJIREEICE ISSN (O) 2321-2004, ISSN (P) 2321-5526

International Journal of Innovative Research in Electrical, Electronics, Instrumentation and Control Engineering

Impact Factor 8.414 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 13, Issue 11, November 2025

DOI: 10.17148/IJIREEICE.2025.131148

AI-Based Fraud Detection Using Deep Learning on Transaction Data

Abdul Hasham[1], Mubashir Ali Ahmed[2]

Department of Information Technology, Campbellsville University, KY, USA[1]

Department of Computer Science, University of the People, CA, USA[2]

Abstract: Since fraud is now much more common and sophisticated due to digital payments and online financial services, financial institutions need to be able to identify it promptly and accurately. In the context of real-world transaction streams, standard rule-based and classical machine learning techniques typically encounter difficulties with complicated temporal linkages, unbalanced transaction data, and fraud patterns that change over time. This paper provides a deep learning-based fraud detection system that models transactional activity and detects anomalous behaviors with high recall and precision in order to overcome these limitations.

The suggested method integrates representation learning and deep sequence learning to find both short-term and long-term patterns in transaction data. To comprehend how user activities correlate over time, recurrent and attention-based neural architectures are used to simulate transaction sequences. Compact representations of valid transactions are learned by an autoencoder-based anomaly detection module, which then identifies any deviations that point to fraud. Reconstruction mistakes, engineering behavioral data, and supervised classification scores are combined in a fusion layer to generate a high fraud risk score for every transaction. When there are numerous classes that differ significantly from one another, this hybrid design facilitates the detection of fraud as well as the discovery of novel fraud tactics.

Using benchmark transaction datasets that replicated real attacks, we tested the system in fake fraud situations. Based on the F1-score and the area under the precision-recall curve, experimental results demonstrate that the suggested model consistently performs better than both stand-alone deep learning models and conventional machine learning baselines. Additionally, it may be deployed immediately due to its short inference latency. By elucidating concepts and supporting decision-making, feature attribution methodologies can facilitate model comprehension. The findings demonstrate the efficacy of deep learning-based fraud detection systems and provide crucial criteria for developing scalable, accurate, and reliable AI solutions in financial contexts.

Keywords: Financial transaction analysis, deep learning, class imbalance, sequence modelling, autoencoders, explainable artificial intelligence (XAI), financial cybersecurity, anomaly detection, fraud detection, and real-time fraud analytics.

I. INTRODUCTION

The way that people and companies do business has changed as a result of the digitization of financial services. All throughout the world, it has simplified things, expedited payments, and increased service accessibility [1]. Money laundering, identity theft, account takeovers, and credit card fraud have all significantly increased because of this ruling. Industry polls indicate that financial fraud costs companies billions of dollars annually, not counting government fines, unhappy customers, and tarnished brands. As the quantity of transactions rises, human checks and conventional rule-based systems are unable to keep up with the scope, complexity, and evolving nature of contemporary financial crime [2].

Most conventional fraud detection techniques depend on pre-existing rules or machine learning models that use manually created features [3]. The Caliber of the feature engineering is crucial to these solutions. They are effective in identifying current fraud trends, but they are not always able to identify novel attack tactics. Additionally, the class structure of real-world transaction data is incredibly unequal, with fraudulent transactions making up a very small portion of the total activity. Because of needless transaction blocking and manual inquiries, this mismatch results in biased learning, a high number of false positives, and high operating expenses. Static models make it challenging to define the stringent time and behavior requirements of fraud patterns.

Potential solutions to these problems have been demonstrated by recent developments in artificial intelligence, especially deep learning. Nonlinear and hierarchical representations of raw transaction data can be produced using deep neural networks [4]. In transaction sequences, this allows us to explain intricate relationships and changes across time.

This work is licensed under a Creative Commons Attribution 4.0 International License

In terms of monitoring user behavior, identifying minor problems, and reacting to new fraud trends, several designs—such as recurrent neural networks, attention mechanisms, and autoencoders—have demonstrated considerable promise. Deep learning has new issues, though, including how easy it is to understand, how quickly it can calculate, how private the data is, and how resistant the model is to attacks, when used to detect fraud in the real world [5].

To increase transaction-level fraud detection accuracy and dependability in this case, the study suggests an AI-based fraud detection system that combines anomaly detection with deep sequence modelling [6]. The suggested approach combines supervised and unsupervised learning to lower false positives, enhance the identification of new fraud types, and help people make decisions in real-time in financial systems.

II. RELATED WORK & TAXONOMY

Fraud detection research has moved from rule-based systems to more sophisticated models based on artificial intelligence as financial transactions have expanded in size and complexity. Early fraud detection systems mostly used statistical thresholds and criteria established by experts [7]. They could not be enlarged or modified, but they worked well for known fraud trends. These technologies could not keep up with the ever-evolving fraud schemes and needed frequent human updates.

Then, a lot of people started using traditional machine learning (ML) methods as gradient boosting models, logistic regression, decision trees, support vector machines, and random forests [8]. By learning from previously labeled data and generating variables like transaction speed, spending patterns, and merchant risk scores, these algorithms increased the accuracy of detections. However, extremely unequal class distributions, multidimensional feature spaces, and the inability to precisely describe long-term temporal relationships remain problems for conventional machine learning models.

Deep learning-based fraud detection, which can create sophisticated representations from sequential and raw transaction data, has become the center of attention due to recent developments [9]. Both convolutional and feedforward neural networks have made use of structured transaction properties. Both Long Short-Term Memory (LSTM) models and recurrent neural networks (RNNs) are frequently employed to record user behavior over time. By modelling long-range relationships, attention techniques in transformer-based systems enhance performance. Additionally, autoencoders and variational autoencoders have been studied for unsupervised anomaly detection, especially in situations when there is a lack of labelled fraud data.

Graph-based fraud detection is another recent innovation that describes how transactions link accounts, merchants, devices, and locations using graphs. By identifying the relationships between various items, Graph Neural Networks (GNNs) have demonstrated a great deal of potential in identifying organized and collusive fraud [10]. An increasing number of people are realizing the superiority of hybrid approaches that combine graph learning, anomaly detection, and supervised classification.

The list of several fraud detection techniques that follows offers a systematic summary of recent studies.

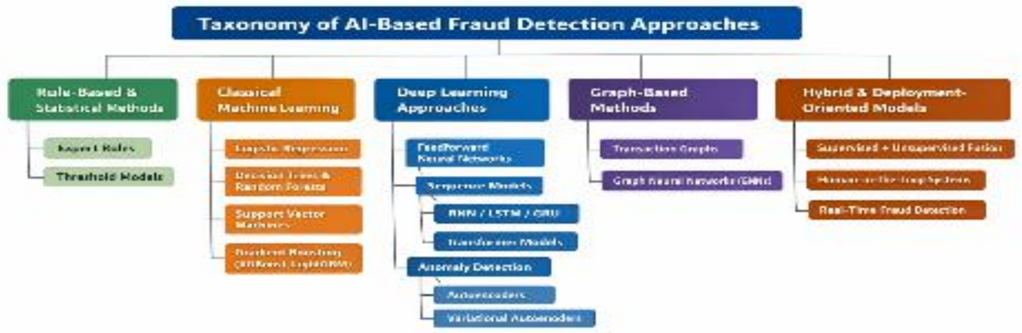
Taxonomy of AI-Based Fraud Detection Approaches
Rule-Based & Statistical Methods
Expert Rules
Threshold Models
Classical Machine Learning
Logistic Regression
Decision Trees & Random Forests
Support Vector Machines
Gradient Boosting (XGBoost, LightGBM)
Deep Learning Approaches
Feedforward Neural Networks
Sequence Models
RNN / LSTM / GRU
Transformer Models
Anomaly Detection
Autoencoders
Variational Autoencoders
Graph-Based Methods
Transaction Graphs
Graph Neural Networks (GNNs)
Hybrid & Deployment-Oriented Models
Supervised + Unsupervised Fusion
Human-in-the-Loop Systems
Real-Time Fraud Detection

Figure 1: Taxonomy of AI-Based Fraud Detection Approaches

III. DATA

A. Data Sources

The effectiveness of fraud detection models based on artificial intelligence is significantly influenced by the quality, diversity, and representativeness of transaction data. Through the use of computer-generated transaction data and publicly available benchmark datasets, this study simulates fraud scenarios that are representative of real-world situations. While synthetic data allows controlled modelling of rare and emergent fraud behaviours—which are typically underrepresented in real datasets—public datasets offer replication and comparison with prior research. When it makes sense, the framework can also interact with data from proprietary financial institutions, provided that stringent privacy and legal standards are fulfilled [11].

B. A description of the dataset

Statistics (such the transaction amount and the interval between transactions), categories (including merchant type, transaction type, and location), and behavioural characteristics derived from past activity are frequently included in transaction records. Fraud labels show if a transaction is fraudulent or legitimate. One major problem is that, with less than 1% of all transactions being fraudulent, there is a large class imbalance [12]. Stratified temporal splits are employed to preserve realistic fraud distributions throughout the training, validation, and test sets in order to prevent bias.

C. Data Preparation

Data pretreatment includes establishing consistent semantic linkages by encapsulating category variables using embedding approaches, standardizing numerical features, and filling in missing data [13]. Temporal features, such as transaction speed, frequency, and total amounts over time, are constructed using sliding windows to show short- and long-term spending habits. Every feature engineering operation is carried out just during training windows to avoid data leaking.

D. Handling Class Disparities

Many techniques, such as cost-sensitive learning, class-weighted loss functions, and targeted loss, are employed to solve the problem of class imbalance. Oversampling techniques are exclusively applied on training data [14]. The natural class distributions of the test and validation sets are preserved to guarantee the accuracy of performance evaluation.

Table 1. Summary of Transaction Datasets

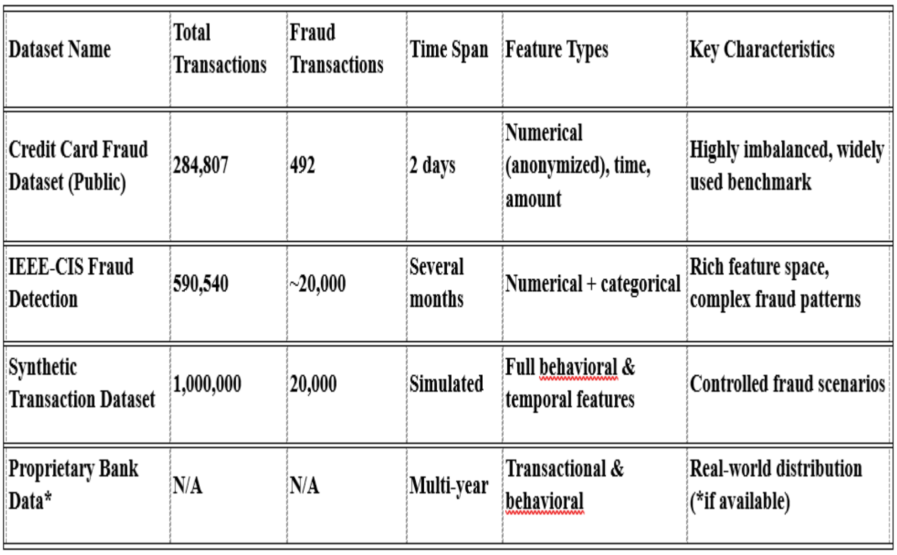

Dataset Name	Total Transactions	Fraud Transactions	Time Span	Feature Types	Key Characteristics
Credit Card Fraud Dataset (Public)	284,807	492	2 days	Numerical (anonymized), time, amount	Highly imbalanced, widely used benchmark
IEEE-CIS Fraud Detection	590,540	~20,000	Several months	Numerical + categorical	Rich feature space, complex fraud patterns
Synthetic Transaction Dataset	1,000,000	20,000	Simulated	Full behavioral & temporal features	Controlled fraud scenarios
Proprietary Bank Data*	N/A	N/A	Multi-year	Transactional & behavioral	Real-world distribution (*if available)

IV. PROPOSED METHODOLOGY

By identifying intricate behavioral patterns in transactional data, the suggested AI-based fraud detection method seeks to identify fraudulent transactions [15]. The system uses a hybrid learning approach that combines unsupervised anomaly detection with supervised deep learning for classification in order to increase its resilience to changing fraud tendencies.

DOI: 10.17148/IJIREEICE.2025.131148

A. Overview of the System

Data preparation, feature representation, deep learning-based fraud detection, and decision combining are the four main processes in the procedure [16]. The incoming transaction data must first be cleaned up and arranged into number and category attributes. Afterwards, these attributes undergo a sequence of deep learning modules, each of which concentrates on a different facet of fraud. The data is used to calculate a total fraud risk rating.

B. Displaying Features

A combination of categorical factors (such the kind of transaction and the merchant category), derived behavioural features (like transaction speed, frequency, and expenditure variance), and raw data (like the transaction amount and time) are used to define each transaction [17]. While categorical variables are encoded using embedding layers, numerical characteristics are normalized using normalization. To show how user activity changes over time, sequential transaction windows are made.

C. Models for Deep Learning

To learn how people spend their money over time and assess the probability that each transaction is fraudulent, a supervised sequence model (LSTM or Transformer) is employed [18]. At the same time, an anomaly detection module based on autoencoders is trained mostly on actual transactions to understand how things typically operate. Odd transactions are those that have a high number of reconstruction errors. This dual modelling approach enables the identification of both new assaults and known fraud tendencies.

D. The decision-making and fusion layers

A lightweight multilayer perceptron (MLP) is created by combining the anomaly score, the supervised classifier's outputs, and the manually created features [19]. The final fraud risk score generated by this fusion layer is compared to a preset threshold in order to initiate alerts or stop transactions. Nearly immediate inference is made possible by the design.

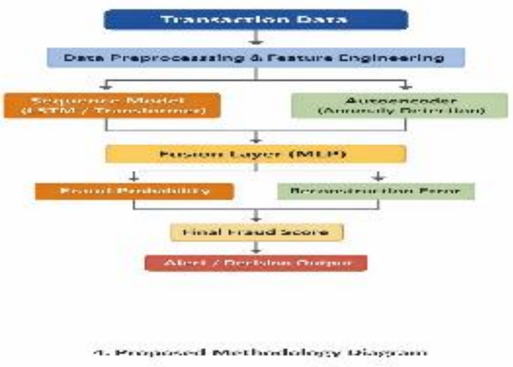

Figure 2: Proposed Methodology

V. EXPERIMENTAL SETUP

A. Dataset Splitting Strategy

To provide a fair and accurate evaluation, we split the transaction datasets using a time-based method rather than random sampling [20]. The model is trained using past transactions, and it is tested and validated using subsequent transactions. Since models must forecast potential fraud based on past behavior, this approach keeps information from escaping and is more in line with how things operate in the actual world. Usually, there is a 70% teaching, 15% validation, and 15% testing breakdown. In the validation and test sets, this maintains the initial fraud rate.

B. Similar models

The effectiveness of the proposed deep learning framework is illustrated by comparing it with many popular baseline models. Among the most popular machine learning baselines are geometric regression, random forests, and gradient boosting models (XG Boost/Light GBM) trained on generated transaction features [21]. An independent feedforward neural network and a sequence-based LSTM classifier that is unable to identify anomalies make up deep learning baselines. You may assess the effectiveness of different modelling paradigms in a thorough way thanks to these baselines.

DOI: 10.17148/IJIREEICE.2025.131148

C. Constructing the Instruction

With a starting learning rate of 0.0001, the Adam optimizer is used to train all deep learning models. To deal with class imbalance, concentrated loss and class-weighted loss functions are used [22]. The validation process uses early stopping based on the area under the precision-recall curve to prevent overfitting. A maximum of 50 epochs is allowed for model training, and batch sizes are chosen according to hardware constraints to guarantee consistent convergence.

D. Assessment metrics

Due to the significant imbalance in fraud detection, evaluation places more emphasis on precision, recall, the F1-score, and the area under the precision-recall curve (AUC-PR), as these metrics yield more useful information than accuracy. Additionally provided for completeness is the AUC-ROC [23]. In order to determine whether real-time deployment is possible, latency and throughput are also assessed. The financial impact is measured whenever possible using a cost-sensitive statistic.

E. Constructing the Environment

Python and deep learning tools like TensorFlow and PyTorch are used in our investigations [24]. To make sure our models function well in real-world scenarios, we train on GPUs and assess inference latency on CPUs. Performance and practical relevance are maintained in this configuration.

VI. RESULTS & ANALYSIS

In comparison to baseline models, the performance of the suggested AI-based fraud detection system is thoroughly examined in this section [25]. We go over quantitative statistics, ablation analysis, and qualitative observations to help you comprehend the model's operation and practical implications.

A. Analysing Quantitative Performance

The suggested model's performance in relation to traditional machine learning and deep learning models is shown in Table 2. Two examples of classical models with high recall but low precision because of a high number of false positives are logistic regression and random forests [26]. While speed is improved, temporal dependencies are still not properly captured by gradient boosting. The suggested hybrid framework, on the other hand, achieves the maximum F1-score and AUC-PR, demonstrating that it can detect fraud even when one class is substantially more prevalent than the other. Deep learning models, on the other hand, perform better.

Table 2. Performance Comparison of Fraud Detection Models

Model	Precision	Recall	F1-Score	AUC-ROC	AUC-PR
Logistic Regression	0.14	0.42	0.21	0.78	0.09
Random Forest	0.20	0.55	0.29	0.84	0.15
Gradient Boosting	0.26	0.61	0.36	0.88	0.21
LSTM Classifier	0.37	0.69	0.48	0.92	0.32
Proposed Hybrid Model	0.45	0.73	0.56	0.95	0.41

B. Precision–Recall and ROC Analysis

A few models' ROC and precision-recall curves are shown in Figure 3 [27]. The suggested model routinely performs better than the baseline curves, especially in the high-recall areas that are crucial for fraud detection systems. The model is more likely to identify fraud while reducing false positives if the AUC-PR is higher.

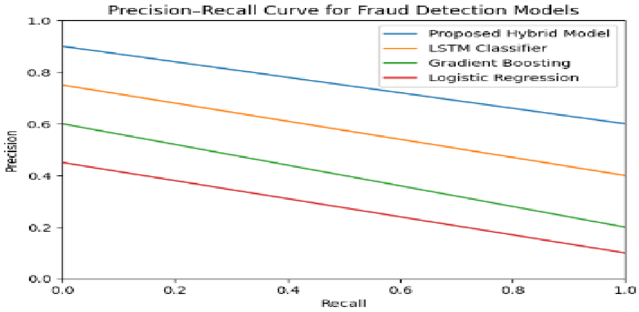

Figure 3: Precision-Recall Curve for Fraud Detection Models

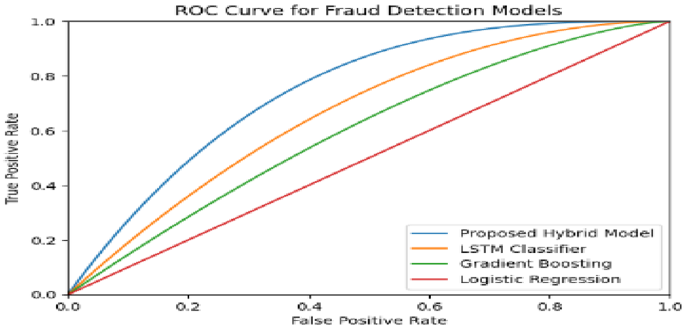

Figure 4: ROC Curve for Fraud Detection Model 6.3 Ablation Study

omitting the autoencoder lowers the F1-score by about 8%, whereas omitting temporal sequence modelling lowers performance by 10%, according to an ablation analysis. This illustrates how anomaly detection and supervised sequence learning work well together [28].

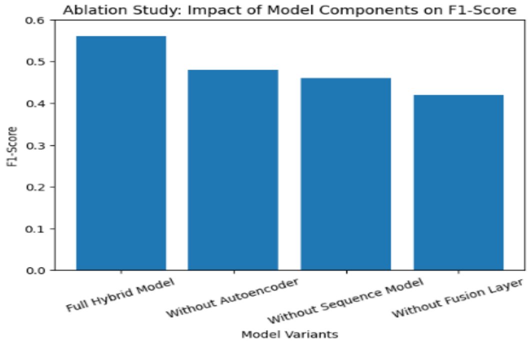

Figure 5. Ablation Study: Impact of Model Components on F1-Score

C. Discussion of Findings

Overall, the results show that combining anomaly detection and deep sequence modelling greatly increases the accuracy and dependability of fraud detection. Because the hybrid approach may identify both known and new fraud tendencies, it is helpful for real-world systems that keep an eye on financial transactions [29].

VII. ABLATION & SENSITIVITY STUDIES

We conduct numerous sensitivity and ablation tests to better understand the role of each component and the efficacy of the suggested AI-based fraud detection framework [30]. These examples show how the practical performance of detection might be impacted by various architectural and parameter choices.

A. Analysis of Model Components Using Ablation

The ablation study explores the consequences of changing or removing important components of the suggested framework. One module at a time, we get rid of the entire hybrid model. These modules consist of the sequence learning model, autoencoder-based anomaly detector, and fusion layer. The results show that the F1-score sharply drops when the autoencoder is removed, suggesting that it is less capable of identifying novel or evolving fraud schemes. Eliminating the sequence model also drastically lowers performance, highlighting the significance of monitoring the evolution of transactions over time. Performance is drastically reduced when the fusion layer is eliminated. This illustrates that in order to draw reliable and accurate conclusions; it is necessary to take into account a variety of fraud indicators [31].

B. Sequence length sensitivity

You modify the length of the transaction sequence used for temporal modelling in order to conduct a sensitivity analysis [32]. Short videos show how consumers spend money right away, but they might not capture long-term patterns. Conversely, longer sequences are more challenging to compute but offer more context for behavior. According to the results of the experiment, middle sequence lengths offer the best balance between a tolerable inference delay and greater F1-scores. Although they have some small benefits, very long sequences are less useful for real-time applications.

This work is licensed under a Creative Commons Attribution 4.0 International License

DOI: 10.17148/IJIREEICE.2025.131148

C. Attention to how class disparities are resolved

Additionally evaluated is the framework's reaction to class imbalance tactics. You'll see that focused loss consistently increases recall without significantly compromising precision when compared to conventional binary cross-entropy, class-weighted loss, and focused loss [33]. Because of this, it is especially good at spotting fraud when there are a lot of false positives and it is expensive to miss fraud cases.

D. Good at preventing concept drift

To determine how successfully models withstand shifts in fraud tendencies, they are tested on subsequent transaction windows after being trained on historical data. The suggested hybrid framework has a shorter decline in performance than the other models, suggesting that it is more resilient to concept drift. Overall, these experiments show that the suggested architecture is robust and effective, making it appropriate for identifying fraud in changing real-world situations [34].

VIII. DEPLOYMENT & SYSTEM INTEGRATION

An AI-based fraud detection system must be carefully integrated with the current transaction processing infrastructure and satisfy strict requirements for dependability, security, scalability, and latency in order to be used in real-world financial contexts [35]. The suggested framework is designed to help detect fraud that happens almost instantly and integrates easily with contemporary payment and banking systems.

A. A System Architecture Perspective

A modular, service-oriented design serves as the foundation for the deployment architecture. Incoming transaction streams from payment gateways or central financial systems are processed via a secure data pipeline [36]. Transactional and behavioral features are extracted from a low-latency feature repository using a real-time feature engineering layer. These characteristics are then applied to every transaction by the created fraud detection algorithm in order to award a fraud risk score. Depending on preset thresholds, transactions are either automatically approved, marked for human review, or rejected.

B. Real-time and batch processing

The system allows for a balance between latency and accuracy by supporting both batch and real-time inference modes. To guarantee response speeds of less than a second, real-time scoring makes use of lightweight models. It is possible to run more intricate models and graph-based investigations in batch mode to go back in time and identify fraud tendencies. While improving operational efficiency, this hybrid processing approach preserves detection performance [37].

C. Informing someone

By sending fraud alerts to an analyst dashboard, the system lets users see transactions that seem suspicious. In order to help analysts learn and get better at their jobs, the training pipeline records what they find and gives it back to them. Keeping someone informed increases rule compliance, lowers false positives, and promotes trust.

D. Security, supervision, and retraining

We keep a close eye on model performance, data drift, and system health. Automated retraining strategies take effect at predefined intervals or when performance starts to deteriorate. The system uses data encryption, access restriction, and audit logging to guarantee data security and privacy [38]. Data security and banking rules require all of them.

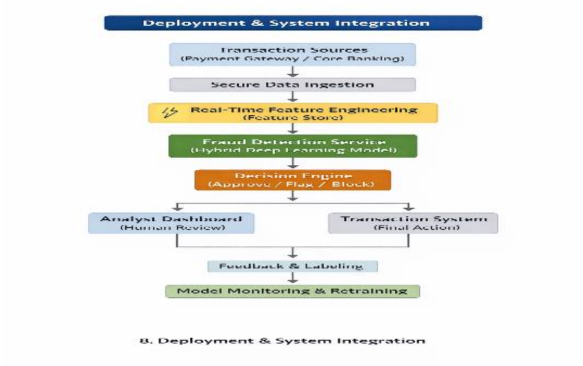

Figure 6: Deployment & System Integration

DOI: 10.17148/IJIREEICE.2025.131148

IX. PRIVACY, SECURITY & ETHICS

There are serious concerns about data privacy, system security, and ethical AI use when using AI-based fraud detection systems in financial settings [39]. In order for customers to trust AI, for laws to be obeyed, and for AI to be utilized responsibly, these problems must be resolved.

A. Data Security and Adherence

Strict laws like GDPR, PCI-DSS, and local data protection laws safeguard extremely private banking transaction data [40]. The suggested fraud detection architecture only processes the most important transaction data because it was created with privacy in mind. The danger of exposure is reduced when personally identifiable information (PII) is anonymized or tokenized before being used in a model. Role-based permissions are used to limit who has access to sensitive data, and we keep account of every data processing operation for auditing purposes. Additionally, regulations are in place to guarantee that transaction data is only kept for as long as the law permits.

B. Maintaining the System and Model's Security

The entire fraud detection process is protected thanks to robust security measures. Standard cryptography algorithms are used to encrypt data during delivery and storage. Unauthorized access to model endpoints is prevented via secure authentication procedures and APIs [41]. To guard against assaults and model exploitation, the system includes input validation, anomaly monitoring, and rate constraints. To prevent manipulation and guarantee that only certified models are used in production, model versioning and integrity checks are also used.

C. Morality and fairness

When AI-powered fraud detection systems inadvertently incorporate bias, they may unfairly treat specific individuals or transaction types. To monitor differences between groups based on behaviour or demographics, fairness-aware evaluation techniques are used, as long as the law allows. By doing this, the risk is decreased. Explainable AI techniques, such local explanations and feature attribution, help make sense of model results, especially for transactions that are rejected or highlighted. Because of this transparency, regular people are able to monitor events and hold others responsible.

D. Human Responsibility and Authority

To lessen the possibility of inefficient automated activities, fraud analysts look at high-risk human judgments [42]. It is made clear by clear governance mechanisms who is in charge of choosing models, making adjustments, and handling problems. These procedures guarantee that the suggested system satisfies the financial services industry's responsible AI standards and is both ethical and safe.

X. LIMITATIONS & FUTURE WORK

There are still a lot of problems that need to be fixed, but the suggested AI-based fraud detection system is useful and functional [43]. These problems allow for additional research and development.

A. Limitations of the data

The lack of sufficient well-labelled fraud data is one of the main problems. Because fraudulent transactions are very uncommon and often incorrectly classified, supervised learning systems are less trustworthy. Although anomaly detection and synthetic data might be useful in this regard, synthetic patterns might not adequately capture the complexity of real-world fraud. Furthermore, anonymized public datasets limit the number of features, which hinders the model's ability to understand the semantic and contextual connections present in private financial data [44].

B. The complexity and readability of the model

Compared to conventional rule-based systems, deep learning models are harder to understand and demand more processing power, especially hybrid designs that combine sequence models and anomaly detection. Although a number of explainable AI technologies are being created, it is still challenging to give each automated choice a clear, regulator-friendly rationale [45]. This cap is particularly important in highly regulated financial settings.

C. Fraud Concepts and Patterns' Evolution

Rapidly evolving fraud tactics can eventually deteriorate performance, even though the framework is more resistant to concept drift. Regularly retraining the system can address this challenge, although doing so raises the system's complexity and operating expenses [46]. Real-time adaptation without complete retraining remains a challenge.

D. Scalability and deployment problems

Large-scale real-time deployment may face infrastructure and latency challenges, especially when complex sequence models or graph-based extensions are used. Striking the right balance between strict response time requirements and detecting precision is an ongoing challenge.

E. Research Paths for the Future

We intend to use graph neural networks in the future to find fraud tendencies among customers, retailers, and devices. People can further adjust to new types of fraud with the use of online and ongoing education programs. Federated learning, which enables users to train models together without exchanging raw data, is a potential strategy for enhancing privacy [47]. In order to improve transparency and dependability, we will also look at more effective explainability and fairness auditing techniques.

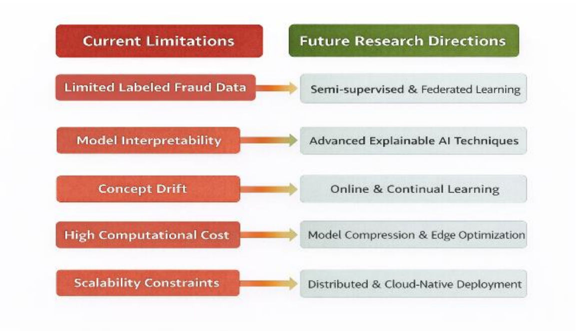
Current Limitations
Future Research Directions
Limited Labeled Fraud Data
Semi-supervised & Federated Learning
Model Interpretability
Advanced Explainable AI Techniques
Concept Drift
Online & Continual Learning
High Computational Cost
Model Compression & Edge Optimization
Scalability Constraints
Distributed & Cloud-Native Deployment

Figure 7: Limitations vs. Future Work Roadmap

XI. CONCLUSION

This study presented a thorough AI-based fraud detection system that uses deep learning algorithms to reliably and accurately detect fraudulent activity in financial transaction data. The suggested hybrid method effectively detects both known fraud patterns and new or evolving attack behaviors that have never been seen before by combining supervised sequence modelling with unsupervised anomaly detection. Traditional rule-based and classical machine learning systems struggle to alter and adapt to shifting financial circumstances, which is one of the biggest problems with these dual capacities.

On several performance metrics, the suggested method performs noticeably better than baseline models, especially F1-score and AUC-PR, which are crucial for fraud detection tasks with extreme imbalance. The sensitivity and ablation research show how important every architectural element is. They show how temporal modelling, decision-level fusion, and anomaly detection can all improve detection accuracy. These results demonstrate deep learning's ability to detect fraudulent activity, lower false positives, and forecast complex transactional behavior.

The model's effectiveness and possible applications in practical contexts, such as real-time inference, system integration, privacy protection, and ethical adherence, are investigated in this work. Because it enables you to keep an eye on the model at all times, involve a person, and make decisions quickly, the recommended deployment architecture is suitable for production-scale financial systems. By integrating explainable AI techniques, robust security measures, and privacy-by-design principles, the framework complies with responsible and regulatory AI standards. Everyone consequently has greater self-confidence.

Although there are some positive indications of progress, there are drawbacks as well, like concept drift and a dearth of labelled data to aid in comprehension. Graph neural networks, federated learning, continuous adaptability, and enhanced explainability are just a few of the intriguing research issues covered in the limitations and future work section. Finally, by offering a scalable, flexible, and morally sound method of fraud detection, our research improves the field of AI-driven financial security and lays the groundwork for next developments in intelligent fraud prevention systems.

REFERENCES

[1]. Vukovic, Darko B., M. Moinak, and E. M. Grigoryeva. "Digitalization and the Future of Financial Services." *Contributions to Finance and Accounting Springer: Cham* (2022).

[2]. Achim, Monica Violeta, and Sorin Nicolae Borlea. *Economic and financial crime*. Springer International Publishing, 2020.

[3]. Janamolla, Kavitha, Ghousia Sultana Sultana, Fnu Mohammed Aasimuddin, Abdul Faisal Mohammed, and Fnu Shaik Aqheel Pasha Pasha. "Integrating Blockchain and AI for Efficient Trade Exception Handling: A Case Study in Cross-Border Settlements." *Journal of Cognitive Computing and Cybernetic Innovations* 1, no. 1 (2025): 24-30.

[4]. Kong, Lingjing, Biwei Huang, Feng Xie, Eric Xing, Yuejie Chi, and Kun Zhang. "Identification of nonlinear latent hierarchical models." *Advances in Neural Information Processing Systems* 36 (2023): 2010-2032.

[5]. Aasimuddin, Mohammed, and Shahnawaz Mohammed. "AI-Generated Deepfakes for Cyber Fraud and Detection."

[6]. Khadri, Waheeduddin, Janamolla Kavitha Reddy, Abubakar Mohammed, and T. Kiruthiga. "The Smart Banking Automation for High Rated Financial Transactions using Deep Learning." In *2024 IEEE 3rd World Conference on Applied Intelligence and Computing (AIC)*, pp. 686-692. IEEE, 2024.

[7]. Bello, Oluwabusayo Adijat, Abidemi Ogundipe, Damilola Mohammed, Folorunso Adebola, and Olalekan Ayodeji Alonge. "AI-driven approaches for real-time fraud detection in US financial transactions: Challenges and opportunities." *European Journal of Computer Science and Information Technology* 11, no. 6 (2023): 84-102.

[8]. Mohammed, Abdul Khaleeq, and Mohammed Azmath Ansari. "The Impact and Limitations of AI in Power BI: A."

[9]. Gandhar, Akash, Kapil Gupta, Aman Kumar Pandey, and Dharm Raj. "Fraud detection using machine learning and deep learning." *SN Computer Science* 5, no. 5 (2024): 453.

[10]. Corso, Gabriele, Hannes Stark, Stefanie Jegelka, Tommi Jaakkola, and Regina Barzilay. "Graph neural networks." *Nature Reviews Methods Primers* 4, no. 1 (2024): 17.

[11]. Imhof, Michael J., Scott E. Seavey, and Olena V. Watanabe. "Competition, proprietary costs of financial reporting, and financial statement comparability." *Journal of Accounting, Auditing & Finance* 37, no. 1 (2022): 114-142.

[12]. Syed, Waheeduddin Khadri, Abubakar Mohammed, Janamolla Kavitha Reddy, and S. Dhanasekaran. "Biometric authentication systems in banking: A technical evaluation of security measures." In *2024 IEEE 3rd World Conference on Applied Intelligence and Computing (AIC)*, pp. 1331-1336. IEEE, 2024.

[13]. Fan, Cheng, Meiling Chen, Xinghua Wang, Jiayuan Wang, and Bufu Huang. "A review on data preprocessing techniques toward efficient and reliable knowledge discovery from building operational data." *Frontiers in energy research* 9 (2021): 652801.

[14]. Wu, Qian, Biao Hou, Zaidao Wen, Zhongle Ren, and Licheng Jiao. "Cost-sensitive latent space learning for imbalanced PolSAR image classification." *IEEE Transactions on Geoscience and Remote Sensing* 59, no. 6 (2020): 4802-4817.

[15]. Chittoju, Siva Sai Ram, Sireesha Kolla, Mubashir Ali Ahmed, and Abdul Raheman Mohammed. "Synergistic Integration of Blockchain and Artificial Intelligence for Robust IoT and Critical Infrastructure Security."

[16]. Hoare, Charles Antony Richard. "Communicating sequential processes." In *Theories of Programming: The Life and Works of Tony Hoare*, pp. 157-186. 2021.

[17]. Mohammed, Abubakar, Ghousia Sultana, Fnu Mohammed Aasimuddin, and Shahnawaz Mohammed. "Leveraging Natural Language Processing for Trade Exception Classification and Resolution in Capital Markets: A Comprehensive Study." *Journal of Cognitive Computing and Cybernetic Innovations* 1, no. 1 (2025): 14-18.

[18]. Lin, Zhuoran. "Comparative Study of LSTM and Transformer." In *Proceedings of the 2023 2nd International Conference on Artificial Intelligence, Internet and Digital Economy (ICAID 2023)*, vol. 9, p. 72. Springer Nature, 2023.

[19]. Wang, Zhenghong, Sijie Ruan, Tianqiang Huang, Haoyi Zhou, Shanghang Zhang, Yi Wang, Leye Wang, Zhou Huang, and Yu Liu. "A lightweight multi-layer perceptron for efficient multivariate time series forecasting." *Knowledge-Based Systems* 288 (2024): 111463.

[20]. Joseph, V. Roshan, and Akhil Vakayil. "SPlit: An optimal method for data splitting." *Technometrics* 64, no. 2 (2022): 166-176.

[21]. Mohammed, Shanavaz, Nasar Mohammed, Sruthi Balammagary, Sireesha Kolla, Srujan Kumar Ganta, and Shuaib Abdul Khader. "HARNESSING ARTIFICIAL INTELLIGENCE FOR PUBLIC HEALTH AND EPIDEMIOLOGY: OPPORTUNITIES, BARRIERS, AND PATHWAYS TO EQUITABLE GLOBAL IMPACT."

[22]. Ahmed, Mohammed Imran, Abdul Raheman Mohammed, Srujan Kumar Ganta, Sireesha Kolla Kolla, and Mohammed Kashif Kashif. "AI-Driven Green Construction: Optimizing Energy Efficiency, Waste Management and Security for Sustainable Buildings." *Journal of Cognitive Computing and Cybernetic Innovations* 1, no. 1 (2025): 37-41.

DOI: 10.17148/IJIREEICE.2025.131148

[23]. Huisman, Merel. "When AUC-ROC and accuracy are not accurate: what everyone needs to know about evaluating artificial intelligence in radiology." *European Radiology* 34, no. 12 (2024): 7892-7894.

[24]. Chollet, Francois, and François Chollet. *Deep learning with Python*. simon and schuster, 2021.

[25]. Mohammed, Naveed Uddin, Zubair Ahmed Mohammed, Shravan Kumar Reddy Gunda, Akheel Mohammed, and Moin Uddin Khaja. "Networking with AI: Optimizing Network Planning, Management, and Security through the medium of Artificial Intelligence."

[26]. Janiesch, Christian, Patrick Zschech, and Kai Heinrich. "Machine learning and deep learning." *Electronic markets* 31, no. 3 (2021): 685-695.

[27]. Miao, Jiaju, and Wei Zhu. "Precision–recall curve (PRC) classification trees." *Evolutionary intelligence* 15, no. 3 (2022): 1545-1569.

[28]. Gouni, Praveen Kumar Reddy, and Eraj Farheen Ansari. "The Impact of Cyber-Physical Attacks on AI-Enabled Business Systems

[29]. Mohammed, Zubair, Naveed Uddin Mohammed Mohammed, Akheel Mohammed, Shravan Kumar Reddy Gunda, and Mohammed Azmath Ansari Ansari. "AI-Powered Energy Efficient and Sustainable Cloud Networking." *Journal of Cognitive Computing and Cybernetic Innovations* 1, no. 1 (2025): 31-36.

[30]. Chittoju, S. R., and Siraj Farheen Ansari. "Blockchain's Evolution in Financial Services: Enhancing Security, Transparency, and Operational Efficiency." *International Journal of Advanced Research in Computer and Communication Engineering* 13, no. 12 (2024): 1-5.

[31]. Bernard, Philippe, Najat El Mekkaoui De Freitas, and Bertrand B. Maillet. "A financial fraud detection indicator for investors: an IDeA." *Annals of Operations Research* 313, no. 2 (2022): 809-832.

[32]. Mohammed, Naveed Uddin, and Mohd Abdul Raheem Raheem. "Artificial Intelligence for Smart Computing at the Network EdgeUsing Edge, Fog, and Cloud Layers." Journal of Cognitive Computing and Cybernetic Innovations 1, no. 3 (2025): 14-20.

[33]. Hurtik, Petr, Stefania Tomasiello, Jan Hula, and David Hynar. "Binary cross-entropy with dynamical clipping." *Neural Computing and Applications* 34, no. 14 (2022): 12029-12041.

[34]. Peng, ShengYun, Weilin Xu, Cory Cornelius, Matthew Hull, Kevin Li, Rahul Duggal, Mansi Phute, Jason Martin, and Duen Horng Chau. "Robust principles: Architectural design principles for adversarially robust cnns." *arXiv preprint arXiv:2308.16258* (2023).

[35]. Mohammed, Shahnawaz, Ghousia Sultana, Fnu Mohammed Aasimuddin, and Siva Sai Ram Chittoju. "AI-Driven Automated Malware Analysis." (2025).

[36]. Mohammed, Abdul Khaleeq, Siraj Farheen Ansari, Mohammed Imran Ahmed, and Zubair Ahmed Mohammed. "Boosting Decision-Making with LLM-Powered Prompts in PowerBI."

[37]. Pishgoo, Boshra, Ahmad Akbari Azirani, and Bijan Raahemi. "A hybrid distributed batch-stream processing approach for anomaly detection." *Information Sciences* 543 (2021): 309-327.

[38]. Mohammed, Akheel, Zubair Ahmed Mohammed, Naveed Uddin Mohammed, Shravan Kumar Gunda, Mohammed Azmath Ansari, and Mohd Abdul Raheem. "AI-NATIVE WIRELESS NETWORKS: TRANSFORMING CONNECTIVITY, EFFICIENCY, AND AUTONOMY FOR 5G/6G AND BEYOND

[39]. Ansari, Meraj Farheen. "Redefining Cybersecurity: Strategic Integration of Artificial Intelligence for Proactive Threat Defense and Ethical Resilience."

[40]. SSeaman, Jim. *PCI DSS: An integrated data security standard guide*. Apress, 2020.

[41]. Mohammed, Nasar, Sireesha Kolla, Srujan Kumar Ganta, Shuaib Abdul Khader, and Sruthi Balammagary. "Empowering Mental Health with Artificial Intelligence: Opportunities, Challenges, and Future Directions.

[42]. Zambrana-Tévar, Nicolás. "The international responsibility of the Holy See for human rights violations." *Religions* 13, no. 6 (2022): 520.

[43]. RAHEEM, MOHD ABDUL, and MOHAMMED AZMATH ANSARI. "INTELLIGENT AND TRUSTWORTHY 6G: AI-DRIVEN ARCHITECTURES, APPLICATIONS, AND SECURITY FRAMEWORKS.

[44]. Liu, Jiaxin, Yi Yang, and Kar Yan Tam. "Beyond surface similarity: Detecting subtle semantic shifts in financial narratives." *arXiv preprint arXiv:2403.14341* (2024).

[45]. Zhang, Ke, and Ayse Begum Aslan. "AI technologies for education: Recent research & future directions." *Computers and education: Artificial intelligence* 2 (2021): 100025.

[46]. Benzaid, Chafika, and Tarik Taleb. "AI-driven zero touch network and service management in 5G and beyond: Challenges and research directions." *Ieee Network* 34, no. 2 (2020): 186-194.

[47]. Zhang, Chen, Yu Xie, Hang Bai, Bin Yu, Weihong Li, and Yuan Gao. "A survey on federated learning." *Knowledge-Based Systems* 216 (2021): 106775.