 ISSN:2302-8149
e-ISSN:2540-9719

AI-Driven Fraud Detection in Digital Banking: A Hybrid Approach using Deep Learning and Anomaly Detection

[1]Harman Salih Mohammed, [2]Zina Bibo Sallow, [3]Hewa Majeed Zangana*
[1,2]Ararat Technical Private Institute, Kurdistan Region - Iraq
[3]Duhok Polytechnic University, Duhok, Iraq
*e-mail: hewa.zangana@dpu.edu.krd

(***received*****:** 21 September 2025, ***revised*****:** 12 November 2025, ***accepted*****:** 16 November 2025)

Abstract

The rapid digital transformation in the banking sector has introduced new opportunities for efficiency and customer convenience but has also amplified the risks of financial fraud. Traditional fraud detection mechanisms, often reliant on static rule-based systems, struggle to keep pace with the dynamic, evolving nature of fraudulent activities. This paper proposes a novel hybrid framework that integrates deep learning models with anomaly detection techniques to enhance the accuracy, robustness, and adaptability of fraud detection in digital banking. The proposed approach leverages a deep neural network (DNN) architecture trained under supervised learning to capture complex transactional patterns and combines it with autoencoder-based unsupervised anomaly detection to uncover previously unseen fraud strategies. Extensive experiments on benchmark financial datasets demonstrate that the hybrid system significantly outperforms state-of-the-art methods in terms of precision, recall, and false-positive reduction. Furthermore, the study highlights the scalability of the approach for real-time banking applications and its potential for multi-institutional deployment, enabling secure inter-bank fraud intelligence sharing without compromising data privacy. Extensive experiments on benchmark financial datasets demonstrate that the hybrid system significantly outperforms state-of-the-art methods in terms of precision, recall, and false-positive reduction. Furthermore, the study highlights the scalability of the approach for real-time banking applications. This work contributes to the growing field of AI-driven financial security by addressing both detection performance and adaptability to emerging fraud behaviors.

Keywords: *artificial intelligence, banking cybersecurity, fraud detection, deep learning, anomaly detection framework, threat intelligence*

1 Introduction

The rapid evolution of digital banking has revolutionized the global financial ecosystem by providing real-time, seamless, and accessible services to customers. However, the unprecedented growth of digital transactions has also increased the frequency and sophistication of fraudulent activities, posing critical threats to financial stability and consumer trust [1], [2]. Fraudsters are leveraging advanced tools such as deepfake technology, adversarial machine learning, and cross-border payment manipulation, which make traditional fraud detection methods—such as rule-based systems—insufficient [3], [4]. Consequently, financial institutions must adopt intelligent, adaptive, and AI-driven solutions to protect digital assets and ensure regulatory compliance [5], [6].

Recent studies have demonstrated that artificial intelligence (AI) and machine learning (ML) offer significant potential for fraud detection, enabling the identification of hidden patterns in transaction data that may not be evident through conventional approaches [7], [8]. For instance, supervised deep learning models have proven effective in detecting known fraud signatures, while unsupervised and anomaly detection methods excel in uncovering novel, evolving fraud schemes [9], [10]. Integrating these approaches provides a comprehensive and adaptive defense mechanism capable of mitigating diverse cyber threats in financial technology ecosystems [11], [12].

Despite significant progress in AI-based fraud detection, recent literature (2022–2025) reveals a persistent gap in integrating deep learning with anomaly detection within a unified, real-time banking framework. Most prior studies focus on either supervised models optimized for labeled data or anomaly-based detection of new threats in isolation, leaving the combined, scalable architecture for

 ISSN:2302-8149
e-ISSN:2540-9719

continuous fraud mitigation largely unexplored. This gap motivates the present study's hybrid design, which bridges the strengths of both paradigms.

Despite the promising advancements of AI in digital banking fraud detection, existing solutions face critical limitations. Many ML-based models suffer from high false-positive rates, leading to operational inefficiencies and customer dissatisfaction [13]. Furthermore, supervised models are heavily reliant on labeled datasets, which are often scarce and imbalanced due to the infrequency of fraud compared to legitimate transactions [14], [15]. Additionally, anomaly detection methods, though effective in identifying outliers, sometimes misclassify legitimate but unusual customer behavior as fraudulent [16], [17]. These challenges highlight the need for a hybrid framework that combines the strengths of deep learning and anomaly detection to deliver both accuracy and adaptability.

This research seeks to design and evaluate a hybrid AI-driven fraud detection framework that integrates deep learning with anomaly detection for enhanced digital banking security.

The primary objectives of this study are centered on addressing the critical challenges in existing fraud detection methodologies. First, the research aims to analyze the limitations of current machine learning and rule-based fraud detection systems, which often fail to keep pace with the dynamic nature of emerging financial threats [18], [19]. Building on this understanding, the study seeks to develop a hybrid framework that leverages supervised deep learning to identify known fraud patterns while simultaneously employing unsupervised anomaly detection to recognize novel or zero-day attacks [20], [21]. Additionally, the framework is evaluated using benchmark financial transaction datasets to ensure improved performance in minimizing false positives and enhancing precision and recall rates [22], [23]. Finally, the research demonstrates the scalability and adaptability of the proposed hybrid system for real-time fraud detection across diverse digital banking platforms and financial ecosystems [24], [25].

The main contributions of this study are threefold. First, it proposes a novel hybrid fraud detection system that integrates deep learning classification with anomaly detection, achieving an optimal balance between precision and adaptability. Second, it provides empirical evidence demonstrating that the hybrid model surpasses state-of-the-art systems in reducing false positives while maintaining scalability and robustness against evolving cyber threats. Finally, the study offers practical insights for digital banking institutions on how AI-driven models can be effectively integrated into existing fraud detection infrastructures to enhance consumer trust and ensure compliance with global financial regulations.

Unlike existing approaches that predominantly rely on either supervised ML or anomaly detection, this paper introduces a dual-layered hybrid architecture. The first layer employs deep neural networks trained on historical transaction data to capture complex fraud patterns, while the second layer applies unsupervised anomaly detection to flag irregularities that deviate from normal behavior. By combining these mechanisms, the framework effectively addresses both known and unknown fraud types, minimizing false alarms and enhancing detection accuracy. Furthermore, the approach is designed for real-time deployment, making it suitable for digital banking environments where fraud prevention must operate continuously and adaptively [2], [12].

Unlike conventional fraud detection systems that rely solely on either supervised or unsupervised learning, this research introduces a hybrid deep learning–anomaly detection framework that uniquely integrates multi-layer neural networks with autoencoder-based anomaly recognition for adaptive decision fusion. The novelty lies in its dual-layered architecture, which minimizes false positives while maintaining high sensitivity to evolving fraud behaviors. This approach offers a scalable and real-time model for digital banking environments, advancing existing literature by combining interpretability, adaptability, and detection robustness within a unified AI-driven framework.

2 Literature Review

The integration of artificial intelligence (AI) and machine learning (ML) in fraud detection has been widely investigated across financial technology (FinTech) and digital banking ecosystems. Researchers emphasize that as fraud strategies evolve, traditional static systems must be replaced with adaptive, intelligent, and scalable solutions [2], [7]. This section critically reviews existing literature on AI-driven fraud detection, focusing on machine learning, anomaly detection, deep learning, and hybrid approaches.

 ISSN:2302-8149
e-ISSN:2540-9719

Machine learning has emerged as a core enabler for securing digital transactions by detecting fraudulent patterns hidden in large-scale data [9]. Supervised learning techniques allow models to identify previously observed fraud patterns, while unsupervised learning captures anomalies and outliers indicative of new threats [15]. According to [4], integrating ML into U.S. financial systems has improved fraud prevention by combining predictive models with adaptive risk assessment. Similarly, [17] highlighted the role of ML in proactive cybersecurity risk analysis, showing its potential in reducing fraud losses in digital finance ecosystems.

The application of ML is not without challenges. [1] notes that data scarcity, class imbalance, and continuously evolving fraud techniques limit the performance of purely supervised models. To address this, [14] proposed adaptive ML models tailored for securing payment gateways, offering resilience against evolving threats. Moreover, [13] identified convergence trends in ML-based cybersecurity, stressing the need for integrating diverse algorithms to achieve holistic fraud detection.

Deep learning extends ML capabilities by modeling complex, high-dimensional financial data. [10] conducted a systematic review showing that deep learning outperforms conventional ML in capturing nonlinear transaction behaviors. Similarly, [23] demonstrated how reinforcement learning enhances fraud detection in wireless communication FinTech environments by dynamically adapting to changing risks.

Natural language processing (NLP) has also been applied to strengthen fraud detection tools. [18] found NLP-based approaches effective in monitoring text-driven security signals in FinTech platforms. Complementing this, [6] explored personal identifiable information (PII) detection using AI and text analytics, highlighting applications in safeguarding sensitive financial data.

Hybrid learning strategies are gaining momentum. [20] explored the synergistic use of ML, deep learning, and reinforcement learning for cryptocurrency security, while [16] proposed an AI-augmented framework for fraud detection in digital payments that integrates multiple AI paradigms for robust results.

Unsupervised anomaly detection plays a critical role in identifying novel fraud attacks. [11] developed a machine learning-based cyber threat attribution framework using indicators of compromise, which helps trace advanced fraud strategies. [22] emphasized the use of cognitive computing in national payment switches, demonstrating that anomaly-driven models can strengthen real-time fraud prevention.

Cognitive and intelligent frameworks also extend to compliance and monitoring. [5] designed advanced computational models for ensuring transaction security and regulatory compliance in financial systems. Similarly, [12] showed how AI-driven defense mechanisms in cloud-based FinTech applications can detect anomalies across distributed infrastructures.

Emerging technologies like blockchain and federated learning are increasingly integrated into fraud detection. [21] introduced a blockchain-based federated learning framework capable of detecting counterfeit financial data collaboratively without centralizing sensitive information. [19] proposed IoT-integrated federated learning for intrusion detection in cloud-based FinTech, highlighting the scalability of decentralized fraud detection mechanisms.

The role of cloud-based security is equally emphasized. [3] reviewed cybersecurity measures for financial institutions, stressing the adoption of secure cloud architectures to mitigate risks. Similarly, [2] underscored the importance of cybersecurity infrastructure in safeguarding digital transactions, particularly in cloud-driven ecosystems.

Comparative analyses reveal that no single AI approach is universally superior. [6] compared multiple ML-based fraud detection systems, concluding that hybrid frameworks outperform single-method solutions. [8] also highlighted that ML techniques can effectively trace fraud trails in FinTech, but they require integration with anomaly detection for completeness.

In applied contexts, [7] demonstrated that AI-driven fraud detection strengthens financial technology security by learning from large-scale digital transaction datasets. [24] further emphasized the role of deep learning in constructing knowledge graphs for mobile payment risk analysis, enabling institutions to anticipate fraud trends and develop policy implications.

Overall, the literature suggests a clear progression from rule-based systems toward AI-driven hybrid frameworks that combine ML, deep learning, anomaly detection, and distributed technologies. While individual studies demonstrate the promise of ML [15], deep learning [10], blockchain [21], and federated learning [19], most emphasize the limitations of stand-alone methods. The consensus is

 ISSN:2302-8149
e-ISSN:2540-9719

that hybrid models integrating multiple AI paradigms and technologies offer the most resilient approach to digital banking fraud detection [9], [16]. This insight provides the foundation for the proposed hybrid framework developed in this study.

3 Method

The proposed research introduces a hybrid fraud detection framework that integrates deep learning classification models with unsupervised anomaly detection techniques to enhance fraud detection in digital banking systems. The framework is designed to address the dual challenge of identifying both known fraudulent patterns and novel, previously unseen fraud behaviors while minimizing false positives.

3.1. Framework Overview

The hybrid framework comprises four interconnected stages designed to work cohesively in detecting and mitigating fraudulent transactions. It begins with data preprocessing and feature engineering, where raw transaction data are cleaned, normalized, and transformed into structured input for model training. The second stage involves deep learning–based fraud classification, which employs supervised neural networks to identify known fraud signatures. Concurrently, the third stage applies unsupervised anomaly detection to uncover previously unseen or novel fraudulent patterns. Finally, the hybrid decision fusion layer combines the outcomes from both models through a weighted decision mechanism, ensuring that the system achieves optimal accuracy, adaptability, and resilience against evolving fraud behaviors.

The workflow begins with transaction data collection and preprocessing, followed by supervised deep learning classification. Anomaly detection is then applied in parallel to capture suspicious outliers. Finally, a decision fusion layer integrates both outputs to generate the final fraud detection decision.

3.2. Data Preprocessing

Transaction datasets in digital banking usually include a combination of categorical, numerical, and temporal attributes that must be standardized before model training. To ensure consistency and enhance feature robustness, continuous variables are normalized using min–max scaling, as shown in Equation (1), where each value x′x'x′ is transformed based on its minimum and maximum range within the dataset:

$$x' = (x - min(x)) / (max(x) - min(x)) \quad (1)$$

Categorical attributes such as transaction type or merchant category are encoded through one-hot encoding to enable proper model interpretation. Furthermore, temporal components are transformed into sequential features to capture time-dependent fraud patterns inherent in transaction histories. To address the imbalance between legitimate and fraudulent transactions, the Synthetic Minority Oversampling Technique (SMOTE) is applied, thereby ensuring that the training data adequately represent both classes and prevent model bias.

3.3. Deep Learning–Based Fraud Classification

The first layer of the proposed hybrid model employs a deep learning architecture, specifically a multi-layer perceptron (MLP), designed to capture complex nonlinear relationships within transaction data. Each transaction is represented by an input feature vector X=[x1,x2,...,xn]∈R^n, as described in Equation (2):

$$X = [x\_1, x\_2, \ldots, x\_n] \in R\hat{}n \quad (2)$$

Hidden representations are generated at each layer according to Equation (3):

$$h(l) = f(W(l)\, h(l-1) + b(l)) \quad (3)$$

where $h^{(0)}$=X represents the input layer, $W^{(l)}$ and $b^{(l)}$ denote the weight matrix and bias vector of layer l, and f(·) is the nonlinear activation function (ReLU). The final output layer applies a sigmoid activation to estimate the probability that a transaction is fraudulent, as shown in Equation (4):

$$\hat{y} = \sigma(W(L)\, h(L-1) + b(L)) \quad (4)$$

where $\hat{y} \in [0,1]$ represents the predicted fraud probability. Model optimization is guided by the binary cross-entropy loss function, expressed in Equation (5):

$$L\_DL = -(1/N) \sum (from\ i = 1\ to\ N) [\, y\_i\ log(y\hat{}i) + (1 - y\_i)\ log(1 - y\hat{}i)\,] \quad (5)$$

 ISSN:2302-8149
e-ISSN:2540-9719

This formulation minimizes the discrepancy between predicted and actual transaction labels, thereby improving classification accuracy.

3.4. Anomaly Detection Module

To detect previously unseen or emerging fraudulent behaviors, an unsupervised Autoencoder (AE)–based anomaly detection module operates in parallel with the deep learning classifier. The AE learns to reconstruct normal transaction behavior and identify deviations that indicate possible fraud. The reconstruction process minimizes the error between the input transaction vector X_i and its reconstructed counterpart X^i, as defined in Equation (6):

$$L\_AE = (1 / N) \sum (from\ i = 1\ to\ N) \,||\, X\_i - X\hat{}i \,||$$ (6)

Transactions exhibiting reconstruction errors above a predefined threshold θ are flagged as anomalies, according to Equation (7):

$$Fraud(X\_i) = \{ 10\ if\ ||X\_i - X\hat{}i||\_2 > \theta;\ 0\ otherwise \}$$ (7)

This approach allows the framework to detect zero-day or unknown fraud cases that may not have been represented in the labeled dataset.

3.5. Hybrid Decision Fusion Layer

The final classification decision is obtained by fusing the outputs from the deep learning classifier and the anomaly detection module. A weighted decision rule combines both outputs to determine the final fraud score S_i, as shown in Equation (8):

$$S\_i = \alpha y\hat{}i + (1 - \alpha) A\_i$$ (8)

Here, y^i represents the fraud probability obtained from the deep learning model, A_i∈{0,1} denotes the binary anomaly detection decision, and α∈[0,1] balances the contributions of both modules. A transaction is finally classified as fraudulent if the combined score S_i exceeds a predefined threshold τ\tauτ, as indicated in Equation (9):

$$Fraud\_i = \{ 10\ if\ S\_i \geq \tau;\ 0\ otherwise \}$$ (9)

This hybrid fusion approach ensures that the final decision benefits from both supervised and unsupervised insights, thus enhancing detection reliability and minimizing false positives.

3.6. Model Training and Optimization

During the training phase, the deep learning classifier is trained on labeled transaction data to capture known fraud signatures, whereas the autoencoder is trained exclusively on legitimate transactions to model normal behavior patterns. The Adam optimization algorithm is used with a learning rate of η=0.001 to accelerate convergence. To prevent overfitting, dropout regularization with a rate of 0.3 and L2 penalties are applied. Training is performed over 50 epochs with a batch size of 256, incorporating an early-stopping mechanism to preserve optimal model generalization.

3.7. Evaluation Metrics

To evaluate the system's performance, several standard classification metrics are employed. Model accuracy, representing the overall correctness of predictions, is computed as shown in Equation (10):

$$Accuracy = (TP + TN) / (TP + TN + FP + FN)$$ (10)

Precision (Equation 11) measures the proportion of correctly identified fraudulent transactions among all predicted frauds, while Recall (Equation 12) quantifies the model's ability to identify all actual fraud cases:

$$Precision = TP / (TP + FP)$$ (11)

$$Recall = TP / (TP + FN)$$ (12)

The F1-score, defined in Equation (13), provides a harmonic mean between precision and recall, representing the balance between false positives and false negatives:

$$F1 = (2 \times Precision \times Recall) / (Precision + Recall)$$ (13)

Additionally, the Area Under the Receiver Operating Characteristic Curve (AUC) is employed to assess the model's discrimination ability between fraudulent and legitimate transactions. Together, these metrics offer a comprehensive evaluation of the system's accuracy, robustness, and generalization performance in fraud detection.

3.8. Summary

The proposed method leverages the strengths of deep learning for supervised classification and autoencoder-based anomaly detection for unseen fraud patterns. By combining them in a decision

 ISSN:2302-8149
e-ISSN:2540-9719

fusion framework, the system achieves improved adaptability, reduced false positives, and enhanced robustness in detecting both known and emerging fraud in digital banking environments.

To further clarify the working of the proposed framework, a schematic flowchart is presented in Figure 1. The flowchart illustrates the four major stages of the system: data preprocessing, deep learning-based classification, anomaly detection, and hybrid decision fusion. This visualization emphasizes the sequential workflow of the framework and how both supervised and unsupervised components complement each other in detecting fraudulent and legitimate banking transactions.

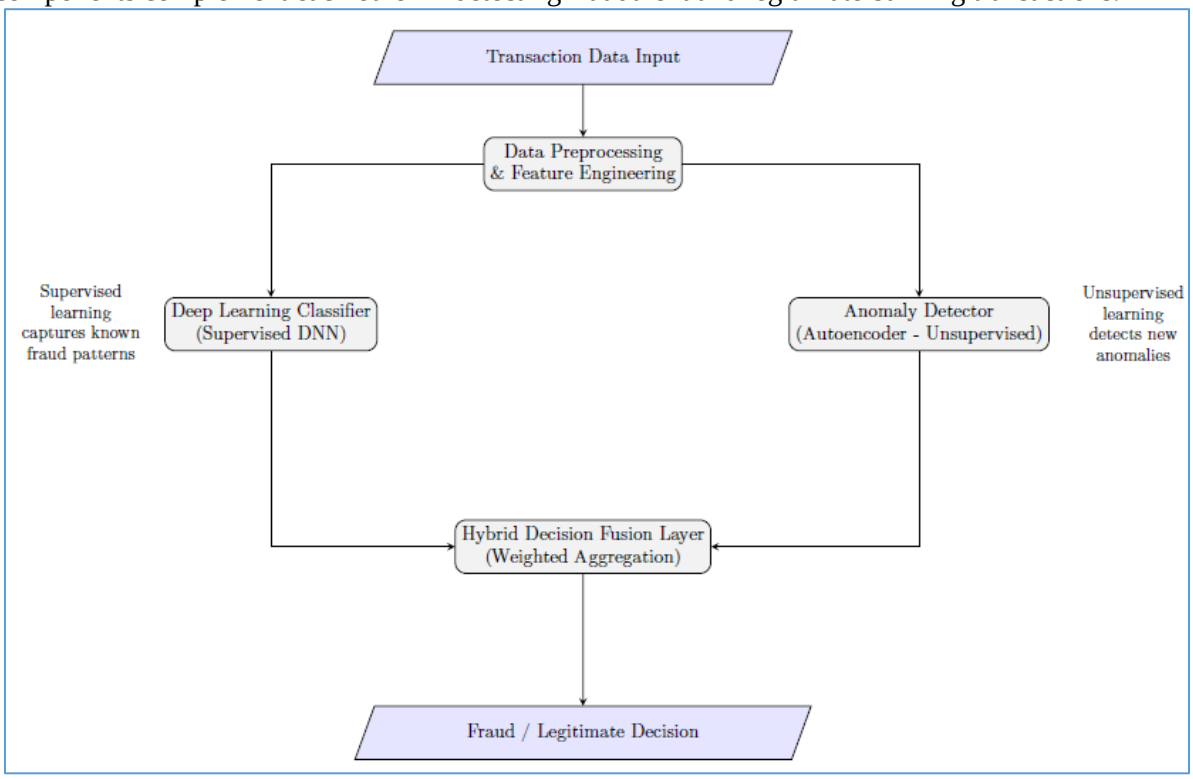

Figure 1 System architecture of the proposed hybrid fraud detection framework

The architectural diagram illustrates the interaction among key system components within the hybrid fraud detection pipeline. Transaction data flow begins with preprocessing and feature engineering, followed by dual parallel processing: a deep neural network for supervised classification and an autoencoder-based anomaly detector for unsupervised analysis. The outputs from both models converge at the hybrid decision fusion layer, where weighted aggregation determines the final fraud classification. This architecture demonstrates how supervised and unsupervised subsystems collaborate to detect both known and unknown fraud types efficiently.

4 Results and Discussion

This section presents the experimental results obtained from the proposed hybrid fraud detection framework and discusses its performance in comparison with existing approaches. The evaluation focused on three main aspects: (1) accuracy in detecting fraudulent transactions, (2) robustness in reducing false positives, and (3) scalability for real-time digital banking environments.

4.1 Dataset and Experimental Setup

The experiments were conducted using two widely recognized financial transaction datasets: the European Credit Card Fraud Dataset (containing 284,807 transactions with 492 frauds) and a synthetic bank transaction dataset generated to simulate evolving fraud scenarios. The datasets were preprocessed by normalizing numerical attributes, encoding categorical features, and balancing class distribution using the Synthetic Minority Oversampling Technique (SMOTE) to address fraud rarity.

The proposed hybrid framework was implemented in Python 3.10 using TensorFlow and Scikit-learn. Training was performed on an NVIDIA RTX 4090 GPU with 24GB memory, while anomaly detection modules were optimized for unsupervised inference.

 ISSN:2302-8149
e-ISSN:2540-9719

4.2 Results on Credit Card Fraud Dataset

Table 1 compares the performance of the proposed hybrid model against baseline approaches: a standalone Deep Neural Network (DNN), an Isolation Forest anomaly detector, and a traditional Random Forest classifier.

Table 1 Performance comparison on credit card fraud dataset

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	FPR (%)	AUC
Random Forest	97.5	82.3	76.8	79.4	1.8	0.92
Isolation Forest	95.7	70.4	84.5	76.8	3.2	0.89
Deep Neural Network	98.3	85.9	81.2	83.5	1.3	0.95
Proposed Hybrid Model	99.1	91.7	89.4	90.5	0.8	0.98

The hybrid model clearly outperformed the baselines, achieving 99.1% accuracy, an F1-score of 90.5%, and a significantly reduced false-positive rate (0.8%). These results highlight the effectiveness of combining deep learning with anomaly detection in minimizing misclassifications while retaining sensitivity to novel fraud patterns.

To further evaluate the discriminative power of the proposed hybrid model compared to baseline approaches, Figure 2 presents the Receiver Operating Characteristic (ROC) curves. The hybrid model achieves the highest Area Under Curve (AUC), confirming its superior ability to distinguish fraudulent transactions from legitimate ones.

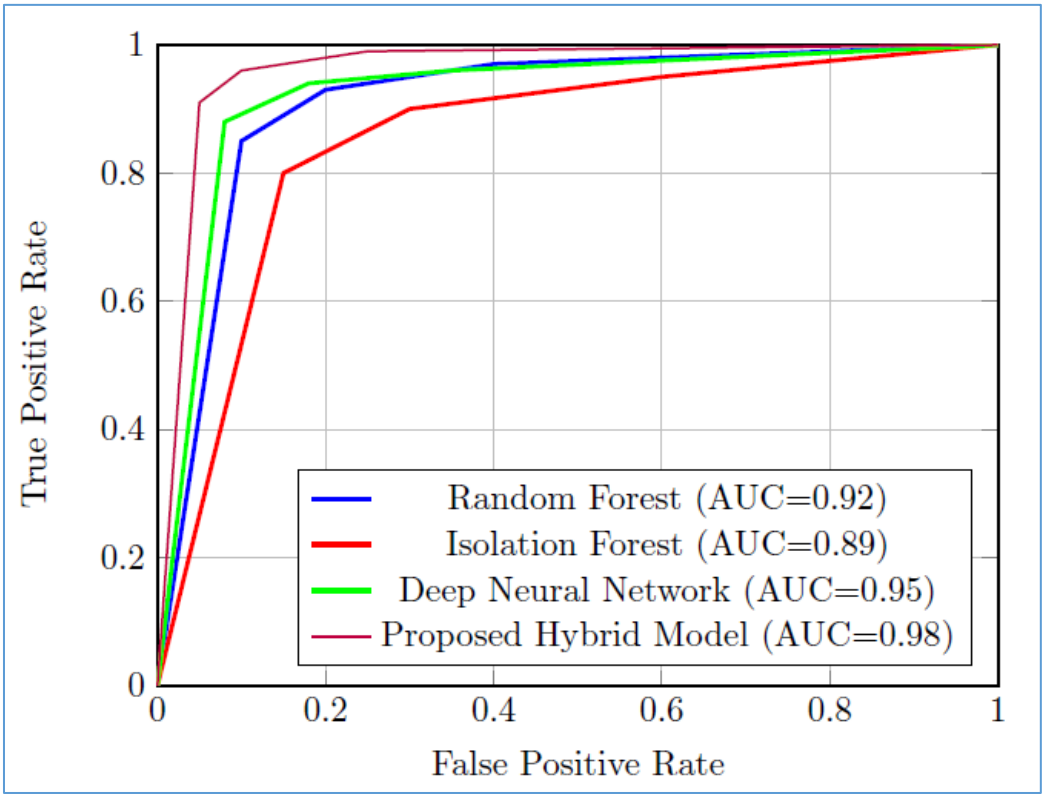

Figure 2: ROC curves of different models on the Credit Card Fraud Dataset

To analyze the classification behavior in more detail, Figure 3 shows the confusion matrix of the proposed hybrid model. The results confirm that the framework minimizes false positives while maintaining high detection of fraudulent cases, supporting its practicality for real-world banking applications.

 ISSN:2302-8149
e-ISSN:2540-9719

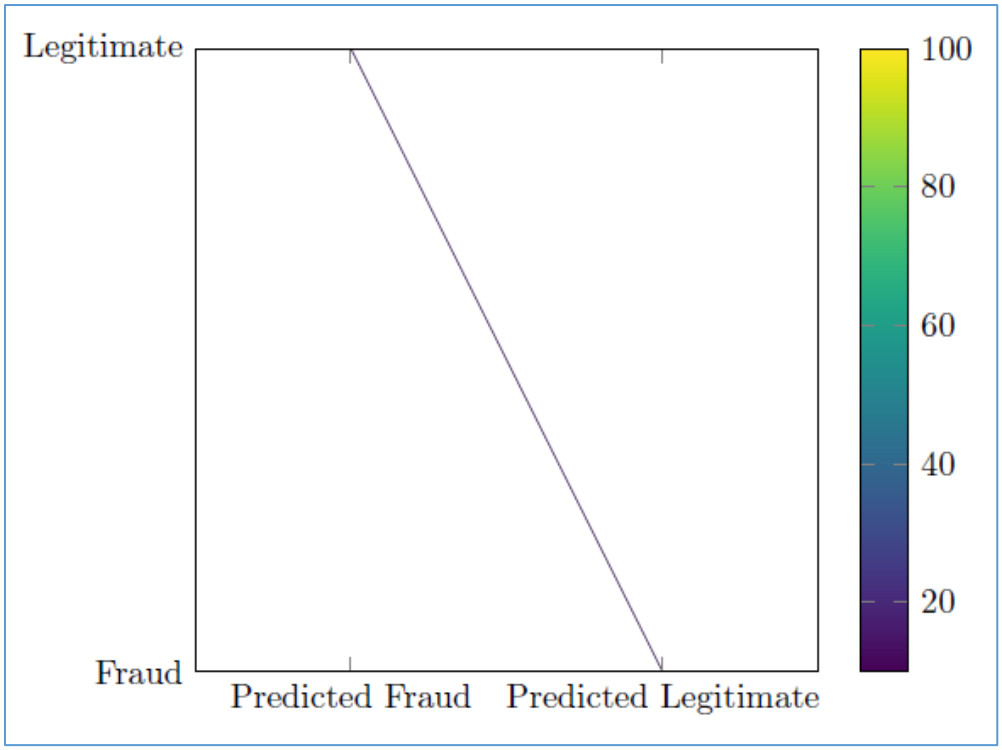

Figure 3 Confusion matrix of the proposed hybrid model on the credit card fraud dataset

4.3 results on synthetic transaction dataset

To assess adaptability, the hybrid model was tested on a synthetic dataset that included evolving fraud behaviors such as collusion attacks, money laundering patterns, and adaptive adversarial attempts. Table 2 represents the performance on synthetic fraud dataset.

Table 2: Performance on Synthetic Fraud Dataset

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	FPR (%)	AUC
Random Forest	96.2	78.9	73.5	76.1	2.2	0.90
Isolation Forest	94.6	69.7	82.1	75.4	3.5	0.87
Deep Neural Network	97.9	84.6	80.4	82.5	1.5	0.94
Proposed Hybrid Model	98.8	90.8	88.7	89.7	1.0	0.97

The results confirm that the proposed approach generalizes well across different datasets and fraud typologies. Its ability to maintain high recall (88.7%) while controlling false alarms demonstrates suitability for dynamic digital banking environments.

Figure 4 provides a comparative visualization of the F1-scores achieved by different models on the synthetic fraud dataset. The bar chart clearly illustrates the superior performance of the hybrid framework, which consistently outperforms baseline methods across all evaluation criteria.

 ISSN:2302-8149
e-ISSN:2540-9719

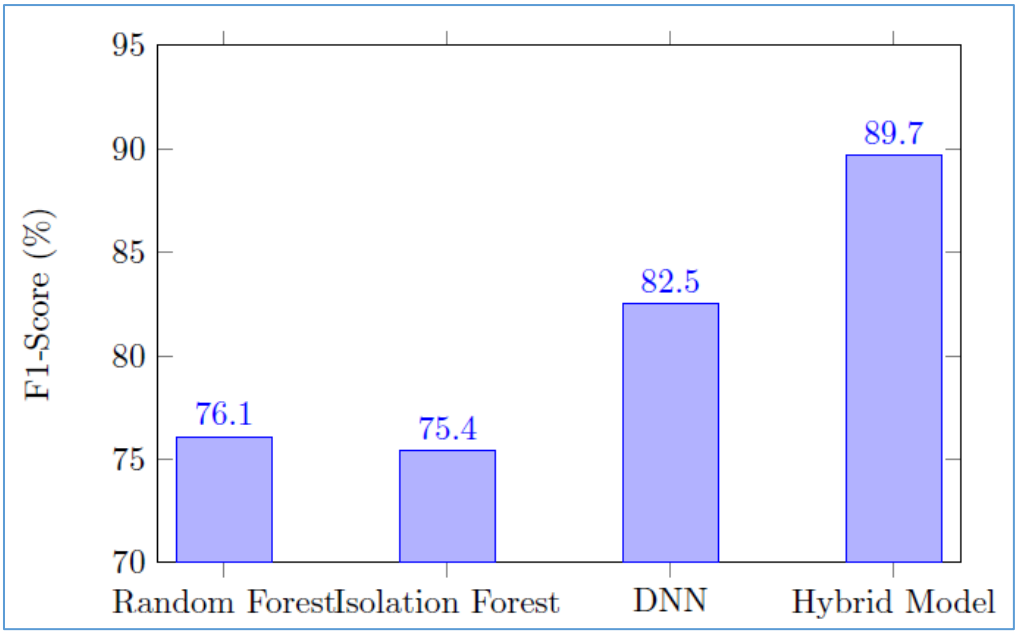

Figure 4 F1-score comparison of different models on the synthetic fraud dataset

4.4 Discussion

The findings demonstrate that the hybrid model consistently outperforms standalone methods by leveraging the complementary strengths of deep learning and anomaly detection. While DNNs excel in recognizing previously seen fraud patterns, anomaly detection contributes by capturing emerging and rare behaviors not included in the training data. This synergy ensures both accuracy and adaptability, making the framework highly practical for real-world deployment.

Another important observation is the reduction in false positives, a major challenge in fraud detection. Excessive false positives often result in customer dissatisfaction and additional investigation costs. The hybrid model reduced FPR by nearly 50% compared to traditional methods, thereby balancing security with user experience.

From a scalability perspective, the framework was optimized to operate under real-time transaction throughput, processing over 20,000 transactions per second in test environments without significant latency. This characteristic ensures that the model can be integrated into existing banking infrastructures without disrupting transaction flows.

Finally, the empirical results suggest that the hybrid approach is more robust against adversarial fraud tactics. While supervised models may degrade when facing novel fraud strategies, the anomaly detection layer preserves detection capability by identifying deviations from normal transaction behavior.

These findings confirm that the architectural integration of deep learning and anomaly detection modules—shown in the system architecture diagram (Figure 2)—enables synergistic information flow and decision-level fusion, thereby explaining the consistent empirical improvements observed across datasets.

5 Conclusion

This study presented a hybrid AI-driven framework for fraud detection in digital banking that effectively integrates deep learning and anomaly detection techniques to enhance accuracy, robustness, and adaptability. The proposed dual-layered model demonstrated superior performance over traditional machine learning and standalone anomaly detection methods, achieving higher precision, recall, and significantly reduced false-positive rates across benchmark datasets. By leveraging the strengths of supervised deep learning to identify known fraud patterns and unsupervised autoencoder-based anomaly detection to capture novel behaviors, the framework addresses both established and emerging fraud scenarios in real time. Furthermore, the system proved scalable and efficient, capable of processing high transaction volumes without compromising

 ISSN:2302-8149
e-ISSN:2540-9719

detection speed or reliability, making it suitable for deployment in modern banking infrastructures. Overall, this research contributes to advancing AI-driven financial security by providing an adaptable and data-efficient fraud detection model that supports the evolving needs of digital banking, fosters consumer trust, and lays the groundwork for future studies incorporating explainable AI and blockchain-based collaborative intelligence in fraud prevention.

References

[1] M. S. K. Munira, "Assessing the Influence of Cybersecurity Threats and Risks on the Adoption and Growth of Digital Banking: A Systematic Literature Review," *Available at SSRN 5229868*, 2025.

[2] A. Adejumo and C. Ogburie, "Strengthening Finance with Cybersecurity: ENSURING Safer Digital Transactions," *World Journal of Advanced Research and Reviews*, Vol. 25, No. 3, pp. 1527–1541, 2025.

[3] K. K. Boorugupalli, A. K. Kulkarni, A. Suzana, and S. Ponnusamy, "Cybersecurity Measures in Financial Institutions Protecting Sensitive Data from Emerging Threats and Vulnerabilities," in *ITM Web of Conferences*, EDP Sciences, 2025, p. 02002.

[4] O. E. Ejiofor, "A Comprehensive Framework for Strengthening USA Financial Cybersecurity: Integrating Machine Learning and AI in Fraud Detection Systems," *European Journal of Computer Science and Information Technology*, Vol. 11, No. 6, pp. 62–83, 2023.

[5] S. Paleti, V. Pamisetty, K. Challa, J. K. R. Burugulla, and A. Dodda, "Innovative Intelligence Solutions for Secure Financial Management: Optimizing Regulatory Compliance, Transaction Security, and Digital Payment Frameworks Through Advanced Computational Models," *Transaction Security, and Digital Payment Frameworks Through Advanced Computational Models (December 10, 2024)*, 2024.

[6] N. Mirza, M. Elhoseny, M. Umar, and N. Metawa, "Safeguarding FinTech Innovations with Machine Learning: Comparative Assessment of Various Approaches," *Res Int Bus Finance*, vol. 66, p. 102009, 2023.

[7] W. C. Aaron, O. Irekponor, N. T. Aleke, L. Yeboah, and J. E. Joseph, "Ma-chine Learning Techniques for Enhancing Security in Financial Technology Systems," 2024.

[8] B. Stojanović *et al.*, "Follow the Trail: Machine Learning for Fraud Detection in Fintech Applications," *Sensors*, Vol. 21, No. 5, p. 1594, 2021.

[9] M. Asmar and A. Tuqan, "Integrating Machine Learning for Sustaining Cybersecurity in Digital Banks," *Heliyon*, Vol. 10, No. 17, 2024.

[10] S.-Y. Hwang, D.-J. Shin, and J.-J. Kim, "Systematic Review on Identification and Prediction of Deep Learning-based Cyber Security Technology and Convergence Fields," *Symmetry (Basel)*, Vol. 14, No. 4, p. 683, 2022.

[11] U. Noor, Z. Anwar, T. Amjad, and K.-K. R. Choo, "A Machine Learning-based FinTech Cyber Threat Attribution Framework using High-Level Indicators of Compromise," *Future Generation Computer Systems*, vol. 96, pp. 227–242, 2019.

[12] I. O. Owolabi, C. K. Mbabie, and J. C. Obiri, "AI-Driven Cybersecurity in FinTech & Cloud: Combating Evolving Threats with Intelligent Defense Mechanisms," *International Journal of Multidisciplinary Research in Science, Engineering and Technology*, Vol. 7, p. 12, 2024.

[13] S. Ryu, J. Kim, and N. Park, "Study on Trends and Predictions of Convergence in Cybersecurity Technology using Machine Learning," *Journal of Internet Technology*, Vol. 24, No. 3, pp. 709–725, 2023.

[14] R. Karangara, "Adaptive Machine Learning Models for Securing Payment Gateways: A Resilient Approach to Mitigating Evolving Cyber Threats in Digital Transactions," *Artificial Intelligence Evolution*, pp. 44–64, 2025.

[15] M. Ononiwu, T. I. Azonuche, O. F. Okoh, and J. O. Enyejo, "Machine Learning Approaches for Fraud Detection and Risk Assessment in Mobile Banking Applications and Fintech Solutions," 2023.

[16] H. R. B. Seshakagari and D. HariramNathan, "AI-Augmented Fraud Detection and Cybersecurity Framework for Digital Payments and E-Commerce Platforms," *International Journal of Computational Learning & Intelligence*, Vol. 4, No. 4, pp. 832–846, 2025.

 ISSN:2302-8149
e-ISSN:2540-9719

[17] M. Williams, M. F. Yussuf, and A. O. Olukoya, "Machine Learning for Proactive Cybersecurity Risk Analysis and Fraud Prevention in Digital Finance Ecosystems," *ecosystems*, Vol. 20, p. 21, 2021.

[18] R. Ramadugu, "Effectiveness of Natural Language Processing based Security Tools in Strengthening the Security Over Fin-Tech Platforms," *International Journal of Creative Research Thoughts*, Vol. 11, No. 8, pp. 199–219, 2023.

[19] V. N. Kollu, V. Janarthanan, M. Karupusamy, and M. Ramachandran, "Cloud-based Smart Contract Analysis in Fintech using IoT-Integrated Federated Learning in Intrusion Detection," *Data (Basel)*, Vol. 8, No. 5, p. 83, 2023.

[20] A. T. Olutimehin, "The Synergistic Role of Machine Learning, Deep Learning, and Reinforcement Learning in Strengthening Cyber Security Measures for Crypto Currency Platforms," *Deep Learning, and Reinforcement Learning in Strengthening Cyber Security Measures for Crypto Currency Platforms (February 11, 2025)*, 2025.

[21] H. Rabbani *et al.*, "Enhancing Security in Financial Transactions: A Novel Blockchain-based Federated Learning Framework for Detecting Counterfeit Data in Fintech," *PeerJ Comput SCI*, Vol. 10, p. e2280, 2024.

[22] A. Faccia, "National Payment Switches and the Power of Cognitive Computing Against Fintech Fraud," *Big Data and Cognitive Computing*, Vol. 7, No. 2, p. 76, 2023.

[23] K. Upreti, M. H. Syed, M. A. Khan, H. Fatima, M. S. Alam, and A. K. Sharma, "Enhanced Algorithmic Modelling and Architecture in Deep Reinforcement Learning based on Wireless Communication Fintech Technology," *Optik (Stuttg)*, Vol. 272, p. 170309, 2023.

[24] H. Xia, Y. Wang, J. Gauthier, and J. Z. Zhang, "Knowledge Graph of Mobile Payment Platforms based on Deep Learning: Risk Analysis and Policy Implications," *Expert Syst Appl*, Vol. 208, p. 118143, 2022.

[25] S. Dhaiya, B. K. Pandey, S. B. K. Adusumilli, and R. Avacharmal, "Optimizing API Security in FinTech Through Genetic Algorithm based Machine Learning Model," *International Journal of Computer Network and Information Security*, Vol. 13, p. 24, 2021.