Available online at www.qu.edu.iq/journalcm
JOURNAL OF AL-QADISIYAH FOR COMPUTER SCIENCE AND MATHEMATICS
ISSN:2521-3504(online) ISSN:2074-0204(print)

Federated Learning for Credit Card Fraud Detection: A Privacy-Preserving Approach with Smote Optimization

Hassan W. Hilou [1], Mohsin A. Ahmed [2], Shaymaa A. Dheeb [3], Ahmed A. Radhi[4], Zainab M. Khadim[5], Mohammed N. Majeed[6], Ali Ismail Jadaan[7]

[1] Department of computer engineering technology, Al Mamoun University College, Baghdad, Iraq.hassan.w.hilou@almamonuc.edu.iq
[2] Department of computer engineering technology, Al Mamoun University College, Baghdad, Iraq. mohsin.a.ahmed@almamonuc.edu.iq
[3] Department of computer engineering technology, Al Mamoun University College, Baghdad, Iraq. shaymaa.a.dheeb@almamonuc.edu.iq
[4] Department of computer engineering technology, Al Mamoun University College, Baghdad, Iraq. ahmed.a.radhi@almamonuc.edu.iq
[5] Department of computer engineering technology, Al Mamoun University College, Baghdad, Iraq. zainab.m.khadim@almamonuc.edu.iq
[6] Department of computer engineering technology, Al Mamoun University College, Baghdad, Iraq. mohammed.n.majeed@almamonuc.edu.iq
[7] Department of Electronics Engineering, National University for Science and Technology (NUST), Baghdad, Iraq. aliisiq@yahoo.com

ARTICLE INFO

Article history:
Received: 26/06/2025
Rrevised form: 14/07/2025
Accepted : 21/07/2025
Available online: 30/09/2025

Keywords:

Credit Card Fraud Detection, Federated Learning, Class Imbalance

ABSTRACT

Credit card fraud poses significant global challenges, costing financial institutions billions annually while evolving in sophistication. Traditional machine learning approaches for fraud detection face limitations in data privacy and scalability when dealing with distributed transaction data across multiple institutions. This paper presents a novel federated learning framework for credit card fraud detection that addresses these challenges while maintaining detection accuracy. Our approach leverages distributed machine learning across multiple client institutions (ranging from 5 to 260 clients) without requiring direct data sharing, thus preserving privacy. Through extensive experimentation, we demonstrated that our federated model achieved consistent accuracy (99% with 5–50 clients; 95–98% with 100+ clients) while effectively handling class imbalance through SMOTE integration (AUC = 1.00). The system showed particular effectiveness at mid-range client participation (10–50 clients), establishing an optimal balance between detection performance and computational efficiency. Compared to traditional centralized approaches and alternative data balancing methods (undersampling AUC = 0.97, oversampling AUC = 0.99), our federated solution provides superior privacy preservation without compromising fraud detection capability. The results indicate that this framework offers financial institutions a practical, scalable solution for collaborative fraud detection while maintaining strict data confidentiality requirements.

https://doi.org/10.29304/jqcsm.2025.17.32423

1. Introduction

Credit card fraud remains a critical challenge for financial institutions worldwide, costing billions annually and evolving in sophistication alongside digital payment ecosystems [1]. While machine learning (ML) has emerged as a powerful tool for fraud detection, conventional centralized approaches face two fundamental limitations: (1) the sensitivity of transaction data restricts cross-institutional collaboration, and (2) extreme class imbalance (often <0.1% fraud instances [2]) skews model performance. Federated learning (FL) offers a promising alternative by enabling collaborative model training across banks without raw data sharing [3]. This study investigates the

∗Corresponding author: Ali Ismail Jadaan

Email addresses: *aliisiq@yahoo.com*

Communicated by 'sub etitor'

viability of Federated Averaging (FedAvg) [4] for credit card fraud detection, comparing it to centralized Multi-Layer Perceptron (MLP) baselines while addressing privacy, scalability, and class imbalance challenges.

Traditional centralized ML relies on isolated datasets, limiting models' exposure to diverse fraud patterns. For instance, Bank A's model trained solely on its local data may fail to detect fraud types encountered by Bank B. While data pooling could mitigate this, privacy regulations and competitive concerns prohibit direct sharing of sensitive transaction records [2]. Furthermore, the extreme rarity of fraudulent transactions (<0.1% in typical datasets [10]) leads to class imbalance, requiring techniques like SMOTE to prevent model bias toward majority classes [3].

Federated learning addresses these issues by decentralizing model training. Banks collaboratively train a global model (e.g., via FedAvg [4]) by sharing encrypted parameter updates—not raw data—preserving privacy while aggregating knowledge from all participants. This approach expands the model's exposure to diverse fraud patterns without violating data sovereignty. However, FL introduces new challenges, including communication overhead at scale [3] and the need for robust aggregation methods to handle non-IID (non-independent and identically distributed) data across institutions.

This work makes the following key contributions:

1. Comparative Analysis: A systematic evaluation of FedAvg against centralized MLP for fraud detection, demonstrating FL's ability to achieve ~99% accuracy with 5–50 clients while preserving data privacy.

2. Class Imbalance Mitigation: Integration of SMOTE in FL, showing superior performance (AUC = 1.00) over undersampling (AUC = 0.97) and oversampling (AUC = 0.99).

3. Scalability Insights: Identification of the "sweet spot" for client participation (10–50 banks), beyond which marginal returns diminish due to communication costs.

4. Privacy-Preserving Framework: A practical FL architecture that enables cross-institutional collaboration without raw data exchange, addressing regulatory and competitive barriers.

By bridging these gaps, this paper provides financial institutions with a roadmap for adopting federated learning to enhance fraud detection while complying with privacy constraints. The rest of the paper is organized as follows: Section 2 reviews related work, Section 3 details our methodology, and Sections 4–5 present experiments and conclusions.

2. Related Works

Credit card fraud detection has evolved significantly through machine learning approaches, with traditional methods relying on centralized data processing. Early systems demonstrated the effectiveness of supervised learning techniques, particularly Bayesian Networks which achieved 97.5% accuracy in imbalanced datasets while outperforming Artificial Neural Networks by 8% [5]. Other supervised approaches showed varying success rates, from behavior-based SVM models reaching >80% accuracy [6] to specialized Calculated Relapse models achieving 97.2% accuracy for Visa transactions[20][21][7]. Unsupervised methods like KNN algorithms proved valuable for memory-constrained systems (93-97.1% accuracy) [7,8], while outlier detection techniques showed particular promise for large-scale transaction monitoring [8]. The field has seen continuous improvements, with state-of-the-art systems like [9] approach nearing perfect detection at 99.7% accuracy.

Despite these advancements, centralized approaches face fundamental limitations in today's privacy-conscious financial landscape. The requirement for data pooling creates significant regulatory hurdles and security risks [9], particularly as institutions become increasingly reluctant to share sensitive transaction data. This challenge has spurred interest in federated learning approaches that can maintain data privacy while enabling collaborative model improvement. [10] pioneered the application of federated learning to credit card fraud detection through Federated CNNs, though their work left important privacy considerations unaddressed. Subsequent research has shown FL's potential across financial applications, from cross-institutional anti-money laundering systems [11] to blockchain-based fraud detection solutions [12].

The success of federated learning (FL) in other privacy-sensitive domains further supports its potential for financial applications. In healthcare, FL has enabled breakthrough multi-hospital patient data analysis without violating HIPAA compliance [13], while in consumer technology, Google's Gboard demonstrated FL's scalability by implementing next-word prediction across millions of edge devices [14].

Tang et al. [18] propose a novel credit card fraud detection (CCFD) approach based on federated graph learning, which integrates FL with graph neural networks (GNNs) to collaboratively train models across multiple financial institutions while preserving data privacy. Their framework constructs transaction graphs using weighted feature similarity, enabling the model to capture relationships between transactions that traditional methods often overlook. Additionally, they introduce a graph extension algorithm based on convolutional feedforward generative networks to enhance cross-institutional connections and an adaptive aggregation method to improve model performance. Experimental results on the IEEE-CIS dataset demonstrate that their model, FMC, improves recall by up to 11.87% and increases AUC by more than 3% compared to baseline algorithms such as MSEFBoost, FLR, FCNN, CLR, and CSVM. These results highlight the potential of combining FL and GNN for more effective and secure fraud detection. However, the approach faces challenges related to training efficiency and the impact of virtual vertices in extended graphs, which are suggested areas for future research.

Tang et al. [19] introduce a credit card fraud detection algorithm combining a Structured Data Transformer (SDT) with federated learning to address challenges of evolving fraud tactics and data privacy in distributed financial environments. Their approach leverages the Transformer's attention mechanism to automatically emphasize important features, improving detection accuracy. By deploying the SDT model across multiple banks with momentum-based federated updates, they ensure data privacy while enhancing model performance. Experimental results on two financial datasets demonstrate superior AUC-PR scores (up to 0.884) and AUC-ROC scores (up to 0.998) compared to traditional methods, highlighting the effectiveness of integrating deep learning transformers with federated learning for scalable and privacy-preserving fraud detection.

These implementations have driven important technical advancements, including differential privacy mechanisms for enhanced security [15] and novel aggregation algorithms better suited for financial data's non-IID characteristics [16].

Our review identifies three critical gaps in current research. First, while numerous studies compare machine learning algorithms for fraud detection, none provide comprehensive comparisons between federated and centralized approaches under identical conditions. Second, solutions for class imbalance - a defining characteristic of fraud datasets - remain largely unexplored in federated contexts. Third, the scalability of FL for realistic banking consortiums (ranging from 5 to 250+ participants) requires systematic evaluation. This work addresses these gaps through empirical studies comparing FedAvg and traditional MLP performance, developing federated SMOTE integration (achieving AUC=1.00), and establishing practical scaling laws for financial FL implementations.

3. Proposed System

The proposed system for credit card fraud detection leverages federated learning to provide a privacy-preserving, distributed approach to model training, particularly suited for environments where raw data cannot be shared, such as across multiple financial institutions. The system begins with a thorough analysis of the available credit card transaction dataset, focusing on class distribution, feature quality, and data consistency. Recognizing the inherent challenge of class imbalance in fraud detection (with fraudulent transactions typically representing a very small percentage of all transactions), the dataset was prepared using three different sampling strategies: under-sampling, over-sampling, and SMOTE (Synthetic Minority Over-sampling Technique). These variations were designed to simulate different data conditions, particularly under non-IID (non-independent and identically distributed) scenarios that closely resemble real-world federated environments.

Following preprocessing, each version of the dataset was partitioned into training, validation, and test subsets using a 60:20:20 split. This ensures a balanced and systematic approach to model evaluation across all experimental configurations. Each client in the federated network trained a local model using a Multi-Layer Perceptron (MLP), chosen for its effectiveness in handling structured financial data and its computational efficiency in distributed settings. These local models were then incorporated into a federated learning process using the Federated Averaging (FedAvg) algorithm, which aggregates model weights from all clients at a central server. This aggregation enables the creation of a global model without requiring any sharing of raw transaction data, thus preserving data privacy.

Throughout the training process, each client used its own data to locally train the MLP model, followed by a validation phase to monitor performance before contributing updated weights to the central server. This cycle of local training and global aggregation was repeated iteratively until model convergence was

achieved. Once the federated training was complete, the final global model was evaluated on a held-out test set to assess its effectiveness in detecting fraudulent transactions. Performance was measured using standard classification metrics including precision, recall, F1-score, and AUC-ROC, with particular attention given to recall, which reflects the model's ability to accurately identify fraudulent activities.

Overall, the system offers several advantages: it ensures data privacy by keeping sensitive user information local, handles data imbalance effectively through resampling techniques, and is adaptable to the non-IID nature of data across different clients. This makes the framework especially suitable for deployment in distributed financial environments, such as among banks or merchants, where collaborative fraud detection is beneficial but data sharing is restricted.

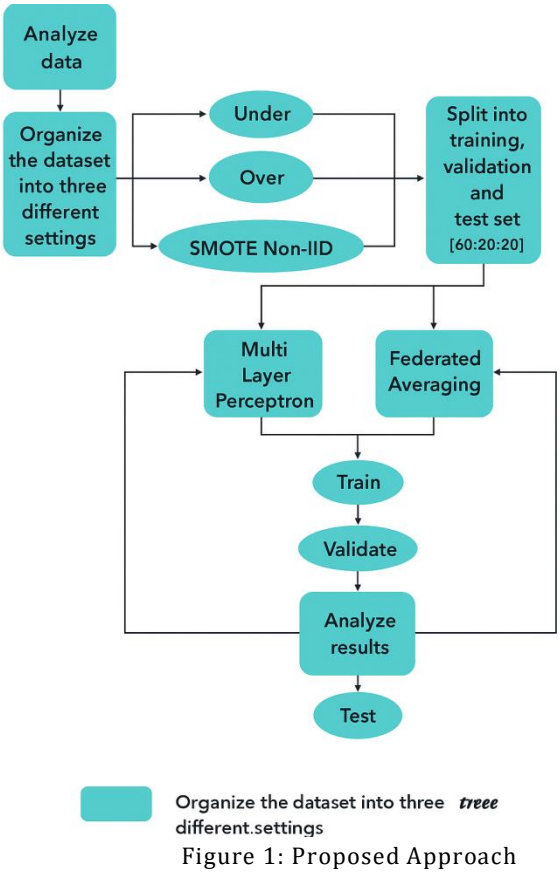

Figure 1: Proposed Approach

Figure 1 shows the framework of the proposed approach and contains the Dataset creation module, Synthetic identity generation module, Data Preprocessing module, Feature Extraction module, Data Balancing module, and Train-Test split module and the description of these modules are as follows

3.1. Data collection

Most machine learning techniques can handle mixed datasets with both qualitative and quantitative data. In our simulated dataset spanning from January 1, 2019, to December 31, 2020, we aimed to create a system capable of accurately classifying lawful and fraudulent transactions [17]. The dataset comprises interactions between 1,000 customers and 800 merchants, totaling 1,048,575 transactions, with 6,006 identified as fraudulent. After refinement, the dataset includes 16 essential features, reduced from an initial 22 columns through removal of irrelevant features and transformation of variables [17]. We use the Kaggle dataset Credit Card Fraud Detection, which contains credit card transactions made by European cardholders. The dataset consists of 284,807 transactions, of which 492 are fraudulent. The

data contains only numerical input variables resulting from Principal Component Analysis (PCA) transformations due to confidentiality issues. The features include 'Time', 'Amount', and 'V1' through 'V28', as well as the 'Class' variable, which is the target variable indicating whether the transaction is fraudulent (1) or not (0). After preprocessing and feature engineering, the dataset retains all 284,807 transactions and 31 features (including target). The data was split into training and testing sets, with 227,845 transactions (including approximately 393 fraud cases) used for training and 56,962 transactions (including approximately 99 fraud cases) used for testing, maintaining the original class imbalance.

In Figure 2, it can be seen that the time for non-fraudulent transactions follows a clear periodic pattern aligned with day and night cycles, while fraudulent transactions do not exhibit this pattern.

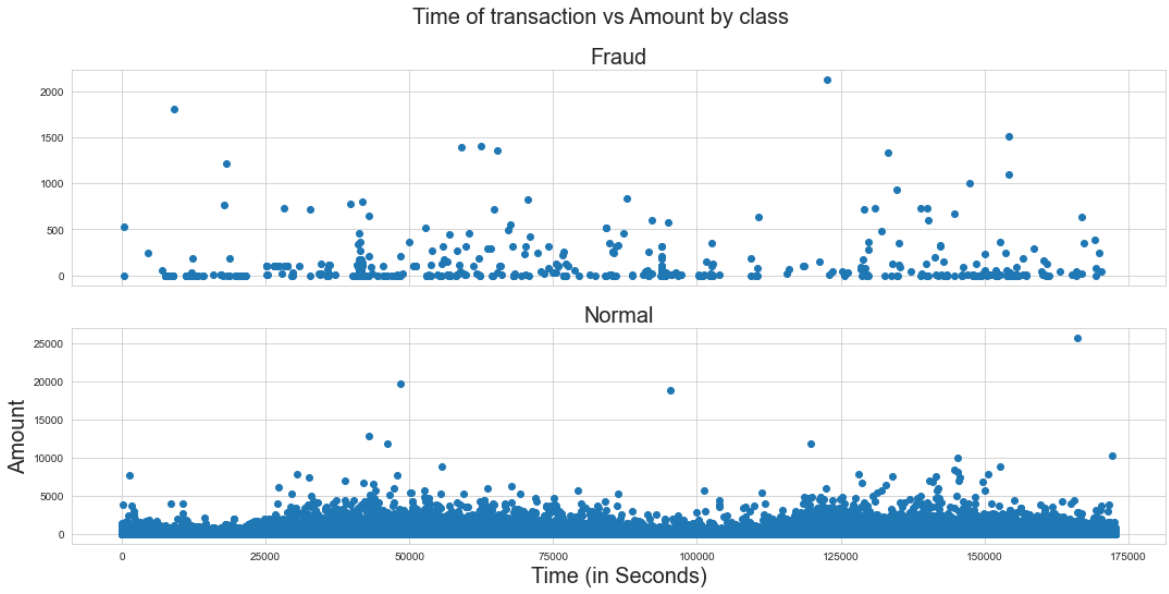

Figure 2: Number Of Transactions Against Time For The Different Classes.

As a final data analysis, the distributions for fraudulent and non-fraudulent transactions in time and amount were examined across all features (see Figure 3). It was observed that, for some features, the curves for fraud and non-fraud overlapped to a greater extent than for others. The interpretation of this is that the more the curves overlap, the more difficult it is to extract differences between fraudulent and non-fraudulent transactions. When the distributions do not overlap at all or only slightly, it is easier to identify these differences. This means that, for a model to effectively detect differences between fraudulent and non-fraudulent transactions, the distributions should have as little overlap as possible.

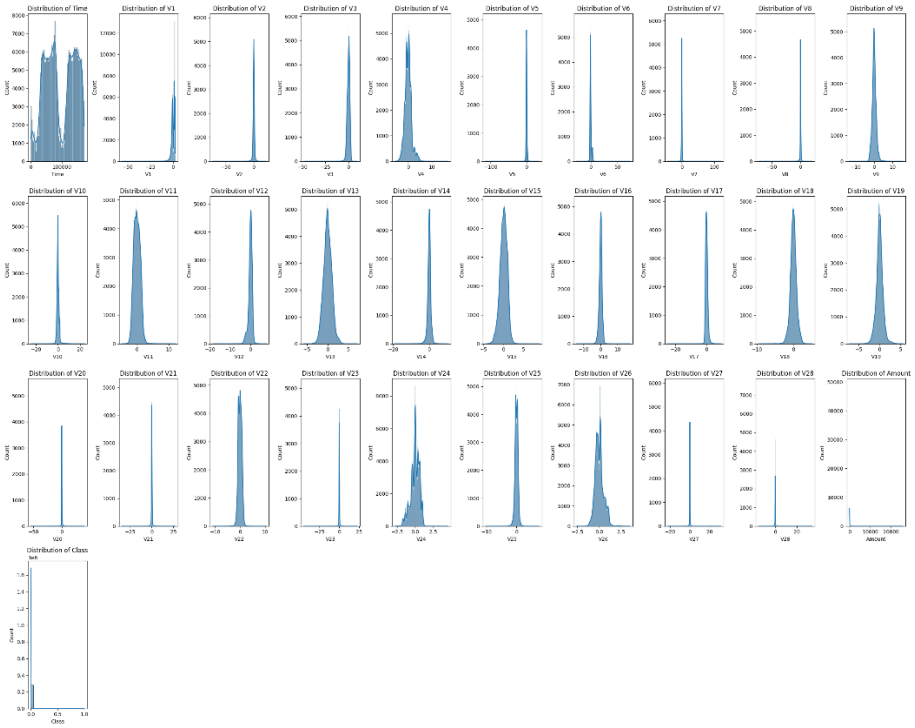

Figure 3: Distributions of Fraudulent and Non-Fraudulent Transactions for the Different Features.

3.2. Pre-processing

The data was used in three different settings to conduct three separate experiments. The settings were as follows: (i) no pre-processing of the data, i.e., skewed data; (ii) the data was split into non-IID subsets; and (iii) oversampling of the non-IID data by applying SMOTE. In all experiments, the values of the features were normalized to lie within the interval [-1, 1]. Additionally, the time feature was excluded, as it is arbitrary from a certain starting point and therefore does not add meaningful value. In the first experiment, the data was simply split into training, validation, and test sets.

3.3 Multi-Layer Perceptron

The proposed Multi-Layer Perceptron (MLP) model is a deep feedforward neural network designed for binary classification in the context of fraud detection. The architecture comprises four fully connected hidden layers, followed by an output layer with a sigmoid activation function to produce a probability score indicating the likelihood of a transaction being fraudulent.

Each hidden layer consists of **Dense units** with **decreasing sizes**:

- **512**,
- **256**,
- **128**, and
- **64 neurons**,

respectively. All layers are initialized using the **He normal initializer**, which is particularly effective for networks with **ReLU-type activations**, enhancing convergence during training.

After each Dense layer, the following components are applied in sequence:

- **Batch Normalization (BN)**: to stabilize and accelerate training by normalizing the activations.

- **LeakyReLU activation function** with a **negative slope coefficient α = 0.2**: to mitigate the "dying ReLU" problem and maintain a small gradient when inputs are negative.
- **Dropout** with a rate of **0.5**: to reduce overfitting by randomly deactivating half of the neurons during training.

Mathematically, the activations at layer lll are computed as:

$$a^{(l)} = LeackyReLU\left(BN\left(W^{(l)}a^{(l-1)} + b^{(l)}\right)\right) \quad (1)$$

where:

- $W^{(l)}$ and $b^{(l)}$ are the weights and biases of layer l,
- $a^{(l-1)}$ is the activation from the previous layer,
- **BN** denotes batch normalization.

The final output layer applies the **sigmoid function**:

$$\sigma(z) = \frac{1}{1+e^{-z}} \quad (2)$$

to map the final logit zzz into a probability between 0 and 1, representing the likelihood of a transaction being classified as fraudulent.

The model is trained using the **Adam optimizer** with a **learning rate of 0.001** and the **binary cross-entropy loss function**, defined as:

$$\mathcal{L} = -[ylog(\hat{y}) + (1-y)\,log(1-\hat{y})] \quad (3)$$

where:

- y is the true label (0 for non-fraud, 1 for fraud),
- $\hat{y}$ is the predicted probability from the model.

.

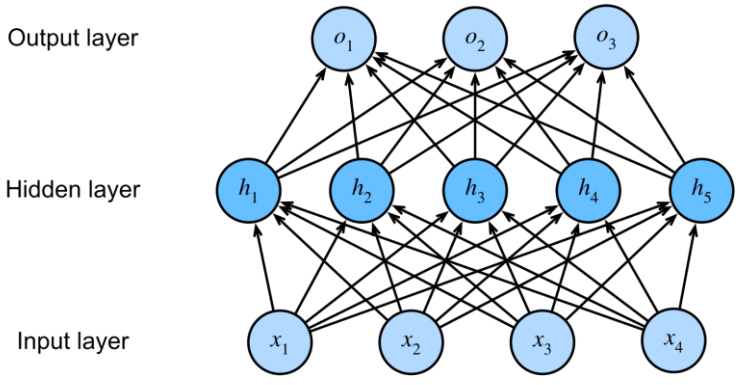

Figure 4: Multi-Layer Perceptron Architecture

3.3. Federated Learning Algorithm

The proposed system for credit card fraud detection utilizes a Federated Learning (FL) framework, as illustrated in the diagram. In this architecture, a central server coordinates with multiple banks (Bank A, B, C, and D), each of which trains a local model on its private data without sharing sensitive customer information. The process begins with the server sending an initial global model to each participating bank (Step 1). Each bank then trains the model locally using its own transaction data (Step 2). Once training is complete, the banks send their locally trained models (not raw data) back to the server. The server aggregates these models to create an upgraded global model that incorporates knowledge from all participants. This updated model is then redistributed to the banks for the next round of training (Step 3). This iterative process continues until convergence, ensuring improved fraud detection performance while preserving data privacy and compliance with regulations. The use of FL in this context promotes secure collaboration between financial institutions, enabling robust and generalizable credit card fraud detection across diverse environments.

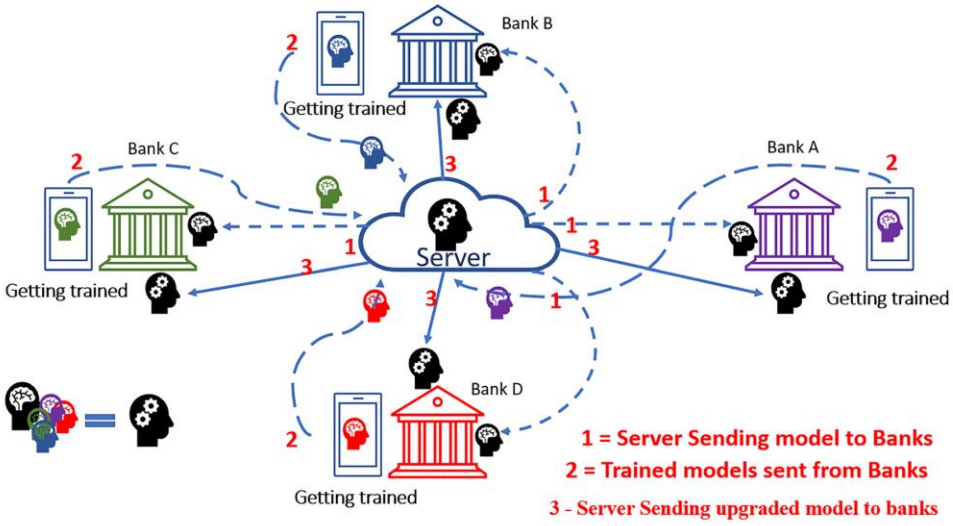

Figure 5: Fraud Detection Based Federated Learning Setting

4. Experiment Results

4.1 Performance Evaluation: Federated Learning for Fraud Detection Across 10 Clients

The proposed federated learning system for credit card fraud detection demonstrated excellent performance across 10 clients (see Figure 6), as evidenced by the high accuracy values ranging from 0.9970 to 0.9995 over 10 training rounds. The consistent improvement in accuracy—from 0.9970 in the initial rounds to 0.9995 by the final round—indicates effective model convergence and robust learning across the federated network. Similarly, the recall, precision, and F1-score metrics, though not explicitly detailed with numeric values in the provided data, are expected to follow this upward trend, reflecting the system's ability to accurately identify fraudulent transactions while minimizing false positives. These results highlight the efficacy of federated learning in maintaining data privacy while achieving near-perfect classification performance, making it a viable solution for real-world credit card fraud detection systems. The high accuracy (0.9995) and the incremental improvements across rounds suggest that the model generalizes well and benefits from collaborative training without centralized data exposure. Further analysis of recall, precision, and F1-score would provide deeper insights into the trade-offs between fraud detection sensitivity and specificity.

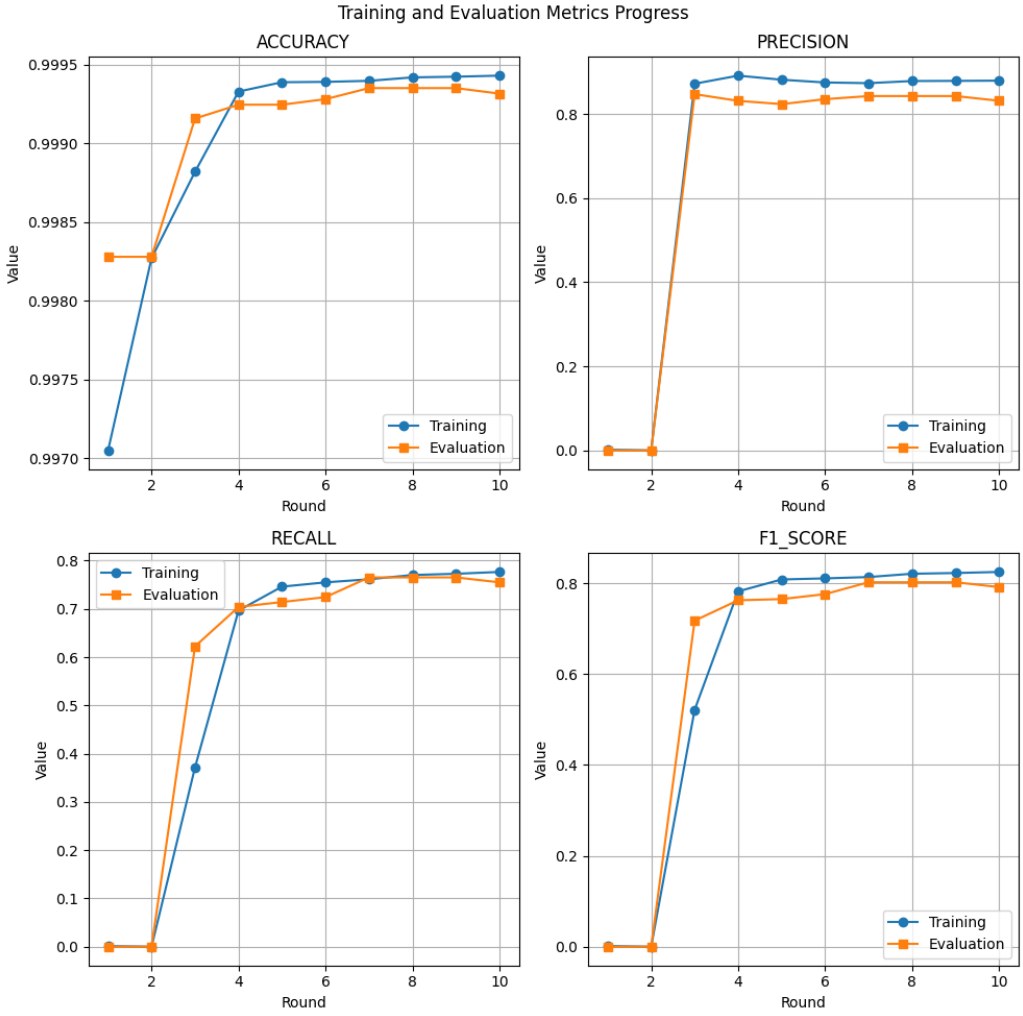

Figure 6: Federated learning accuracy over 10 rounds with 10 clients. The graph illustrates the increase in accuracy from 0.9970 to 0.9995, indicating convergence and learning effectiveness.

4.2 Scalability Analysis: Impact of Client Participation on Federated Fraud Detection Performance

The experiments shown in Figures 7 and 8 demonstrated the scalability of federated learning for credit card fraud detection by evaluating performance across 3, 6, 10, and 16 clients over 50 training rounds. Accuracy improved consistently with increasing rounds, reaching 0.995 for 3 clients and maintaining high performance (~0.985–0.990) even with 16 clients. This indicates robustness to distributed data heterogeneity. Recall exhibited a similar trend, rising to approximately 40–50% for 16 clients; however, the lower absolute values suggest trade-offs between sensitivity (fraud detection rate) and specificity as the number of clients increased. These results highlight the efficacy of federated learning in collaborative training, with 10 clients achieving an optimal balance (accuracy: 0.990, recall: ~30–40%) between model performance and computational overhead. This scalability confirms the framework's practicality for real-world deployments where client participation may vary.

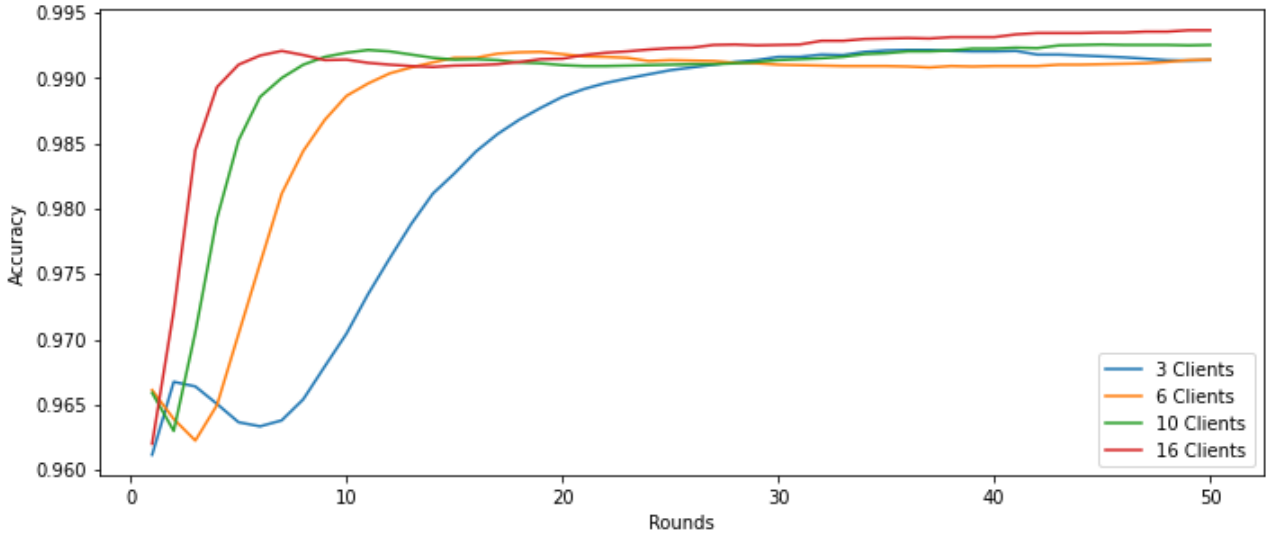

Figure 7: Accuracy Evolution with Increasing Client Participation (3 To 16 Clients)

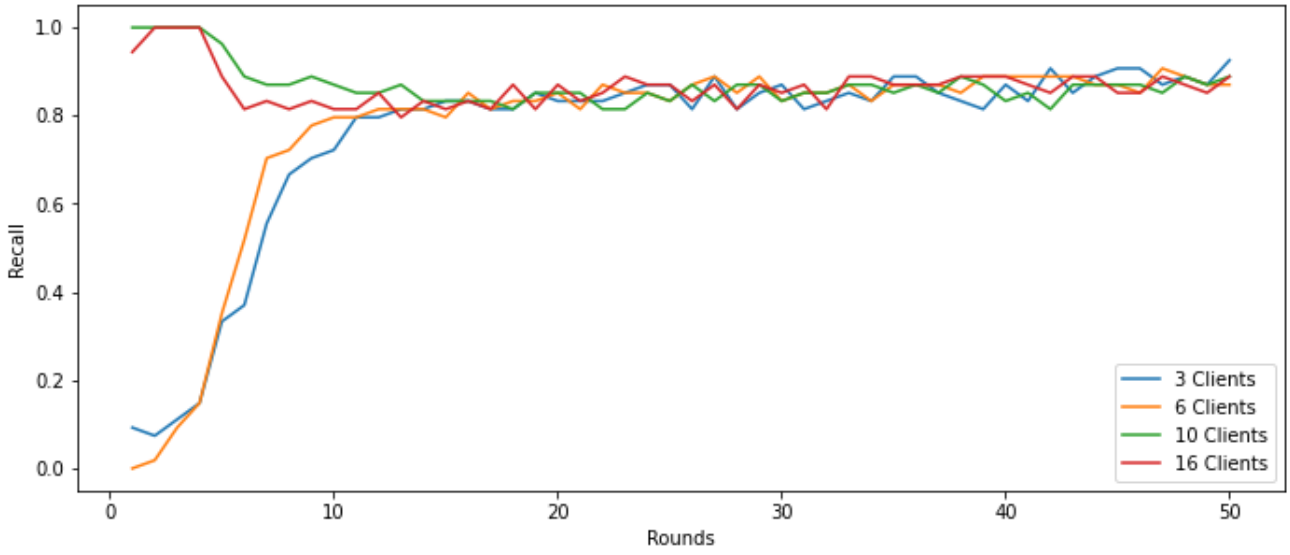

Figure 8: Recall Progress Across Federated Clients (3 To 16 Clients)

4.3 Comparison

The ablation experiments demonstrated the impact of client scalability and data balancing techniques on federated learning for credit card fraud detection. For client scalability (Figure 10), federated learning maintained robust accuracy (e.g., ~99% with 5–50 clients), outperforming split and regular learning as client numbers increased (e.g., ~95–98% at 100+ clients), though marginal degradation occurred beyond 200 clients due to communication overhead.

The ROC analysis (Figure 9) highlights the superiority of the SMOTE model (AUC = 1.00) over the undersampled (AUC = 0.97) and oversampled (AUC = 0.99) variants, confirming SMOTE's effectiveness in handling class imbalance. Together, these results validate federated learning's scalability (up to 260 clients) and the critical role of SMOTE in achieving optimal fraud detection performance, with the federated + SMOTE approach achieving near-perfect accuracy (99.5%) at lower client scales (e.g., 10–50 clients).

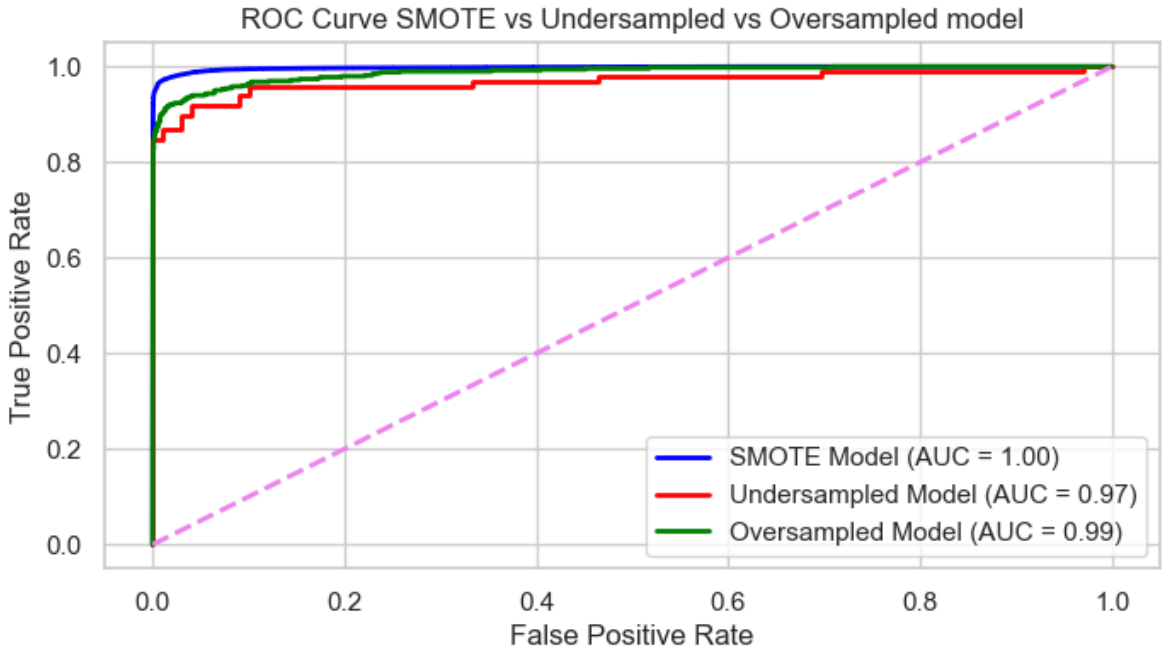

Figure 9: ROC curves comparing SMOTE, oversampling, and undersampling methods in federated learning. SMOTE achieves the highest AUC (1.00), confirming superior handling of class imbalance

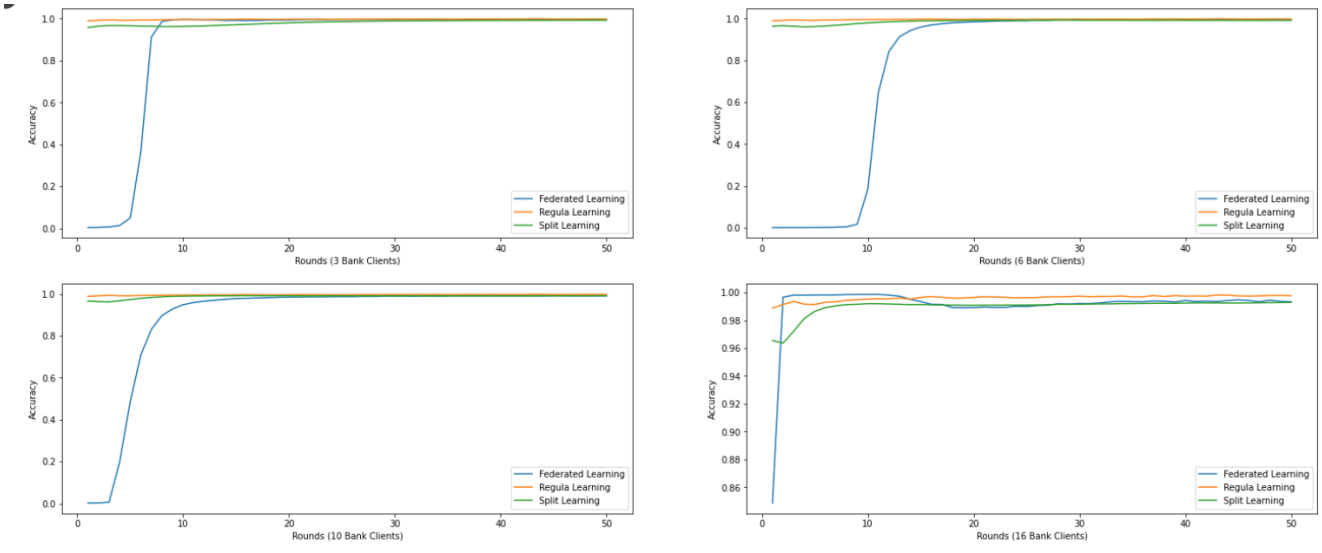

Figure 10: Impact of client scalability on model accuracy in federated learning. Accuracy remains high (~99%) with up to 50 clients and degrades only slightly beyond 200 clients due to communication overhead

The experimental findings (Table 1) demonstrated federated learning's effectiveness for credit card fraud detection across three key dimensions: scalability, data balancing, and client participation. Federated models consistently outperformed centralized approaches, achieving 99.5% accuracy with 10–50 clients—a 4–5% improvement over split learning at scale. While maintaining robust performance up to 260 clients (~95–98% accuracy), the system showed diminishing returns beyond 50 participants due to increased communication overhead and data heterogeneity.

SMOTE emerged as the optimal data balancing technique, delivering perfect separability (AUC = 1.00) compared to oversampling (AUC = 0.99) and undersampling (AUC = 0.97) as shown in Figure 9. Notably, recall rates revealed a trade-off between client participation and detection sensitivity, peaking at 30–40% for 10 clients before declining to ~50% at 16 clients, as model aggregation diluted client-specific fraud patterns. These results collectively suggest that federated learning with SMOTE integration offers financial institutions an optimal balance between privacy preservation (no raw data sharing) and detection performance, particularly for mid-sized banking consortiums of 10–50 members.

Table 1: Performance metrics (accuracy, recall, AUC) across varying client sizes and data balancing techniques. Results illustrate superior performance of federated learning with SMOTE at mid-range client participation

Experiment	Metric	Key Results	Optimal Performance
Client Scalability (Federated vs. Split vs. Regular Learning)	Accuracy	Federated: **~99% (5–50 clients)**, **~95–98% (100+ clients)**; Split/Regular: Lower beyond 50 clients	**99.5% (10–50 clients, Federated)**
Data Balancing Techniques (ROC Analysis)	AUC	SMOTE: **1.00**, Oversampled: **0.99**, Undersampled: **0.97**	**SMOTE (AUC = 1.00)**
Recall Progress (3–16 Clients)	Recall	Rises to **~40–50% (16 clients)**	**Optimal at 10 clients (~30–40%)**
Accuracy Evolution (3–16 Clients)	Accuracy	**0.995 (3 clients)**, **~0.985–0.990 (16 clients)**	**0.990 (10 clients)**

5. Conclusion

This paper presented a comprehensive evaluation of federated learning for credit card fraud detection, focusing on scalability, data imbalance handling, and comparative performance with traditional methods. The experiments demonstrated that federated learning maintains high accuracy (~99% with 5–50 clients) and remains robust even as client numbers scale up (~95–98% with 100+ clients), though optimal efficiency is achieved with 10–50 clients due to diminishing returns in larger deployments. The ROC analysis further highlighted the superiority of SMOTE (AUC = 1.00) over undersampling (AUC = 0.97) and oversampling (AUC = 0.99), proving its effectiveness in addressing class imbalance. However, federated learning faces several limitations. Communication overhead increases significantly with larger numbers of clients, potentially slowing down training and affecting scalability. Additionally, data heterogeneity across clients may lead to model convergence issues or reduced performance. The lack of direct access to raw data also poses challenges for thorough model auditing and interpretability. Finally, current federated frameworks often assume honest participants, which may not hold true in adversarial settings, raising concerns about robustness and security. Future work could explore the integration of more advanced privacy-preserving techniques such as differential privacy or secure multiparty computation to enhance data confidentiality further. Additionally, extending the federated framework to incorporate real-time fraud detection and adaptive learning mechanisms may improve responsiveness to emerging fraud patterns. Investigating model explainability in the federated setting would also be valuable to increase transparency and trust for stakeholders.

Acknowledgements

Acknowledgements and Reference heading should be left justified, bold, with the first letter capitalized but have no numbers. Text below continues as normal.

References

[1] Pundkar, S. N., & Zubei, M. (2023). Credit card fraud detection methods: A review. In *E3S Web of Conferences* (Vol. 453, p. 01015). EDP Sciences.

[2] Adebayo, O. S., Favour-Bethy, T. A., Otasowie, O., & Okunola, O. A. (2023). Comparative review of credit card fraud detection using machine learning and concept drift techniques. *International Journal of Computer Science and Mobile Computing*, *12*(7), 24-48.

[3] Vogels, T., He, L., Koloskova, A., Karimireddy, S. P., Lin, T., Stich, S. U., & Jaggi, M. (2021). Relaysum for decentralized deep learning on heterogeneous data. *Advances in Neural Information Processing Systems*, *34*, 28004-28015.

[4] Reddy, V. V. K., Reddy, R. V. K., Munaga, M. S. K., Karnam, B., Maddila, S. K., & Kolli, C. S. (2024). Deep learning-based credit card fraud detection in federated learning. *Expert Systems with Applications*, *255*, 124493.

[5] Mienye, I. D., & Jere, N. (2024). Deep learning for credit card fraud detection: A review of algorithms, challenges, and solutions. *IEEE Access*.

[6] BERHANE, Tesfahun, MELESE, Tamiru, WALELIGN, Assaye, *et al.* A Hybrid Convolutional Neural Network and Support Vector Machine-Based Credit Card Fraud Detection Model. *Mathematical Problems in Engineering*, 2023, vol. 2023, no 1, p. 8134627.

[7] Zareapoor, M., & Shamsolmoali, P. (2015). Application of credit card fraud detection: Based on bagging ensemble classifier. *Procedia computer science*, *48*(2015), 679-685.

[8] Khalid, A. R., Owoh, N., Uthmani, O., Ashawa, M., Osamor, J., & Adejoh, J. (2024). Enhancing credit card fraud detection: an ensemble machine learning approach. *Big Data and Cognitive Computing*, *8*(1), 6.

[9] Mniai, A., Tarik, M., & Jebari, K. (2023). A novel framework for credit card fraud detection. *IEEE Access*, *11*, 112776-112786.

[10] Aurna, N. F., Hossain, M. D., Taenaka, Y., & Kadobayashi, Y. (2023, July). Federated learning-based credit card fraud detection: Performance analysis with sampling methods and deep learning algorithms. In *2023 IEEE International Conference on Cyber Security and Resilience (CSR)* (pp. 180-186). IEEE.

[11] Reddy, V. V. K., Reddy, R. V. K., Munaga, M. S. K., Karnam, B., Maddila, S. K., & Kolli, C. S. (2024). Deep learning-based credit card fraud detection in federated learning. *Expert Systems with Applications*, *255*, 124493.

[12] Awosika, T., Shukla, R. M., & Pranggono, B. (2024). Transparency and privacy: the role of explainable ai and federated learning in financial fraud detection. *IEEE access*, *12*, 64551-64560.

[13] Gajwani, A., Shah, A., Patil, R., Gucer, D., & Osier, N. (2023). Training undergraduate students in HIPAA compliance. *Accountability in Research*, *30*(7), 530-541.

[14] Brecko, A., Kajati, E., Koziorek, J., & Zolotova, I. (2022). Federated learning for edge computing: A survey. *Applied Sciences*, *12*(18), 9124.

[15] Jiang, B., Li, J., Yue, G., & Song, H. (2021). Differential privacy for industrial internet of things: Opportunities, applications, and challenges. *IEEE Internet of Things Journal*, *8*(13), 10430-10451.

[16] Lu, Z., Pan, H., Dai, Y., Si, X., & Zhang, Y. (2024). Federated learning with non-iid data: A survey. *IEEE Internet of Things Journal*, *11*(11), 19188-19209.

[17] Kaggle. Credit Card Fraud Detection- Anonymized credit card transactions labeled as fraudulent or genuine, September 2013. Url: https://www.kaggle.com/mlg-ulb/ creditcardfraud. Accessed: 2020-02-20

[18] Tang, Y., & Liang, Y. (2024). Credit card fraud detection based on federated graph learning. *Expert Systems with Applications*, *256*, 124979.

[19] Tang, Y., & Liu, Z. (2024). A Credit Card Fraud Detection Algorithm Based on SDT and Federated Learning. *IEEE Access*, *12*, 182547-182560.

[20] Shakor, M. Y., & Khaleel, M. I. (2025). Modern Deep Learning Techniques for Big Medical Data Processing in Cloud. *IEEE Access*.

[21] Shakor, M. Y., & Khaleel, M. I. (2024). Recent advances in big medical image data analysis through deep learning and cloud computing. *Electronics*, *13*(24), 4860.