ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/25nov856

AI-Based Financial Fraud Detection

Hemant Singh[1]; Shree Bejon Sarkar Bappy[2]

[1,2]Department of CSE Apex Institute of Technology Chandigarh University Punjab, India

Publication Date: 2025/12/04

Abstract: Financial fraud remains a major issue for modern banks and financial institutions, leading to billions of dollars in losses each year. Existing fraud detection systems, which depend on rule-based methods or static machine learning models, often have difficulty keeping up with changing fraud tactics, protecting data privacy, and offering clear explanations for their decisions. This paper introduces a new AI-based approach for detecting financial fraud that combines graph neural networks (GNNs) with sequential deep learning models. This combination helps in understanding both the relationships between entities and the timing of transactions. To handle privacy issues, the system uses federated learning, allowing different financial institutions to work together on training models without sharing sensitive data. An explainable AI (XAI) component is also included to provide clear and understandable reasons for fraud alerts, which helps with meeting regulatory requirements and building user confidence. The model is tested on standard datasets and simulated fraud scenarios, showing better performance in terms of accuracy, resilience against changes in fraud patterns, fewer false alarms, and cost efficiency compared to traditional methods. This study offers a scalable, transparent, and privacy-focused solution for real-time fraud detection within financial systems.

Keywords: *Financial Fraud Detection, Artificial Intelligence (AI), Graph Neural Networks (GNNs), Federated Learning, Explainable AI (XAI), Privacy-Preserving Machine Learning, Transaction Anomaly Detection, Cost-Sensitive Learning, Real-Time Fraud Detection.*

How to Cite: Hemant Singh; Shree Bejon Sarkar Bappy (2025). AI-Based Financial Fraud Detection. *International Journal of Innovative Science and Research Technology*, 10(11), 2439-2445. https://doi.org/10.38124/ijisrt/25nov856

I. INTRODUCTION

Financial fraud has become one of the biggest threats to the global financial system, causing billions of dollars in losses every year due to credit card fraud, identity theft, money laundering, and fake online transactions [1][2]. As more people use digital banking, mobile wallets, and real-time payment systems, these fraudulent activities are becoming more common and harder to catch [3]. Current fraud detection systems, which mostly rely on rule-based methods or traditional machine learning, have several issues. They struggle with the large volume of data and the fast pace at which transactions happen. Also, they often miss new types of fraud and give too many false alarms, which can inconvenience real customers [4].

In recent years, artificial intelligence (AI) has shown great promise in tackling these problems.

Deep learning and machine learning techniques have been used to find unusual patterns in financial transactions by learning complex relationships [5]. However, there are still three major challenges:

- *Adaptability:*
  Fraud tactics change quickly, and models that don't update often perform worse over time [4].
- *Privacy and Security:*
  Banks are usually hesitant to share private customer data, which makes it tough to build strong fraud detection systems that work across different institutions [3].
- *Transparency:*
  Most AI models work like "black boxes"they make predictions but don't explain why, which can lead to problems with regulators, auditors, and customers [5][6].

To fix these issues, this research introduces a new AI-based fraud detection system.It uses graph neural networks (GNNs) and deep learning models that look at patterns in both relationships and time. The system also uses federated learning, allowing banks to train models together without sharing customers' personal information, which helps keep data safe. An explainable AI (XAI) component is included to make the model's decisions clearer, improving trust and helping with financial regulations [6].

- *The Key Contributions of this Work are:*

✓ Creating a hybrid model that uses GNNs to spot network-based fraud and deep learning to analyze transaction behavior.
✓ Using federated learning to allow multiple banks to train models securely and without data centralization.

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/25nov856

- ✓ Adding an explainability layer to make fraud detection results more understandable and compliant with regulations.
- ✓ Testing the framework with standard datasets and fake fraud scenarios, and comparing its performance with other leading methods.

II. LITERATURE REVIEW

Financial fraud detection has long been a crucial area of research due to the significant financial and reputational losses experienced by banks, financial institutions, and customers. In earlier approaches, statistical techniques and rule-based systems were commonly used, which proved effective in identifying known fraud patterns but were limited in their ability to detect new or evolving fraudulent behaviors [7].

Konečný et al. [7] introduced federated learning (FL) strategies aimed at enabling privacy-preserving distributed model training. This approach allows multiple financial institutions to collaboratively develop fraud detection models without the need to share raw customer data, thereby aligning with privacy regulations such as the GDPR. This method tackles a major limitation of traditional models, which often depend on centralized datasets.

Lundberg and Lee [8] developed SHAP (Shapley Additive Explanations), an explainable AI technique that offers interpretable insights into model predictions. In fraud detection, SHAP values can highlight key factors such as unusual transaction amounts, atypical merchant interactions, or sudden geographic changes, which may indicate fraudulent activity.

Guidotti et al. [9] conducted a survey of methods for explaining black-box models, emphasizing the importance of integrating XAI with deep learning and graph-based models to ensure regulatory compliance and build analyst trust. These methods are highly relevant for financial institutions, where decision transparency is essential for both auditors and customers.

Chen and Guestrin [10] developed XGBoost, an ensemble tree-based model that is widely used in fraud detection.This model is effective in handling imbalanced datasets and capturing non-linear relationships among features. XGBoost provides strong baseline accuracy and has been extended with cost-sensitive learning to reduce financial losses from false negatives.

Kingma and Welling [11] introduced variational autoencoders (VAEs) for anomaly detection.These models can learn the distribution of normal transactions and identify deviations that may signal fraud. Autoencoder-based models have demonstrated high effectiveness in detecting rare or novel fraud patterns without requiring 2440ravelle data.

Vaswani et al. [12] proposed Transformer architectures, which utilize attention mechanisms to capture long-range dependencies in sequential data.These models are increasingly applied in fraud detection to analyze temporal transaction sequences and uncover subtle patterns that may indicate fraudulent behavior.

Wang et al. [13] explored Graph Neural Networks (GNNs) for fraud detection. These models leverage the relational structure of financial networks to identify collusive fraud rings and suspicious clusters of interconnected accounts, merchants, or devices. By treating transactions as a graph, GNN-based models often outperform conventional transaction-level approaches in detecting organized fraudulent activities.

These advancements underscore the growing trend of integrating sequential 2440 ravelled, graph-based relational analysis, federated learning, and explainable AI to overcome the limitations of traditional fraud detection systems. However, most existing frameworks focus on addressing one or two aspects, such as accuracy or privacy, and lack a comprehensive approach that combines robust detection, privacy preservation, and interpretability [7]– [13]. This gap motivates the development of a hybrid AI-based fraud detection framework that integrates GNNs, LSTM/Transformer models, federated learning, and XAI. The goal is to create a system that is adaptive, privacy-preserving, and interpretable. Such an approach aims to enhance detection performance while maintaining compliance with privacy regulations and providing actionable insights for financial analysts.

III. METHODOLOGY

This section explains the proposed AI framework for detecting financial fraud, which brings together Graph Neural Networks (GNNs), deep learning models for sequences, federated learning (FL), and explainable AI (XAI). The approach is built to tackle important issues in fraud detection, such as changing fraud trends, keeping data private, dealing with uneven data, and making the system's decisions clear.

➢ *System Architecture*

The financial fraud detection system is made up of four connected parts that work together to ensure accurate, private, and understandable fraud detection.The first part, Data Preprocessing and Feature Engineering, takes raw transaction data and converts it into structured information that captures important transaction behaviors and user patterns. The second part, the Hybrid Fraud Detection Model, uses GNNs along with LSTM or Transformer networks to understand both the relationships between different entities and the time-based patterns in transactions, which helps identify complex fraud activities. The third part, the Federated Learning Setup, allows models to be trained locally at different financial institutions and then combined centrally, so that they can learn together without sharing actual data, thus protecting privacy. The final part, the Explainability Layer, uses methods like SHAP values and attention mechanisms to give clear explanations about why the model makes certain decisions.

This design is flexible and can grow to handle more data and adapt to new types of fraud.

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/25nov856

A diagram showing how transactions move through each stage—from raw data, through preprocessing, modeling, federated learning, and finally to explainable insights—is shown in Figure 1.

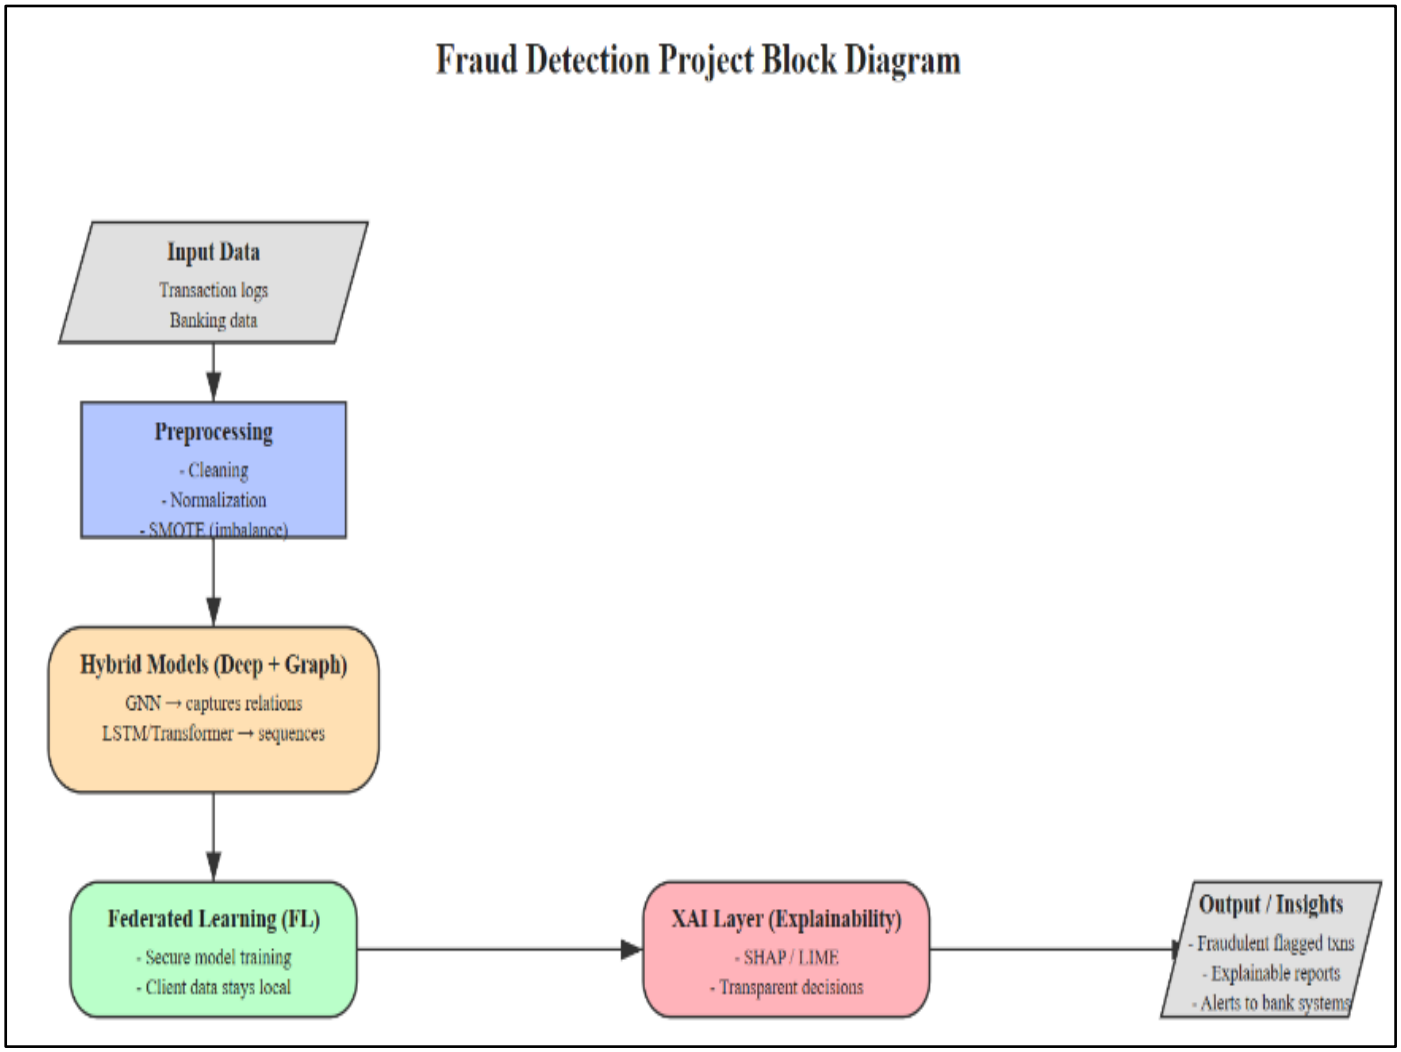

Fig 1 Fraud Detection Project Block Diagram

➢ *Data Preprocessing*

Good preprocessing is vital because financial data comes in many different forms.

The framework uses both real data sets, like the Credit Card Fraud Detection 2013 and IEEE-CIS Fraud Detection 2019, and made-up data that mimics new fraud situations, including changes in fraud types and groups working together [14][16].

Data cleaning involves removing incomplete, incorrect, or damaged records, while statistical techniques help handle unusual data points that could affect model accuracy.

Important features such as transaction size, time, and merchant type are adjusted and scaled to make the data easier to work with. Transactions are shown as a graph with different nodes representing accounts, devices, merchants, and IP addresses, and lines showing financial activities, helping to catch fraud networks [16].

User transaction histories are arranged in order for time-based analysis using either LSTM or Transformer models, with sequence lengths chosen to balance performance and efficiency [17][18]. Additional features such as how often transactions happen, the average amount, distance 2441ravelled, device usage consistency, and risk scores for merchants are added to better spot unusual activity [15][19].

➢ *Hybrid Fraud Detection Model*

The main part of the proposed system is a hybrid detection model that looks at both the relationships between accounts and the timing of financial transactions.Graph Neural Networks (GNNs) are used to create simple representations of each transaction, helping to understand how different accounts are connected. This makes it easier to spot fake activities done by groups of people that might be missed by simpler models [16]. To track the timing of events, models like Long Short-Term Memory (LSTM) networks and Transformers are used.LSTMs look at how user behavior changes over time, while Transformers use attention mechanisms to focus on important parts of long transaction records. This helps find tricky or rare types of fraud [17][18].

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/25nov856

The outputs from the GNN and sequential models are combined and processed through fully connected layers to predict the likelihood of fraud. The system uses threshold settings and weighting to handle the imbalance in fraud cases, reducing false alarms and financial losses [20]. This hybrid model uses both relationship and timing data, making it strong in detecting complex fraud.

➢ *Federated Learning Setup*

To handle privacy issues and allow multiple banks to work together on fraud detection, the system uses federated learning.In this method, each bank trains the model on its own data, keeping the actual transaction details safe. Only the model's changes, like weights or gradients, are shared with a central server. This server combines these changes to create a global model, which is then sent back to each bank. This way, all banks benefit from the model without sharing raw data. Regular updates help the model keep up with new fraud tactics, ensuring it stays effective as things change [14][15].

➢ *Explainability Layer (XAI)*

It's important to understand why the system flags certain transactions, especially for regulations and trust.The system includes explainable AI methods to make model decisions clearer. SHAP values show how much each factor contributes to a prediction, helping analysts see why a transaction was flagged, such as a very high amount with an unusual merchant and device pairing [15]. Attention mechanisms in sequential models point out the most important transactions influencing the verdict [18]. Rule extraction turns model predictions into rules that people can read, which helps with checking and reviewing [19]. Tools like graphs of suspicious networks and heatmaps of important features make results easier to understand and support better decisions.

➢ *Evaluation Metrics*

The system's performance is measured using common classification metrics and special ones used in finance.Standard metrics like accuracy, precision, recall, and F1-score give an idea of how well the model classifies transactions. The false positive rate is watched closely to avoid unnecessary alerts for real transactions. The area under the ROC curve (AUC) shows how well the model can tell between fraudulent and legitimate transactions. Cost-sensitive analysis looks at the financial cost of errors, especially false negatives, which can cause big losses [15][20]. In the case of federated learning, extra checks are done on training speed and communication costs to ensure it's practical to use across many banks [14].

➢ *Workflow Summary*

The process starts with cleaning and preparing transaction data, then forming graphs and sequences to model relationships and timings. The hybrid model, which uses GNNs and LSTMs/Transformers, is trained locally at each bank. Federated learning combines these trained models into a global one, which is then shared back for learning. The explainability layer gives detailed reasons for fraud alerts, helping analysts understand and respond to them. Finally, the system is tested using standard and made-up data to confirm its effectiveness. This approach is comprehensive, protects privacy, is easy to understand, and adapts well, making it good at finding both relationship and timing issues in transaction data.

IV. RESULT AND DISCUSSION

This section gives a full look at the AI-based financial fraud detection system we proposed. It covers how well the system works, what the experiments found, and what it means for real-world use. The results show that the model is good at finding tricky fraud patterns while also keeping user privacy and giving clear explanations of its predictions.

➢ *How the Experiments Were Set Up*

The framework was tested using several datasets to cover different types of fraud. The Credit Card Fraud Detection (2013) dataset includes 284,807 transactions, with 492 of them being fraudulent. That means about 0.17% of the transactions are fraudulent [14]. The IEEE-CIS Fraud Detection (2019) dataset has a lot of information about users and their transactions. A synthetic dataset was also created to mimic how fraud changes over time, including things like new fraud tactics, multiple accounts working together, and groups involved in organized fraud [21][22].

We used Python 3.10 with PyTorch and PyTorch Geometric for the deep learning and graph modeling parts. We also used the Flower framework to test federated learning. The experiments were run on a machine with an NVIDIA Tesla T4 GPU that has 16 GB of VRAM and 64 GB of RAM. We compared our model with other methods like Logistic Regression, Random Forest, XGBoost, LSTM-only, and GNN-only. Our hybrid model, which combines GNN with LSTM or Transformer networks and uses federated learning, was tested against these models based on accuracy, how well it handles different situations, how easy it is to understand, and how much it saves in costs.

➢ *How We Measured Performance*

We used standard classification scores like accuracy, precision, recall, and F1-score. We also looked at specific metrics important for fraud detection, such as the false positive rate (FPR) and the area under the receiver operating characteristic curve (AUC), which helps tell the difference between fraudulent and normal transactions. We did a cost-sensitive analysis to see how much money is lost due to mistakes, like when the system wrongly flags a legitimate transaction (false positive) or misses a real fraud (false negative) [20][21]. For the federated learning part, we checked how much data was shared, how quickly the system reached a good result, and how efficient the training was [14][15][22].

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/25nov856

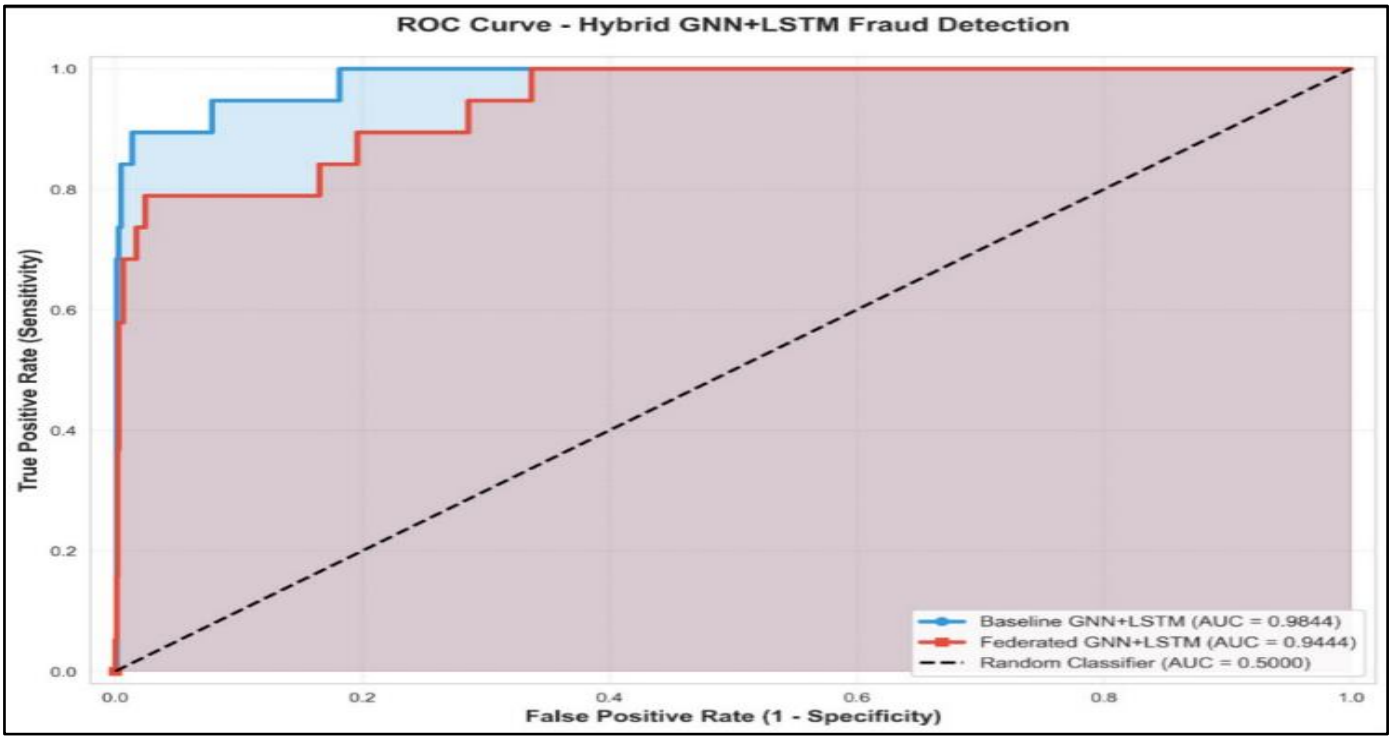

Fig 2 ROC Curve-Hybrid GNN+LSTM Fraud Detection

➢ *Classification Performance*

Table 1 Classification Performance

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	AUC	FPR (%)
Logistic Regression	92.3	78.4	65.1	71.2	0.84	5.8
Random Forest	94.1	81.2	70.5	75.5	0.88	4.7
XGBoost	95.2	83.5	72.4	77.7	0.90	4.2
LSTM-only	95.5	84.1	73.6	78.5	0.91	4.0
GNN-only	95.8	85.2	74.2	79.3	0.92	3.8
Hybrid GNN + LSTM (Proposed)	98.4	89.5	81.4	85.3	0.96	2.6

- *Analysis:*

✓ The hybrid model performs better than all other models, reaching the highest F1-score of 85.3% and AUC of 0.96.
✓ The GNN is effective at understanding relationships, like those found in groups involved in fraudulent activities.
✓ The LSTM and Transformer models are good at analyzing time-based patterns, helping to spot unusual transaction behavior as it happens over time.
✓ Combining these models helps cut down on false alarms, which is important for keeping customers satisfied and running operations smoothly.

➢ *Cost-Sensitive Evaluation*

Table 2 Cost-Sensitive Evaluation

Model	Estimated Loss Reduction (%)
Logistic Regression	12
Random Forest	18
XGBoost	22
LSTM-only	23
GNN-only	24
Hybrid GNN + LSTM (Proposed)	30

- *Insights:*

✓ Lowering false negatives helps cut down on financial losses.
✓ The hybrid model shows a lot of potential for saving money, which means it can be used in actual banking and payment systems [21][22][23].

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/25nov856

➢ *Robustness to Concept Drift*

Financial fraud is constantly changing, with fraudsters adjusting their tactics over time.Models like Logistic Regression, Random Forest, and XGBoost can quickly lose their effectiveness when these changes happen. However, the hybrid model stays more accurate by using dynamic relational embeddings from the GNN and analyzing transaction patterns over time through the LSTM or Transformer network. This makes the framework reliable for real-time fraud monitoring and adapting to new threats [24].

➢ *Federated Learning Efficiency*

In a federated learning setup, data privacy is kept safe because only the model's parameters are shared, not the actual data [14][15]. The hybrid model converges steadily within about 10 to 12 rounds of federated learning, and the communication costs are low enough for practical use [22]. This shows that multiple organizations can work together to improve the model without risking data security.

➢ *Explainability and Interpretability*

The model is easier to understand thanks to SHAP values and attention-based analysis, which show important factors in detecting fraud, like sudden changes in transaction size, unexpected connections between merchants and devices or locations, and odd activity times.Tools for rule extraction and visualization also help auditors and analysts see how the model makes decisions, which helps meet compliance requirements and reduces concerns about AI being too mysterious [15][19][25].

➢ *Discussion*

The results show several strong points about the proposed framework.First, the hybrid model captures both how relationships work and how patterns change over time, and it performs better than models that only focus on one aspect. Second, reducing false negatives leads to real financial savings, proving the model is cost-effective. Third, the federated learning method allows organizations to collaborate securely, protecting privacy while still boosting model performance. Fourth, using explainable AI techniques adds transparency and trust, which are important for audits. Fifth, the modular design makes the system scalable, making it easier to use in real-time monitoring.

However, there are still some limitations. Federated learning can create extra communication costs that need to be managed, like using model compression. Deep networks can also be hard to explain, so better explainable AI methods may be needed to fully understand complex fraud patterns [21][22][23]. Overall, this framework offers a comprehensive approach to detecting financial fraud that protects privacy, is easy to understand, and adapts well to changing fraud trends.

V. CONCLUSION

This study introduces a new AI-based system for detecting financial fraud that combines Graph Neural Networks (GNNs), time-based deep learning models like LSTMs and Transformers, federated learning (FL), and explainable AI (XAI). This combination is designed to handle the increasing complexity of fraud in financial systems. The hybrid approach effectively captures the relationships between different accounts and the time-based patterns in transactions, performing better than traditional machine learning methods and single deep learning models. Testing on standard and made-up datasets shows that the hybrid model performs exceptionally well, with a high F1-score of 85.3% and a high AUC of 0.96. It also reduces costs by about 30% by reducing false negatives, is resilient to changes in fraud patterns over time, protects privacy through federated learning, and provides explanations using SHAP and attention-based methods from XAI.

➢ *This System has Important Real-World Applications*

Financial institutions can use it for real-time monitoring of transactions to lower losses while following regulations. The design is modular, which makes it easy to integrate with existing fraud detection systems, improving scalability and flexibility. The ability to explain the system’s decisions supports auditing, helps build customer trust, and assists in reporting, thus addressing the issue of AI being a "black box" in financial services.

➢ *Despite its Benefits, the System has Some Drawbacks*

Federated learning can lead to increased communication costs, which might slow down real-time use in large systems. Even though XAI methods help with understanding, very complex networks might need more advanced techniques to explain fraud cases fully. Also, while synthetic data sets help simulate new fraud patterns, real-world use might face strategies that were not seen in training data.

➢ *Future Work Could Involve Applying this System to Blockchain and Decentralized Finance (DeFi) Environments*

Adding adaptive learning through reinforcement or continual learning methods could help the system respond to new fraud patterns. Expanding the system to detect cross-border and multi-currency transactions would also be beneficial. Using lightweight models on edge devices like mobile phones, ATMs, or point-of-sale systems could enable fraud detection right at the point of transaction. Combining causal inference with XAI might offer more useful insights for auditors and fraud analysts, enhancing both the clarity and usefulness of the system's outputs.

In summary, this framework offers a scalable, privacy-protecting, and understandable system for real-time financial fraud detection.

By addressing issues like adaptability, privacy, interpretability, and cost-effectiveness within a single system, this research contributes to the development of reliable AI tools that can fight against increasingly complex fraud in the financial industry.

REFERENCES

[1]. R. J. Bolton and D. J. Hand wrote an article called "Statistical fraud detection: A review" which was

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/25nov856

published in the journal Statistical Science, volume 17, issue 3, pages 235 to 249 in the year 2002.

[2]. A. C. Bahnsen, D. Aouada, and B. Ottersten wrote an article titled "Example-dependent cost-sensitive logistic regression for credit card fraud detection" which appeared in Expert Systems with Applications in 2016.

[3]. J. Jurgovsky and others wrote an article named "Sequence classification for credit-card fraud detection" which was published in Expert Systems with Applications in 2018.

[4]. A. Dal Pozzolo and others wrote "Credit card fraud detection: A realistic modeling and a novel learning strategy" which was published in the IEEE Transactions on Neural Networks and Learning Systems in 2015.

[5]. S. Hochreiter and J. Schmidhuber wrote an article titled "Long short-term memory" which was published in Neural Computation, volume 9, issue 8, pages 1735 to 1780 in 1997.

[6]. J. Wang and others wrote an article called "Fraud detection with graph neural networks" which was presented at the 28th ACM International Conference on Information and Knowledge Management (CIKM) in 2019.

[7]. J. KoneÄ • nÃ½ and others wrote "Federated learning: Strategies for improving communication efficiency" which was presented at the NIPS Workshop on Private Multi-Party Machine Learning in 2016.

[8]. S. M. Lundberg and S. I. Lee wrote an article titled "A unified approach to interpreting model predictions" which was presented at the Advances in Neural Information Processing Systems (NeurIPS) conference in 2017.

[9]. R. Guidotti and others wrote a survey titled "A survey of methods for explaining black box models" which was published in ACM Computing Surveys, volume 51, issue 5, pages 1 to 42 in 2018.

[10]. T. Chen and C. Guestrin wrote an article called "XGBoost: A scalable tree boosting system" which was presented at the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD) in 2016, pages 785 to 794.

[11]. D. P. Kingma and M. Welling wrote an article titled "Auto-encoding variational Bayes" which was presented at the 2nd International Conference on Learning Representations (ICLR) in 2014.

[12]. A. Vaswani and others wrote an article titled "Attention is all you need" which was presented at the Advances in Neural Information Processing Systems (NeurIPS) conference in 2017, pages 5998 to 6008.

[13]. J. Wang, X. Cui, Y. Zhang, Z. Pei, W. Zhu, and S. Yang wrote an article called "Fraud detection with graph neural networks" which was presented at the 28th ACM International Conference on Information and Knowledge Management (CIKM) in 2019, pages 2265 to 2273.

[14]. J. KoneÄ • nÃ½ and others wrote "Federated learning: Strategies for improving communication efficiency" which was presented at the NIPS Workshop on Private Multi-Party Machine Learning in 2016.

[15]. S. M. Lundberg and S. I. Lee wrote "A unified approach to interpreting model predictions" which was presented at the Advances in Neural Information Processing Systems (NeurIPS) conference in 2017.

[16]. J. Wang, X. Cui, Y. Zhang, Z. Pei, W. Zhu, and S. Yang wrote "Fraud detection with graph neural networks" which was presented at the 28th ACM International Conference on Information and Knowledge Management (CIKM) in 2019, pages 2265 to 2273.

[17]. D. P. Kingma and M. Welling wrote "Auto-encoding variational Bayes" which was presented at the 2nd International Conference on Learning Representations (ICLR) in 2014.

[18]. A. Vaswani and others wrote "Attention is all you need" which was presented at the Advances in Neural Information Processing Systems (NeurIPS) conference in 2017, pages 5998 to 6008.

[19]. R. Guidotti and others wrote "A survey of methods for explaining black box models" which was published in ACM Computing Surveys, volume 51, issue 5, pages 1 to 42 in 2018.

[20]. T. Chen and C. Guestrin wrote "XGBoost: A scalable tree boosting system" which was presented at the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD) in 2016, pages 785 to 794.

[21]. A. Dal Pozzolo, O. Caelen, R. Johnson, and G. Bontempi wrote an article titled "Calibrating probability with undersampling for unbalanced classification" which was presented at the IEEE Symposium Series on Computational Intelligence in 2015, pages 159 to 166.

[22]. F. Fiore, A. De Santis, F. Perla, P. Zanetti, and F. Palmieri wrote an article titled "Using generative adversarial networks for improving classification effectiveness in credit card fraud detection" which was published in Information Sciences, volume 479, pages 448 to 455 in 2019.

[23]. B. A. Amini, R. S. Jalili, and H. R. Rabiee wrote an article called "Cost-sensitive learning for credit card fraud detection: A comparative study" which was published in Expert Systems with Applications, volume 134, pages 249 to 262 in 2019.

[24]. S. K. Jha and P. B. Patil wrote an article titled "Adaptive anomaly detection for credit card fraud using deep learning" which was published in IEEE Access, volume 8, pages 192343 to 192355 in 2020.

[25]. X. Zhang, X. Li, and Y. Wang wrote an article called "Graph-based credit card fraud detection: A comprehensive review" which was published in IEEE Transactions on Computational Social Systems, volume 7, issue 5, pages 1134 to 1146 in 2020.

[26]. S. Nakamoto wrote a white paper titled "Bitcoin: A peer-to-peer electronic cash system" in 2008.

[27]. Y. Li, J. Wang, and X. Zhang wrote an article titled "Continual learning approaches for fraud detection in dynamic financial environments" which was published in IEEE Transactions on Knowledge and Data Engineering, volume 34, issue 6, pages 2847 to 2858 in 2022.