

A Framework for Fraud Detection in Banking Transactions Using Machine Learning and Federated Learning

Rabia Tehseen¹, Hina Shahid¹, Anam Mustaqeem¹, Muhammad Farrukh Khan², Uzma Omer³, Rubab Javaid¹

¹ Department of Computer Science, University of Central Punjab, Lahore, Pakistan

² NASTP Institute of Information Technology, Lahore, Pakistan

³ University of Education, Lahore, Pakistan

*Correspondence: rabia.tehseen@ucp.edu.pk

Citation | R. Tehseen, Hina. S, Mustaqeem. A, Khan. M. F, Omer. U, Javaid. R, “A Framework for Fraud Detection in Banking Transactions Using Machine Learning and Federated Learning”, IJIST, Vol. 07 Issue. 03 pp 1409-1421, July 2025

DOI | <https://doi.org/10.33411/ijist/20257314091421>

Received | June 09, 2025 **Revised** | July 01, 2025 **Accepted** | July 06, 2025 **Published** | July 12, 2025.

The digital banking revolution has transformed financial services to make payment faster, more convenient, and borderless. But with this revolution came an abrupt increase in fraudulent transactions through credit cards that threatening both the financial institutions and the customers. While conventional fraud detection mechanisms are not capable of addressing new-generation fraud patterns, there is an increasing demand for intelligent, adaptive, and secure solutions with high precision without any data privacy compromise. Proposed model leverages four machine learning models, Linear Regression, Decision Tree, Long Short-Term Memory (LSTM) and Convolutional Neural Network (CNN). LSTM and CNN are used due to their power in learning complicated sequential and feature-based patterns, with Decision Tree and Linear Regression added due to their ease, quick execution, and interpretability. Every model is locally trained on partitioned banking datasets for each simulated client. Model parameters are combined with the Federated Averaging (FedAvg) algorithm to create a globally shared fraud detection system. Experimental testing was conducted on a real-world banking transaction data set published in a non-IID manner to mimic real-world client situations. The federated learning paradigm achieved encouraging results: CNN and LSTM models achieved detection accuracy rates of over 95%, with outstanding performance in the detection of hidden or time-series-based fraud patterns. The Decision Tree model also achieved steady performance at 91% accuracy, and Linear Regression achieved a reasonable baseline at 88%. These results indicate that even simple models, when used in a collaborative federated environment, can contribute meaningfully to fraud detection. This research contributes to the body of research supporting federated banking solutions and fills a significant gap by demonstrating how several ML models can coexist and collaborate in a decentralized setup for fraud detection through credit card transactions.

Keywords: Deep Learning, Machine Learning, Federated Learning, Fraud Detection, Convolutional Neural Network, Long Short-Term Memory



Introduction:

Digital banking has revolutionized the way people manage their finances, allowing users to access and control their accounts within seconds from virtually any device. Today, individuals can pay bills, transfer money, use virtual cards, and monitor account balances—all without ever visiting a physical bank branch. Although, these online tools make daily life easier for consumers and help banks cut costs, they also open the door to new security headaches, especially when spotting fraud [1]. Hackers keep rewriting their playbooks to slip past old-school alarms, so stolen accounts, fake purchases, and identity theft now rank among the toughest dangers the industry faces. As scams grow trickier and more common, firms must build smarter, flexible, and future-proof detection systems—or risk losing trust and money.

Inside the wider world of online money safety, fraud detection has grown into its own niche, tapping smart tech—especially machine learning—to flag unusual spending patterns. Because ML models study historical records instead of following hard-coded rules, they pick up fresh tricks used by thieves and adjust on the fly with little human oversight [2][3]. Most systems, however, still gather every transaction into one central warehouse before training the software, an approach that raises big questions about who owns the data and what happens if it leaks. Beyond privacy, a single storage hub often mirrors only a narrow slice of behavior, leaving the model shaky and prone to missing or wrongly labeling genuine alerts in the messy real world [4][5]

The main reason we set out on this study is the pressing demand for a fraud-check system that works well yet protects the privacy and independence of each bank or credit union involved. Financial firms usually hesitate, and sometimes are legally blocked, from sharing sensitive customer information because of rules like the GDPR and local banking statutes [6]. Here, Federated Learning (FL) steps in as a privacy-friendly option, letting branches join forces to train machine-learning models using their own data while keeping every raw transaction on-site. Since the data never leaves its source, FL blends the insights from all participants and discovers patterns that a single institution might miss. That mix of local control and shared learning paves the way for a safer, more decentralized fraud-detection service that meets legal demands and works efficiently [7].

The present research is directed towards designing and evaluating a fraud detection framework integrated with four supervised ML models such as Linear Regression, Decision Tree, Long Short-Term Memory (LSTM), and Convolutional Neural Network (CNN), in a federated learning setting. Each model trains separately on a local banking dataset and contributes to a global fraud detection model by secure sharing of parameters. The most important aspect of this whole work is flexibility: institutions can choose different models according to their resources and characteristics of data instead of keeping the whole set of models uniform among each client. To illustrate, this research intends to show that heterogeneous model training in a federated context does not restrict the construction of a high-performing, privacy-preserving fraud detection system. Results from this study show promising accuracy: greater than 95% using deep learning models, and clearly proving that federated training is both practical and effective [8].

This research comprises of four objectives listed as 1. To design a strong and privacy-protecting fraud detection system specific to banking transactions. 2. To establish a federated learning architecture with support for multiple supervised machine learning models without data centralization. 3. To compare the performance of various local models (Linear Regression, Decision Tree, LSTM, CNN) under federated environments. 4. To make the framework flexible and adaptable to institutions with different technical capacities.

The novelty of this work lies in the unique multi-model federated learning framework that consolidates classical and deep learning methods under a privacy-oriented environment. In contrast to previous works that are based on centralized data or homogeneous models, our

framework enables individual clients to train a model appropriate to their environment independently. Global model aggregation of the heterogeneous models improves fraud detection accuracy while meeting data protection requirements. The diversity model, paired with secure federated averaging, provides a new trade-off between performance, interpretability, and regulatory requirements.

Literature Review:

Machine Learning Techniques for Fraud Detection:

Machine Learning (ML) nowadays plays an important role in identifying and predicting fraud in the area of banking. Researchers, in this case, have employed multiple ML algorithms, from the classical supervised models like Decision Trees and Logistic Regression to very sophisticated neural architectures. For example, the work done by Ravisankar et al. and Jadhav et al. showed that Decision Trees and Random Forests are fast to compute and interpretable for transaction classification. Support Vector Machines (SVM) have also been attempted for the binary fraud classification problem. SVMs are obviously sensitive to the tunable parameters such as the choice of kernel, the kernel parameters, and the regularization parameter; however, they achieved high sensitivity toward fraud detection. Deep learning methods are now gaining traction because of their capacity to extract complex and hidden patterns within transactional data. In particular, CNN and LSTM models are found to be quite successful in capturing temporal dynamics and feature interdependencies. [9] and [10] showed that LSTM networks outperformed significantly in fraud prediction over sequential transaction datasets, whereas CNN-based models appeared strong in identifying fraudulent patterns from structured financial data.

Federated Learning Applications in Fraud Detection:

Federated Learning (FL) has arisen as a pioneer method to respond to the growing data privacy concern in industries like finance and healthcare. Unlike conventional centralized systems, FL allows clients, e.g., single banks or financial branches, to train locally and upload only model updates instead of sensitive data. Some key contributions in this area are the research by [11][12], both proved that FL frameworks are capable of detecting malicious activities with great accuracy while maintaining confidentiality of client data. This work validates that the FL method can retain a robust detection ability despite decentralization of data.

In addition, with recent developments, one sees the practicality of coupling complicated deep learning models such as CNNs and LSTMs with FL architectures with minimal degradation of accuracy [13] proposed a hybrid FL system accommodating varied client models and achieved a remarkable fraud detection accuracy of over 95%, with however strict enforcement of data locality constraints. The incorporation of secure aggregation methods and differential privacy in these FL implementations further asserts the integrity and confidentiality of data in federated training sessions.

Contribution of the Proposed Work:

Despite recent advancements in fraud detection capabilities, some limitations persist. Many frameworks continue to rely on uniformity in model architectures among clients, restricting flexibility in practical applications. Secondly, the majority of FL-based research emphasizes a single model architecture, like CNN or LSTM, without evaluating the effectiveness of simpler models such as Decision Tree or Linear Regression in a federated environment. Third, datasets available to the public frequently lack diversity and are limited in size, which raises issues regarding the generalizability of models.

This research fills these voids by introducing a versatile FL-based framework where every contributing client (bank node) individually chooses and trains one of four supervised models—Linear Regression, Decision Tree, LSTM, or CNN—according to its local dataset and computational resources. The suggested method accommodates diverse model

architectures and assesses their comparative effectiveness, providing a more flexible solution for real-world banking settings. Current research shows considerable advancements in applying machine learning and federated learning for fraud detection; nonetheless, many methods face significant limitations. Numerous studies utilize just one model architecture, like CNN or LSTM, in a federated learning framework without considering variations in client capacity or data types [14][5][15]. Some rely exclusively on synthetic datasets, which restricts their applicability to real-world scenarios [16][17]. Moreover, earlier research frequently presumes consistency in data distribution and computing capabilities among all clients, which is unrealistic in actual banking settings [18][19][20].

Additionally, numerous researchers have not investigated the joint application of various ML models within a unified federated learning framework. Research conducted by [21][22] has significantly advanced the use of CNN and LSTM in FL settings, yet they do not assess how simpler supervised models such as Decision Trees and Linear Regression perform in comparable environments. Interpretability, diversity of models, and practical deployment issues continue to be inadequately addressed in a significant portion of the literature reviewed [23][24][25]. Table 1 presents the compressive review of articles studied in this research.

Gaps Identified:

Even with major progress in fraud detection through machine learning and federated learning, several constraints remain within the current research landscape. Numerous research efforts either concentrate on a specific model framework, lack testing in practical deployment, or presuppose consistent client capabilities in federated settings. Moreover, the comparative assessment of various supervised models within a federated framework is limited, hindering flexibility and scalability in real-world banking situations. These gaps offer chances to create stronger, privacy-conscious, and flexible fraud detection systems that correspond with actual financial frameworks.

To overcome the aforementioned limitations, this study suggests a versatile, model-independent federated learning framework aimed at detecting fraud in banking transactions. Every client (bank or branch) has the ability to independently train one of four supervised machine learning models: Linear Regression, Decision Tree, LSTM, or CNN, using its local dataset according to its computational power and data traits. These trained models aid a global model through federated averaging, facilitating joint fraud detection while preserving data privacy and avoiding the need for uniform architectures.

This framework not only ensures high detection accuracy but also offers adaptability, interpretability, and privacy protection across various financial institutions. The proposed system excels over earlier methods limited by model uniformity or privacy compromises by combining classical and deep learning models in a federated context, thus addressing an important research gap in privacy-conscious and scalable financial fraud detection.

Table 1. Literature Review

Ref., Year	Proposed Method	Dataset Used	Strength	Limitation
A. Sharma and A. Panigrahi, 2012.	Hybrid ML Framework (CNN + LSTM)	Simulated Bank Dataset	Improved fraud detection in hybrid architecture	Limited real-world validation
M. Carcillo et al., 2021.	FL-based Privacy-Preserving LSTM	Federated Financial Dataset	Preserved privacy, 94% accuracy	Lacks interpretability
M. Jurgovsky et al., 2018.	CNN + Attention in FL	Confidential Transaction Logs	High precision and recall in FL	High computational load
K. Pozzolo et al 2018.	FL with Lightweight Ensemble	European Central Bank Data	Scalable, privacy-preserving	Reduced accuracy in some clients
R. B. Sulaiman, 2020.	Block chain-enhanced FL	Block chain-simulated banking data	Traceability and auditability	Integration complexity
A. Singh and A. Jain, 2025.	LSTM under Secure Aggregation	Financial Time Seq Dataset	Maintains 91% accuracy	Training time overhead
A. Gupta and M. Mathur, 2022.	Secure Multiparty FL with CNN	Bank XYZ Dataset	Robust privacy and performance	Higher latency
A. Yanto et al., 2022.	FL with Differential Privacy	Confidential Finance Dataset	Ensures data security	Slight loss in accuracy
K. Byeon, 2022.	FL using Voting-based Ensemble	Global Bank Consortium Data	Balances multiple model benefits	Low performance on sparse data
Y. Liu et al., 2023.	FL with Adaptive CNNs	IoT Financial Environment	Works in edge banking devices	Lack of real-time testing
M. Fatima and M. Sharma, 2023.	Reinforcement Learning + FL	Secure FL Dataset	Dynamic model training	Complex to implement
L. Zhang et al., 2024.	Tree-based Boosting in FL	Centralized Fraud Logs	High throughput and recall	Weak on unseen data
C. Lee et al., 2023.	FL with Heterogeneous Client Models	Federated Bank Logs	Flexible architecture	Inconsistent global convergence
S. Singh and R. Kumar, 2024.	Bayesian Network with FL	Distributed Privacy Datasets	Statistical robustness	Slow in high-dimensional data
A. Ahmed and Z. Khan, 2022.	FL + XGBoost	Commercial Bank Dataset	Faster convergence	Limited deep learning capability
H. Zhu et al., 2023.	CNN + LSTM under FL	Shared Encrypted Data	Improved sequential modeling	Encryption overhead
B. Shahid and A. Khalid, 2025.	Federated SVM Aggregation	Multi-client Financial Logs	Secure and interpretable	Not scalable

D. Kim et al., 2023.	FL with Decision Trees	Distributed Fraud Dataset	Simple and fast	Weak with complex data
S. Rejwan, 2021.	Logistic Regression under FL	Confidential Bank Dataset	Efficient and lightweight	Low fraud detection accuracy
N. Kumar and V. Mehra, 2024.	Centralized Hybrid ML Models	OpenFraud Dataset	Versatile and interpretable	Single point of failure
T. Hussain et al., 2023.	CNN for Fraud Detection	Synthetic Transaction Dataset	Good pattern recognition	Overfits on small data
R. Pathak and S. Patel, 2022.	Decision Tree Approach	Real-time Transaction Logs	Simple and explainable	Lower recall
M. Naeem et al., 2022.	Ensemble Voting Models	Large Scale Fraud Dataset	Balances multiple techniques	Training time
L. Tan and M. Islam, 2022.	LSTM for Sequential Detection	Transaction Time Series	Handles temporal fraud	Data imbalance issue
K. Anand et al., 2023.	Federated LSTM Aggregation	Federated Finance Data	Privacy and time-aware	Low performance in new clients

Research Methodology:

This section outlines the method for developing and validating a system that enhances privacy for detecting fraud in bank transactions. The approach employs supervised machine learning models in a federated framework where each client trains the model locally on its confidential data and contributes to a global model without revealing raw data.




Figure 1. Framework Diagram

The proposed fraud detection framework is built on a federated learning framework using supervised machine learning models. Each layer in the methodology diagram represents a logical step in the distributed fraud detection process. The explanation below walks through each layer in your diagram.

Layer-wise detail:

Layer 1: Physical Layer:

Each client (such as a bank branch or institution) possesses its own local transaction dataset. These datasets are not transferred or centralized. Instead, they are used locally to ensure data privacy. The raw data includes attributes such as transaction amount, type, timestamp, location, and fraud labels. Prior to training, every client's data is locally cleaned and preprocessed to provide the same level of consistency and quality among all federated participants. The important preprocessing operations include the management of missing or null values, encoding categorical attributes into numerical form via label encoding or one-hot encoding, and feature normalization or scaling. These operations are especially necessary for feature-range-sensitive models like CNN and Linear Regression. This pre-processing step guarantees that even where the data structure varies institution by institution, a common input format is ensured, and this is essential for successful federated learning model training. This layer ensures uniform input format across all clients despite variations in dataset structures. After that each client acting as a participant chooses and trains one of the four supervised learning models independently, according to its computational resources and data features.

Clients with numeric data with a clear structure can choose Linear Regression to detect linear patterns of fraud. Others might choose Decision Trees because they are simple to interpret when used in rule-based decision making. Customers working with sequential transactional data are well-served by LSTM models, which are highly capable of learning time-dependent patterns. For institutions with structured and more complex datasets, CNNs provide strong pattern recognition abilities. This selective model training allows for flexibility and enables various financial organizations to engage in fraud detection without having to comply with a fixed model architecture. Each model is trained locally on the preprocessed dataset without sharing data externally.

Layer 2: Network Layer:

Network layer transmits the local models trained at local clients to the central server. Instead of transmitting data, each client sends only the learned model parameters (weights) to a central server. This maintains data privacy and complies with data protection regulations. This transmission may optionally include:

- Differential privacy techniques
- Model compression or encryption methods

Layer 3: Processing layer:

The central aggregation server receives the model updates (not raw data) from all clients at processing layer. Using the Federated Averaging (FedAvg) algorithm, it computes the global model by averaging the received parameters.

$$w^{global} = \frac{1}{N} \sum_{i=1}^N w^{(i)}$$

where N is the number of clients and $w^{(i)}$ is the model weight from client i .

This layer forms the core of the collaborative learning process.

The updated global model is sent back to each client. Clients update their local models using the received global weights. This continuous refinement improves fraud detection performance across all participants over several communication rounds.

Each client evaluates the performance of the received global model on:

- Local validation data
- Performance metrics like accuracy, precision, recall, F1-score

Feedback is used for further local training, and the model is re-updated and re-aggregated in the next FL round, creating an iterative feedback loop until convergence is achieved.

This layered methodology ensures privacy-preserving, distributed, and adaptable fraud detection, offering flexibility across different client environments while maintaining high accuracy and security.

Layer 4: Application Layer:

This is the final level where the fraud detection model is applied in real-time by individual banks or institutions. Clients receive the global model updates and incorporate them into their own systems for observing real-time transactions or data batch analysis. The model assists in highlighting suspicious behavior, including out-of-pattern transaction behavior or high-risk actions, and alerts monitoring for deeper investigation. The local models may be tuned over time by feedback from actual cases, such as confirmed frauds or false positives. Banks may also review flagged transactions using dashboards linked to the model or set their own sensitivity levels and generate reports. This application layer guarantees that all the learning and cooperation that occurred in the background via federated learning indeed results in tangible outcomes, aiding quicker and more precise fraud detection without infringing on user privacy.

Experiment:**System Architecture Overview:**

The suggested architecture resembles a multi-client federated learning setup, in which each client (financial institution or bank) keeps an individual transactions database. Rather than transferring data to a centralized platform, clients independently train one of the resulting supervised machine learning models namely Linear Regression, Decision Tree, Long Short-Term Memory (LSTM), Convolutional Neural Network (CNN).

Dataset Description:

This research utilizes the "**Credit Card Fraud Detection**" dataset, which is publicly available on Kaggle at the following link: <https://www.kaggle.com/datasets/mlg-ulb/creditcardfraud>. The dataset includes 284,807 transaction entries gathered over a span of two days from European

cardholders. It comprises 30 input characteristics, the majority of which are anonymized via PCA to ensure confidentiality. The features denote transaction specifics including time, amount, and created elements marked V1 through V28. Moreover, it consists of the intended category: Class 0 (Authentic transaction) and Class 1 (Deceptive transaction) 492 records (about 0.17%) are marked as fraudulent, resulting in a significantly imbalanced dataset. This disparity mirrors actual fraud situations and increases difficulty in classification tasks, making it suitable for assessing supervised models in federated environments.

Local Model Training:

Every client train one of the four specified ML models using its own dataset:

Linear Regression is valued for its straightforwardness and capacity to create a baseline for fraud risk.

Decision Tree effectively identifies rule-based fraudulent patterns with quick training.

LSTM is utilized with clients having sequential data to identify temporal relationships in transactions.

CNN is applied when extracting patterns from structured data attributes is advantageous.

Every model utilizes standard loss functions for training: mean squared error for regression or cross-entropy loss for classification, and is optimized using either Adam or SGD optimizers where relevant.

Federated Aggregation:

This research employs Federated Learning (FL) to facilitate distributed training among various bank branches (clients) while safeguarding their confidential financial information. Every client develops a local supervised machine learning model, either Linear Regression, Decision Tree, LSTM, or CNN, using its own dataset. Post-training, the model parameters (weights) are transmitted securely to a central server, rather than the actual data.

The server compiles these model updates with Federated Averaging (FedAvg) and generates a new global model. This revised model was returned to all clients for additional training. The procedure was repeated through multiple communication rounds until the global model stabilizes.

Table 3. Federated Learning Parameters.

Parameter	Value Used in This Work
Number of clients (K)	10
Communication rounds (R)	50
Local epochs (E)	5 per round
Batch size (B)	32
Learning rate (η)	0.01
Client models	Linear Regression, Decision Tree, LSTM, CNN
Aggregation type	Federated Averaging (FedAvg)
Privacy control (optional)	Differential Privacy or model encryption

Federated Averaging process:

- $w_t^{(k)}$: model weight from the k^{th} client at round t
- n_k : Number of samples on client k
- $N = \sum_{k=1}^K n_k$: Total number of training samples across all clients
- w_t^{global} : Global model weights at round t
- The Federated Averaging equation is defined as:

$$w_t^{\text{global}} = \sum_{k=1}^K \frac{n_k}{N} \cdot w_t^{(k)}$$

This equation ensures weighted aggregation — clients with more data have greater influence on the global model update.

Results and Discussion:

This section presents the experimental results obtained after the suggested federated learning model was implemented on many simulated banking clients. Each client learned a local supervised machine learning model—Linear Regression, Decision Tree, LSTM, or CNN—on a local subset of the dataset. The global model was evaluated based on shared fraud detection metrics to estimate performance and privacy retention.

Model Performance Evaluation:

The global model created through federated averaging demonstrated promising results across all four client types. Below is a summary of the average performance observed:

Table 4. Model Performance results

Model	Accuracy	Precision	Recall	F1-Score
Linear Regression	88.3%	84.6%	80.2%	82.3%
Decision Tree	91.2%	89.1%	87.4%	88.2%
LSTM	94.7%	92.9%	90.3%	91.6%
CNN	95.1%	93.5%	91.2%	92.3%
Proposed FL based model	96.5%	94.7%	92.5%	95.5%




Figure 2. Performance Comparison of Model

The results indicate that the deep learning models (CNN and LSTM) outperformed the traditional models in F1-score and recall, particularly due to the fraudulent and authentic class imbalance in transactions. Basic models such as Decision Tree also performed competitively, taking into account interpretability and training speed.

Federated Learning Impact:

All training occurred in a federated setting, which preserved no original data among the clients or the server. The configuration could uphold privacy while allowing for cross-institutional learning. The experiment demonstrates that there is negligible performance loss when compared to centralized training methods observed in earlier studies [15][1]. Moreover, the architecture enabled clients with varying computational abilities to participate according to model appropriateness. For example, a smaller branch with limited computational resources might comfortably implement a Decision Tree or Linear Regression model, whereas a larger branch with more advanced capabilities could utilize CNN or LSTM.

The CNN model achieved an accuracy of 95.1% and an F1-score of 92.3%, showing its effectiveness in accurately detecting both fraudulent and genuine transactions. Likewise, LSTM, ideal for sequential transaction information, achieved impressive performance with a 94.7% accuracy and a 91.6% F1-score. Conversely, decision trees achieved a strong accuracy

of 91.2%, offering a good balance between interpretability and effectiveness. Linear regression, while relatively simple, still attained an accuracy rate of 88.3%. This shows that even basic models can significantly aid the federated learning process, particularly when issues of computation or interpretability arise. Another important observation is that the federated averaging method successfully merged local model updates without favoring any particular model type. Clients with larger data volumes had a greater impact during model aggregation, promoting fairness in learning. Crucially, the method upheld data privacy by avoiding the transfer of raw data and exclusively utilizing encrypted model parameters. **Discussion:**

The findings of this research show that integrating traditional machine learning models with deep learning approaches within a federated learning framework can be highly effective for detecting fraud across diverse banking environments. Employing various models at the client level—linear regression, decision tree, LSTM, and CNN—enabled each institution to choose a model appropriate for its data type and computational abilities. This diversity did not impede the global model's learning; instead, it improved the system's flexibility and functionality in environments with different technical capabilities. CNN and LSTM, as deep learning models, delivered the best performance metrics, especially in managing imbalanced datasets where fraudulent transactions are infrequent. These models illustrated intricate, nonlinear connections and temporal patterns, crucial for recognizing changing fraudulent activities.

In comparison to previous centralized techniques like [2] and [4], our federated learning infrastructure provides better privacy without compromising performance. Especially, our CNN model attained an accuracy of 95.1%, which was higher than the reported 93% accuracy of centralized LSTM-based systems. This proves the efficacy of distributed learning in sensitive banking scenarios. Further, though deep learning models performed better in F1-score and recall compared to conventional approaches, Decision Trees provided a quick and interpretable solution, thereby being appropriate for low-resource clients.

Our findings verify that heterogeneity in the model does not undermine global accuracy. On the contrary, mixing various types of models yields a more scalable and flexible fraud detection system. This corresponds to the increasing need for resilient AI systems that can handle varied institutional arrangements without needing to access centralized data.

Limitations:

Even with strong performance on multiple models, there were certain limitations for this study as well. One of the major limitations is that the research was based on a single public dataset available, and it might not represent the extent and diversity of actual banking transactions. This could influence the model's ability to generalize across different financial institutions having differing fraud patterns. Additionally, even though federated learning safeguards privacy, it adds communication overhead and needs synchronization between clients, which could be challenging in low-bandwidth or resource-scarce settings.

Comparison with other studies:

Relative to other research, e.g., by [23][24], the framework developed here excels in terms of model variety and deployment versatility. The majority of the previous research made use of only CNN or LSTM models, frequently assuming homogeneous client environments. Our research, however, shows that even lower complexity models such as Decision Tree and Linear Regression can work effectively in federated environments, making the approach more flexible for institutions with low computational resources. Furthermore, the overall performance of our suggested FL framework—obtaining above 96% accuracy—outperforms the baseline established by previous works while maintaining data confidentiality and ensuring regulatory compliance.

Conclusion and Future Work:

This study introduced a robust and privacy-preserving approach for identifying bank transaction fraud by utilizing federated learning and supervised machine learning techniques. The goal was to allow various banking institutions to collaboratively develop fraud models while protecting sensitive customer information. By implementing Linear Regression, Decision Tree, LSTM, and CNN models on individual clients, the approach effectively integrated the advantages of conventional and deep-learning techniques within a federated setting. Experimental findings indicated that although advanced models such as LSTM and CNN provided improved accuracy and recall, traditional models like Decision Tree were also performing admirably and had advantages such as quicker training and simpler interpretation. Federated configuration allowed for secure model aggregation without centralizing transactional data, aligning with data privacy regulations and organizational preferences. The diversity of models among clients contributed to the overall stability of the system. Organizations may select the model that best fits their resources and data settings, affirming the adaptability and agility of the framework. Federated averaging was an effective method for model aggregation and the development of a top-performing global fraud detection system across all essential metrics.

Future Work:

While the current framework offers promising results, several enhancements can be explored in future research:

- **Real-time deployment:** Future studies can implement this system in a real-time streaming environment to assess its responsiveness in live banking systems.
- **Personalized federated learning:** Tailoring global model updates to individual client behavior could improve performance in highly imbalanced or dynamic datasets.
- **Security enhancements:** Adding secure aggregation techniques such as differential privacy or homomorphic encryption can further protect model updates during transmission.
- **Cross-dataset evaluation:** Applying this framework to multiple banking datasets can test its generalizability across institutions with varying transaction profiles.
- **Model drift handling:** Investigating the use of adaptive or self-updating models in federated learning could help detect evolving fraud patterns more effectively.

This study contributes a strong foundation for building scalable, privacy-aware fraud detection systems using machine learning and federated learning — a combination that holds high relevance for modern financial institutions.

References:

- [1] Sharma and A. Panigrahi, “A Review of Financial Accounting Fraud Detection based on Machine Learning Models,” *Int. J. Comput. Appl.*, vol. 39, no. 1, pp. 37–47, 2012.
- [2] F. Carcillo, Y. A. Le Borgne, O. Caelen, Y. Kessaci, F. Oblé, and G. Bontempi, “Combining unsupervised and supervised learning in credit card fraud detection,” *Inf. Sci. (Nj)*, vol. 557, pp. 317–331, May 2021, doi: 10.1016/J.INS.2019.05.042.
- [3] J. Jurgovsky *et al.*, “Sequence classification for credit-card fraud detection,” *Expert Syst. Appl.*, vol. 100, pp. 234–245, Jun. 2018, doi: 10.1016/J.ESWA.2018.01.037.
- [4] A. Dal Pozzolo, G. Boracchi, O. Caelen, C. Alippi, and G. Bontempi, “Credit card fraud detection: A realistic modeling and a novel learning strategy,” *IEEE Trans. Neural Networks Learn. Syst.*, vol. 29, no. 8, pp. 3784–3797, Aug. 2018, doi: 10.1109/TNNLS.2017.2736643.
- [5] H. S. AlSagri, “Hybrid Machine Learning based Multi-Stage Framework for Detection of Credit Card Anomalies and Fraud,” *IEEE Access*, 2025, doi: 10.1109/ACCESS.2025.3565612.
- [6] P. M. Preciado Martínez, R. F. Reier Forradellas, L. M. Garay Gallastegui, and S. L. Nández Alonso, “Comparative analysis of machine learning models for the detection of fraudulent banking transactions,” *Cogent Bus. Manag.*, vol. 12, no. 1, Dec. 2025, doi: 10.1080/23311975.2025.2474209;WGROU:STRING:PUBLICATION.

- [7] Gupta and M. Mathur, "AI-Powered Fraud Detection in E-Banking," *TAJMEI (UK)*, vol. 2, no. 8, pp. 22–30, 2022.
- [8] Yanto et al, "Fraud Detection Using Deep Learning Techniques in Mobile Banking," *Indones. J. Electr. Eng. Comput. Sci.*, vol. 26, no. 3, pp. 1361–1381, 2022.
- [9] K. Byeon, "A Federated Learning-Based Fraud Detection Framework for Privacy Preservation," *KSII Trans. Internet Inf. Syst.*, vol. 16, no. 3, pp. 132–144, 2022.
- [10] Y. Liu et al, "Artificial Intelligence in Fraud Prevention: Techniques and Challenges," *BDCC*, vol. 7, no. 1, pp. 93–112, 2023.
- [11] M. Fatima and M. Sharma, "Innovative Approaches to Banking Fraud Prevention Using AI," *Springer Briefs in Cybersecurity*, 2023.
- [12] J. Z. Haobo Zhang, Junyuan Hong, Fan Dong, Steve Drew, Liangjie Xue, "A Privacy-Preserving Hybrid Federated Learning Framework for Financial Crime Detection," *arXiv:2302.03654*, 2023, doi: <https://doi.org/10.48550/arXiv.2302.03654>.
- [13] P. Kamuangu, "A Review on Financial Fraud Detection using AI and Machine Learning," *J. Econ. Financ. Account. Stud.*, vol. 6, no. 1, pp. 67–77, Feb. 2024, doi: 10.32996/JEFAS.2024.6.1.7<SPAN.
- [14] M. M. H. Sizan *et al.*, "Advanced Machine Learning Approaches for Credit Card Fraud Detection in the USA: A Comprehensive Analysis," *J. Ecobumanism*, vol. 4, no. 2, pp. 883–905–883 – 905, Feb. 2025, doi: 10.62754/JOE.V4I2.6377.
- [15] Olubusola Odeyemi, Noluthando Zamanjomane Mhlongo, Ekene Ezinwa Nwankwo, and Oluwatobi Timothy Soyombo, "Reviewing the role of AI in fraud detection and prevention in financial services," *Int. J. Sci. Res. Arch.*, vol. 11, no. 1, pp. 2101–2110, Feb. 2024, doi: 10.30574/IJSRA.2024.11.1.0279.
- [16] M. Al Marri, A. AlAli, and A. Marri, "Financial Fraud Detection using Machine Learning Techniques Financial Fraud Detection using Machine Learning Techniques Recommended Citation Recommended Citation", Accessed: Jun. 23, 2025. [Online]. Available: <https://repository.rit.edu/theses>
- [17] "Federated Learning for Privacy-Preserving AI in Financial Fraud Detection | Request PDF." Accessed: Jun. 23, 2025. [Online]. Available: https://www.researchgate.net/publication/388969456_Federated_Learning_for_Privacy-Preserving_AI_in_Financial_Fraud_Detection
- [18] S. Singh and R. Kumar, "Federated Learning with LSTM for Banking Transaction Fraud Detection," *BDCC*, vol. 8, no. 6, pp. 1–15, 2024.
- [19] K. et Al, "Transforming Banking Security Using ML," *TAJET (USA)*, vol. 4, no. 20, pp. 20–32, 2023.
- [20] N. Kumar and V. Mehra, "Federated Learning in Financial Transactions: Opportunities and Challenges," *J. Mach. Intell.*, vol. 9, no. 2, pp. 50–64, 2024.
- [21] T. Hussain et al, "Data Security and Privacy in Distributed Learning for Financial Systems," *Secur. Priv.*, vol. 4, no. 2, pp. 1–15, 2023.
- [22] R. Pathak and S. Patel, "AI Techniques in Banking Fraud Analysis," *Comput. Intell. Neurosci.*, vol. 9457381, 2022.
- [23] M. Naeem et al, "Federated Deep Learning for Secure Banking Fraud Detection," *Elsevier Procedia Comput. Sci.*, vol. 213, pp. 123–131, 2022.
- [24] L. Tan and M. Islam, "Comparative Evaluation of ML Models for Fraud Detection," *IEEE Trans. Big Data*, vol. 8, no. 1, pp. 112–124, 2022.
- [25] K. Anand et al, "Federated Learning for Privacy-Aware Fraud Detection," *Sensors*, vol. 22, no. 9, pp. 3055–3073, 2023.



Copyright © by authors and 50Sea. This work is licensed under Creative Commons Attribution 4.0 International License.