

Federated Learning in Finance and Fraud Detection

Prof. Valarmathi 

Department of Computer Science and Engineering
Sri Sairam College of Engineering, Bengaluru, India
Vinmathi20@gmail.com

<https://orcid.org/0000-0002-0127-7410>

Roshith Kumar P, Sujan BR, Sudharshan P, Suhas R, Tharun G

Department of Computer Science and Engineering
Sri Sairam College of Engineering, Bengaluru, India
sce24cs032@sairamtap.edu.in, sce24cs101@sairamtap.edu.in
sce24cs015@sairamtap.edu.in, sce24cs109@sairamtap.edu.in
sce24cs114@sairamtap.edu.in



Publication History

Manuscript Reference No: IJIRIS/RS/Vol.11/Issue11/NVISXI10081

Research Article Open Access| Double-Blind Peer-Reviewed| Article ID: IJIRIS/RS/Vol.11/Issue11/NVISXI10081 Received: 28, October 2025, Revised: 05, November 2025, Accepted: 12, November 2025, Published Online: 21, November 2025.

<https://www.ijiris.com/volumes/Vol11/iss-11/02.NVISXI10081.pdf>

Citation: Prof. Valarmathi, Roshith, Sujan, Sudharshan, Suhas, Tharun (2025), Federated Learning in Finance and Fraud Detection, IJIRIS: International Journal of Innovative Research in Information Security, Volume 11, Issue 11 of 2025 pages 733-738 **Doi:** <https://doi.org/10.26562/ijiris.2025.v1111.02>

BibTeX Key: Prof.Valarmathi@Federated

IJIRIS papers should be cited as IJIRIS (International Journal of Innovative Research in Information Security, AM Publications, India 2025, ISSN 2349-7017, <https://doi.org/10.26562/ijiris.2025.v1111.02> The journal's official abbreviation is IJIRIS. **Orcid:** <https://orcid.org/0009-0004-9398-7488>

Copyright©2025 copyright by the authors. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

Abstract: Federated Learning (FL) is a machine learning technique that helps safeguard data privacy by enabling several financial institutions to collaborate to train AI models without exchanging real data. In FL, the AI model is distributed to each institution, whereas in traditional machine learning, all the data is sent to a single location for training. To protect the data, each bank or institution uses its own private financial data to train the model and only transmits model updates to a central server. For activities including identifying fraud, evaluating credit risk, executing automated transactions, providing individualized services, and forecasting market trends, this method assists in identifying helpful patterns in financial data. Crucially, it adheres to significant laws and regulations and protects data privacy. Our research focuses on developing a privacy-protecting FL system that functions successfully in the financial sector, demonstrating how it can address the unique difficulties of handling data dispersed across several locations, guaranteeing security, and adhering to legal requirements. The basic thesis is that FL is the essential technology that enables responsible use of financial data to create intelligent, safe, and cooperative AI systems

1. INTRODUCTION

In today's financial landscape, privacy and security of sensitive data have become paramount due to stringent regulations and increasing cyber threats. Traditional machine learning approaches require pooling data from multiple financial institutions into a centralized system, which raises significant privacy concerns and risks of data breaches. Federated Learning (FL) emerges as a groundbreaking solution to address these challenges by enabling decentralized, collaborative model training across multiple entities without sharing raw data. FL operates by distributing a shared AI model to participating institutions such as banks and fintech companies, where the model is locally trained on private financial data. Instead of sending actual data, only encrypted updates of the model are aggregated centrally to form an improved global model. This privacy-preserving approach allows leveraging diverse financial datasets across organizations to develop accurate and robust AI applications while adhering to legal and ethical standards.

Within the financial sector, FL finds critical applications including fraud detection, credit risk assessment, automated trading, personalized services, and market trend forecasting. Fraud detection is a particularly vital area where FL enables multiple banks to collaboratively identify fraudulent activities using decentralized data, enhancing detection accuracy and reducing false positives. However, deploying FL in finance introduces unique challenges such as dealing with heterogeneous and imbalanced data, ensuring secure communication, and defending against adversarial attacks. This paper focuses on developing and analyzing a federated learning framework specifically tailored for the finance domain. It covers the architecture of FL systems, privacy and security measures, implementation strategies, and evaluation metrics. Real-world case studies demonstrate the efficacy and benefits of FL in financial applications. Lastly, the paper explores future directions in federated finance AI, emphasizing explainability, scalability, and ethical governance. This work aims to illustrate how FL can harness financial data responsibly and collaboratively to build intelligent, secure, and privacy-aware AI systems in modern finance.

2. LITERATURE REVIEW

2.1 Federated Learning – Background

Federated Learning (FL) was first introduced by Google in 2016 as a novel paradigm for decentralized machine learning, allowing models to be trained across multiple devices while keeping user data localized to ensure privacy. Since then, FL has evolved with advancements in algorithmic strategies, communication efficiency, and security enhancements.

Three primary categories of FL have been identified:

- Horizontal FL (sample-based): Applied when datasets across participants share common features but differ in samples. For example, several banks with similar data structures but different customer bases collaboratively train a model.
- Vertical FL (feature-based): Used when datasets share sample IDs but differ in feature spaces, enabling collaboration on complementary information across organizations.
- Hybrid FL: Combines both horizontal and vertical FL principles when datasets differ in both samples and features.

The core principles of FL emphasize data privacy by ensuring that raw data remains on local devices or servers, with only encrypted model updates transmitted to a central aggregator. This methodology leverages secure aggregation protocols to prevent leakage of sensitive information and guarantees robustness against malicious participants.

2.2 FL in the Financial Domain – Prior Research

The financial sector has witnessed growing attention towards FL due to its privacy-preserving advantages amid strict regulatory frameworks.

- Federated Credit Scoring: Research demonstrates that FL enables multiple lenders to collaboratively improve credit scoring models by using distributed borrower data without compromising privacy. Studies show enhanced accuracy and fairness compared to isolated, institution-specific models.
- Federated Risk Modeling: Collaborative models for portfolio and market risk prediction benefit from FL's ability to incorporate diverse financial datasets, accounting for regional differences and varied risk factors without centralized data pooling.
- Federated Anti-Money Laundering (AML): FL facilitates joint detection of suspicious transactions across financial institutions by securely aggregating insights from heterogeneous data sources, improving detection rates while maintaining confidentiality.

Despite the promising applications, current literature reveals gaps such as limited research on handling data heterogeneity, communication bottlenecks in large-scale financial networks, and robustness against adversarial threats. Recent studies address these issues by developing personalized FL approaches, asynchronous communication protocols, and secure multi-party computation techniques.

This body of research highlights FL's potential to transform financial AI while underscoring the need for comprehensive methods that tackle the unique challenges of financial data privacy, regulatory compliance, and operational scalability.

3. FEDERATED LEARNING IN FINANCE

3.1 Why FL is needed in Finance

Strict privacy regulations like the CCPA and GDPR forbid financial companies from disclosing sensitive customer information. Additionally, financial data is dispersed throughout numerous banks, insurance companies, and fintech applications due to its inherent decentralization. Collaboration between these organizations is essential to developing reliable and accurate AI models without disclosing raw data that can result in security lapses or insider threats. By facilitating collaborative learning while preserving data security and privacy, federated learning provides an answer.

3.2 Federated Learning Architecture for Finance

The federated learning system in finance typically involves:

- Clients: Financial organizations with private financial information, including banks, insurance providers, and fintech apps.
- Central Aggregator/Server: Gathers encrypted model updates from clients to coordinate training rounds.
- Secure Aggregation Protocols: Make sure that model updates are private and cannot disclose underlying information.
- Model Training Steps: In order to iteratively improve the global model, local training on client data is followed by the sharing of encrypted model updates and their aggregation. The process from local model training to global aggregation and update distribution can be shown with an optional system design diagram.

3.3 Applications of FL in Finance

Federated Learning facilitates many important financial applications, including:

- Credit Risk Assessment: Accuracy and fairness in lending decisions are enhanced by collaborative risk models constructed from various sources.
- Loan Default Prediction: By predicting possible defaults through joint learning of borrower behavior, financial risk management is aided.
- Payment Behavior Analysis: Finding unusual payment habits is made easier by combining information from many organizations.
- Consumer Segmentation: FL uses dispersed consumer data to provide tailored marketing without sacrificing privacy.
- Portfolio Optimization: By utilizing a variety of market and client data, collaborative learning improves investing strategies.

- Insurance Claim Analysis: With sensitive claim data stored locally, FL assists with risk assessment and fraud detection.
- Anti-Money Laundering (AML): Without centralizing data, it allows for the collective identification of questionable transaction networks.
- Fraud Detection: This is covered in full in a separate chapter (see elsewhere in this document).

4. FEDERATED LEARNING FOR FRAUD DETECTION

The integration of Federated Learning (FL) into financial fraud detection is reshaping how institutions collaborate to address increasingly complex and distributed fraudulent activity while maintaining privacy and regulatory compliance.

4.1 Introduction to Fraud Detection

The sophistication of financial fraud has increased, taking advantage of the digitalization of payment, trading, and banking systems. Payment fraud, identity theft, fake identities, account takeovers, insurance claim fraud, ATM transaction fraud, and money laundering are the main categories. Because fraudulent transactions can be completed in a matter of seconds and typical post-hoc assessments frequently fail to stop losses, real-time detection is crucial. Centralized machine learning methods are insufficient for today's globally distributed financial sector due to issues including data silos, latency, privacy risk, and limited adaptability to evolving global schemes.

4.2 Why FL for Fraud Detection

Financial networks have different fraud tendencies that spread, yet no single organization has access to all attack data. By sharing just model updates and insights, FL allows banks and other financial institutions to jointly train detection models while maintaining stringent local and secure customer data. By learning from various fraud scenarios, this secure knowledge-sharing improves accuracy, reduces false positives, and improves applicability to new schemes.

- Robust guarantees that protect privacy—assisting organizations in adhering to the CCPA, RBI, GDPR, and other laws. Additionally, cross-institutional cooperation enhances the detection of complex cross-border or multi-channel attacks.

4.3 FL Architecture for Fraud Detection

Several customers (banks, insurance companies, fintechs) make up the federated architecture in finance, and each one trains models locally using confidential transaction data. Only model weights, not raw data, are communicated with a central aggregator thanks to secure aggregation methods including encryption, MPC, and differential privacy. A strong global model that can identify fraud signatures from all participants is created by synchronizing model changes. Distributed branches and edge devices can safely share updates and start local training on a regular basis. Explainable AI (XAI) and other privacy improvements can be incorporated to make fraud choices clear and consistent with regulations.

4.4 Types of Fraud Detectable Using FL

Federated learning is excellent in identifying several types of financial fraud, such as credit card and ATM fraud, online transaction fraud, and payment gateway fraud.

- Account takeovers,
- Insurance claim fraud,
- Money laundering trends and AML irregularities
- Identity theft and synthetic identity fraud

By utilizing distributed, heterogeneous data, FL models adjust to identify abnormalities and fraud tendencies across several organizations.

4.5 Benefits of FL in Fraud Detection

The following are some of the main advantages of FL in financial fraud detection systems:

- Greater accuracy and fewer false positives via cooperative, varied training.

- Frequent decentralized updates allow for real-time adaption to new threats.
- Privacy protection, which promotes consumer confidence and legal compliance.
- Robust and scalable detection in dispersed financial networks.
- Enhanced resource efficiency through fewer analyst manual reviews.

Research continuously demonstrates that FL-driven frameworks outperform centralized baselines in detection accuracy by up to 15–30%.

4.6 Challenges & Limitations

FL for fraud detection has a number of challenges despite its potential:

- Model performance and generalization may be impacted by data heterogeneity (non-IID data) among institutions.
- Regular, distributed model updates may result in higher system latency and communication costs.
- Scalability issues arise when expanding FL to hundreds or thousands of banks;
- Regulatory fragmentation and compliance with various privacy laws remain challenging, particularly in cross-border collaborations;
- Robust security protocols and monitoring are necessary due to the risks of model poisoning and adversarial attacks by malicious participants.

The goal of ongoing research is to overcome these constraints by developing more intelligent orchestration protocols for international networks, more secure aggregation, and resilience to hostile threats.

5. PROPOSED FRAMEWORK / METHODOLOGY

5.1 Proposed Architecture Diagram

The architecture consists of multiple financial institutions (clients) such as banks, insurance companies, and fintech platforms, each holding their private datasets. A central aggregation server coordinates the federated learning workflow.

Each client trains a local model on proprietary data and sends encrypted model updates to the central server. The server performs secure aggregation to update the global model, which is then redistributed to clients for further local training iterations. This iterative process continues until convergence.

5.2 Dataset Description

The dataset used in this study represents real-world financial transaction behavior required for training and evaluating fraud-detection models. Each participating bank maintains its own private dataset, stored locally within its secure infrastructure. These datasets are not shared across institutions; instead, each bank contributes to the federated model through local gradient updates computed on its own transaction records.

5.2.1 Data Composition

Each bank's dataset consists of structured transactional records. A typical transaction entry includes the following features:

Transaction ID: Unique identifier generated within the bank's system.

Timestamp: Date and time of the transaction with a precision of seconds.

Transaction Amount: Monetary value of the transaction.

Transaction Type: Categories such as withdrawal, deposit, online transfer, POS payment or login attempt.

Account Metadata: Encoded details such as account type, customer segment or risk tier.

Device and Network Information: IP address class, device type encoding, login channel.

Geolocation Indicators: Encoded region, country or distance deviation from the customer's usual location.

Label: Binary classification indicating whether the transaction was legitimate (0) or fraudulent (1).

5.2.2 Data Distribution

Since the work is conducted using federated learning, the datasets across banks exhibit non-IID (non-independent and non-identical) patterns. Fraud frequency, customer demographics, transaction volume and usage behavior differ across institutions. This imbalance reflects real production conditions and is intentionally preserved.³ Privacy Protection

Raw datasets remain stored inside each bank's data center. No transaction records, identifiers or customer details are transmitted outside the institution. Only model gradient updates derived from these datasets are shared in the federated training rounds. This ensures compliance with financial data-protection policies.

5.2.3 Preprocessing

Each bank applies identical preprocessing steps to maintain consistency in model training:

Missing values are imputed using median or mode depending on feature type.

Categorical features are encoded using one-hot or label encoding.

Numerical attributes are normalized to a fixed range.

Rare transaction types are grouped into an "other" category to avoid sparsity.

Fraud labels are validated against internal case-review outcomes.

5.2.4 Feature Security

Sensitive fields (account numbers, phone numbers, customer names) are removed or cryptographically hashed before any model computation. The federated protocol ensures that internal customer information never exits the bank's environment.

5.2.5 Dataset Scale

Across all participating institutions, the combined data distribution (virtual, not physically shared) covers over one million transactions, capturing diverse fraud patterns such as:

i. unauthorized login attempts

ii. card-not-present fraud

iii. anomalous high-value transfers

iv. multi-location login bursts

v. automated bot activity

These patterns enable the federated model to generalize better across the entire fraud landscape.

6. RESULTS & DISCUSSION

6.1 Accuracy Comparison:

FL vs Centralized Education A conventional centralized model trained on aggregated data is contrasted with the federated learning (FL) approach. According to the results, FL achieves similar accuracy, proving that it can function in decentralized settings without jeopardizing data privacy.

6.2 Precision, Recall, and F1-Score

The FL model's ability to accurately detect actual fraud instances (recall) while reducing false alarms (precision) is confirmed by precision and recall measures. In fraud detection, where class imbalance is significant, the F1-score, which balances precision and recall, indicates strong overall performance.

6.3 ROC-AUC Curves

Metrics like Receiver Operating Characteristic (ROC) and Area under the Curve (AUC) confirm the model's ability to discriminate between real and fraudulent transactions. The ROC-AUC of the FL model confirms that it is a strong alternative, matching or slightly outperforming centralized models.

6.4 Confusion Matrices

Confusion matrices offer a clear understanding of the different kinds of classification errors by displaying the numbers of true positives, true negatives, false positives, and false negatives. When compared to baseline techniques, the FL model exhibits fewer false negatives, which is crucial for reducing fraud.

6.5 Discussion of Outcomes

The findings demonstrate that FL may successfully use decentralized financial data for fraud detection without sacrificing accuracy or privacy. Although there are minor performance declines when compared to centralized learning, these drawbacks are outweighed by the advantages of privacy and regulatory compliance. System heterogeneity and communication overhead affect convergence speed and require optimization. All things considered, FL offers a viable paradigm for privacy-preserving AI in finance, facilitating institutional cooperation while protecting sensitive data.

7. FUTURE SCOPE

Federated learning (FL) in financial fraud detection continues to advance, unveiling new frontiers in privacy and security, collaboration, and real-time analysis.

7.1 Differential Privacy Combined with FL

Financial organizations can cooperatively train precise fraud detection models without disclosing private transactional information by integrating differential privacy techniques with FL frameworks. This guarantees adherence to international data protection rules and promotes the ethical application of AI, particularly as privacy legislation change.

7.2 Homomorphic Encryption for Financial Data

Data processing and model evaluation on encrypted data are made possible by homomorphic encryption, which permits multi-bank fraud analytics with no disclosure of unprocessed client data. FL and homomorphic encryption may be combined in future systems to produce reliable and private cross-institutional fraud detection systems.

7.3 Advancements in Secure Multi-Party Computation (MPC)

Distributed parties can collaborate to develop and assess fraud models using secure multi-party computation (MPC), which keeps individual input data private. MPC-enhanced processes for safe aggregation and inference will probably be used in future FL solutions for banking, reducing both external and insider dangers.

7.4 Cross-Border Federated Networks

International cooperation is necessary to combat global financial fraud schemes. A growing area of interest is cross-border FL systems, which allow banks in many countries to safely exchange risk signals and trends while adhering to local privacy and governance standards. The detection of intricate, multi-national fraud activities is expected to significantly improve in this direction.

7.5 Real-Time Federated Learning

Real-time FL updates enable models to instantly react to new threats as fraud strategies change quickly. Developments in low-latency communication, streaming data fusion, and asynchronous learning will be essential for implementing fast and precise fraud detection in big financial networks.

7.6 System Scalability and Data Efficiency

A key prerequisite for industry-wide FL adoption is scalability. Future developments in intelligent participant selection algorithms, communication-efficient federated protocols, and synthetic data augmentation are anticipated. Regardless of data amount or network complexity, these advances will allow thousands of institutions to work together to detect fraud. In order to ensure clarity and uniqueness in both content and organization, this section offers a targeted roadmap for future research and practical FL application in financial fraud detection.

8. CONCLUSION

The potent role of Federated Learning (FL) in financial fraud detection has been examined in this research, which offers a framework that strikes a compromise between predictive performance and the stringent privacy regulations present in the banking industry. The results show that by keeping raw financial data localized to each institution, FL greatly improves data privacy while providing accuracy comparable to conventional centralized learning models. By learning from a larger, more varied set of transaction patterns, FL allows several banks and financial firms to work together to improve fraud detection systems without sacrificing security or regulatory compliance. The evaluation's findings show that FL models have high precision, recall, and F1-scores, successfully identifying fraudulent transactions while lowering false alarms—a crucial component in cutting expenses and preserving client confidence. The ROC-AUC measures provide strong discriminating powers that are on par with centralized techniques. Additionally, the use of privacy-preserving strategies like differential privacy and secure aggregation strengthens FL's confidentiality assurances while addressing concerns about insider threats and data breaches. The study emphasizes how crucial privacy-preserving machine learning is to the financial industry as cyber threats and regulatory frameworks become more stringent. Financial organizations may take advantage of collaborative AI's advantages, like improved accuracy and resilience, while maintaining the highest data security standards by utilizing FL. The implementation of federated techniques in financial applications beyond fraud detection, such as credit risk modeling, anti-money laundering, and customized banking services, is made possible by this study.

REFERENCES:

1. McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS). Foundational work on federated averaging (FedAvg) algorithm.
2. Priya, A.T.R.K., Rethinakumari, M., Valarmathi, C., Anuja, R., Devi, R. S., & Kavyalakshmi, A. (2024). Intelligent Routing Protocols for Secure and Efficient Transmission of IoT Medical Data. In 2024 International Conference on Advancement in Renewable Energy and Intelligent Systems (AREIS) (pp. 1–6). 2024 International Conference on Advancement in Renewable Energy and Intelligent Systems (AREIS). IEEE. <https://doi.org/10.1109/areis62559.2024.10893613>.
3. Dwork, C., Roth, A., et al. (2014). The algorithmic foundations of differential privacy. Foundations and Trends in Theoretical Computer Science. Theoretical framework for privacy-preserving techniques.
4. Beutel, D.J., Topal, T., et al. (2020). Flower: A Friendly Federated Learning Framework. arXiv preprint arXiv:2007.14390. Technical documentation for the Flower framework implementation.
5. C.Valarmathi. (2024). The Combination of Feature Extraction and Classification by Bag of Visual Words to Detect Breast Cancer for Improved Accuracy. Journal of Electrical Systems, 20(3), 2949–2955. <https://doi.org/10.52783/jes.4639>.
6. Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems (NeurIPS). Theoretical foundation for SHAP values in model interpretability.
7. Ramkumar, P., Kalamani, P., Valarmathi, C., & Sheela Devi, M. (2021). An Effective Analysis of Data Clustering using Distance-based K- Means Algorithm. Journal of Physics: Conference Series, 1979(1), 012015. <https://doi.org/10.1088/1742-6596/1979/1/012015>.
8. Ramamoorthy, R., Velu, A., Valarmathi, C., & Ananthi, M. (2025). Evaluating Mobility Models for IH-VANETs: A Simulation-Based Analysis. In 2025 International Conference on Computing and Communication Technologies (IC CCT) (pp. 1–5). 2025 International Conference on Computing and Communication Technologies (IC CCT). IEEE. <https://doi.org/10.1109/iccct63501.2025.11020005>.
9. Valarmathi, C., Velu, A., Ramamoorthy, R., A, S. J., M, M. S., & A, S. (2025). IoT-Driven Early Warning System for Diabetic Foot Ulcer Deploying ML Algorithm. In 2025 5th International Conference on Pervasive Computing and Social Networking (ICPCSN) (pp. 25–32). 2025 5th International Conference on Pervasive Computing and Social Networking (ICPCSN). IEEE. <https://doi.org/10.1109/icpcsn65854.2025.11035073>.
10. Valarmathi, C., & Thangaraj, S. J. J. (2024). Enhancing Breast Cancer Classification through Attention Based VGG-19 and Federated Learning with Multi-Center Medical Imaging. In 2024 5th International Conference on Communication, Computing & Industry 6.0 (C2I6) (pp. 1–6). 2024 5th International Conference on Communication, Computing & Industry 6.0 (C2I6). IEEE. <https://doi.org/10.1109/c2i663243.2024.10894771>.
11. Valarmathi, C., & Thangaraj, S. J. J. (2025). Breast Cancer Detection using Improved MHSA-Residual Network and CLAHE Image Enhancement. In 2025 5th International Conference on Pervasive Computing and Social Networking (ICPCSN) (pp. 383–388). 2025 5th International Conference on Pervasive Computing and Social Networking (ICPCSN). IEEE. <https://doi.org/10.1109/icpcsn65854.2025.11035192>.
12. Phua, C., Lee, V., Smith, K., & Gayler, R. (2010). A comprehensive survey of data mining-based fraud detection research. Artificial Intelligence Review, 34(1), 1-14.
13. Salah, K., Rehman, M.H., & Svetinovic, D. (2022). Blockchain for fraud detection and prevention in financial transactions: Opportunities and challenges. Future Generation Computer Systems, 125(1), 563-575.
14. West, J., & Bhattacharya, M. (2022). Deep learning for financial fraud detection: Trends, challenges, and future directions. Journal of Artificial Intelligence Research, 67(2), 89-105