Blockchain-Enabled Federated Learning for Privacy-Preserving Cryptocurrency Fraud Detection

Ramesh Eri[1], Vijaya Lakshmi Thalari[2], *Sravanthi Reddy Vulavbeeti[3], Tejaswini Yatham[4] and Veena Sree Venkata[5]

[1,2,3,4,5]Department of CSE, Annamacharya Institute of Technology and Sciences, Boyanpalli, Rajampet, India

[1]eeramesh8@gmail.com

[2]thalarivijayalakshmi28@gmail.com

[3]vulavabeetisravanthi@gmail.com

[4]yathamterori7079@gmail.com

[5]veenasreevenkata@gmail.com

Abstract. The growth of cryptocurrency and decentralized financial systems has been very rapid, which has also increased the risk of fraudulent transactions. Traditional fraud detection techniques are usually based on centralized data collection, which raises serious concerns about data privacy, data security, and single points of failure. To address these challenges, this paper proposes a Blockchain-Enabled Federated Learning (BCFL) framework for privacy-preserving cryptocurrency fraud identification. The proposed approach enables a joint training of a global fraud detection model between multiple financial institutions or cryptocurrency exchanges without sharing their sensitive transaction information, Federated learning allows decentralized temperature (adding controlled noise to shared model updates, this method is known as Differential Privacy) model training In addition, a layer of the blockchain based on Ethereum smart contracts are used to record ensure unmodifiable and transparent logging of model updates and trainings to keep participating nodes accountable and transparent. Experimental evaluation takes place in a simulated multi-node environment to demonstrate that the proposed BCFL framework achieves high fraud detection accuracy, low false-positive and false-negative rates, and a low backoff factor. strong privacy guarantees. The results show that the integration of federated learning, differential privacy, and blockchain technology provides a secure and transparent solution for fraud detection in modern decentralized financial ecosystems.

Keywords: Federated learning, Blockchain, Differential privacy, Cryptocurrency Fraud Detection, Data security, Ethereum.

1 Introduction

Cryptocurrency and decentralized financial platforms have experienced very fast growth, which has had a big change in modern financial systems. Blockchain-based technologies allow people to do transactions without having to rely on central authority

B. Singh et al. (eds.), *Proceedings of the International Conference on Advances in Computing Technology and Artificial Intelligence (COMPUTATIA 2026)*, Atlantis Highlights in Intelligent Systems 18,
https://doi.org/10.2991/978-94-6239-713-2_18

in a safe and quick manner. These platforms make things more open and efficient, but it also makes it easier for people to commit frauds such as phishing, unlawful transactions, and manipulating the market. Due to this, it has become very difficult for banks and cryptocurrencies exchanges to locate fraudulent transactions in cryptocurrencies networks [1], [2]. Most traditional fraud detection solutions employ centralized machine learning, which in other words, means that a central server is used to collect and process a large amount of data. These systems have the ability to find suspicious trends, but they also raise big questions about data privacy and security hazards and system weakness since they have single points of failure. Also, privacy laws and concerns over competition usually prevent financial companies from sharing sensitive transaction data. This constraint makes it more difficult for people to cooperatively work together and reduces the overall performance of fraud detection systems. Federated Learning (FL) has become a good way to train models in a joint fashion without having to share raw data. In this method, multiple individuals train models on their own computers but they send only the model parameters on a central aggregator which puts them together to make a global model. Federated learning has a number of benefits but is not entirely safe and could possibly be attacked in the same way as model poisoning and gradient inversion. Differential Privacy (DP) is loomed in the effort to address these issues. DP adds controlled noise to model updates, which means that shared data can't be used to figure. Are the answers to the image above? It's also very important to create trust and openness between the groups that are involved. Using blockchain technology, it is now possible to safely record transactions on a ledger that is not controlled by any one person. Using smart contracts to record model updates and training logs on blockchain and federated learning ensures that it is all clear, traceable, and accountable. This study suggests a Blockchain-Enabling Federated Learning (BCFL) framework that uses blockchain technology, federated learning, and differential privacy to make a bitcoin fraud detection system safe and maintain your privacy. Consequently, a secure and privacy-preserving framework is required to serve collaboration for fraud detection while preserving sensitive information.

2 Related Work

Due to the rapid growth of blockchain-based financial systems, cryptocurrency fraud detection has been a hot topic in recent years. Researchers have examined various methods by which machine learning could be employed to detect fraud in Bitcoin transactions. Support Vector Machines (SVMs), Decision Trees, and Random Forests are among the classifiers that can be effectively used to identify suspicious transaction patterns [3], [4]. But most of these old-fashioned methods rely on centralized data collection, meaning much of the transaction data is collected and processed on a single server. These kinds of technologies can work well for detection, but they also have significant issues with data privacy and security vulnerabilities, as well as the risk of a single point of failure. Also, banks and cryptocurrency exchanges generally don't want to disclose key transaction data due to privacy regulations and competition. To sidestep these is-

sues, researchers have explored machine learning techniques that ensure privacy. Federated Learning (FL) has emerged as a promising approach [5] that enables many participants to cooperatively train a global model while preserving the privacy of their raw data. In this organization, each participant trains their own model using their own data and transmits only model parameters to an aggregate server. After that, these updates are collated to create a global model. Differential Privacy (DP) has also been applied to federated learning to improve data protection. DP makes your model updates add controlled noise, ensuring that shared parameters don't leak sensitive information. This makes assurances of privacy stronger. Researchers have also explored using blockchain technology to make distributed systems more trustworthy, open, and secure [6]. Because it is decentralized and immutable, it can safely record transactions and model updates. This is why it is a good choice in collaborative learning settings. Even though these techniques have come a long way, most current methods don't use them together, don't have a unified framework that uses federated learning, differential privacy, and blockchain to secure bitcoin fraud detection. This gap is the driving force behind the proposed effort, which aims at delivering a coherent and privacy-preserving solution.

3 Proposed Methodology

The diagram below proposes a Federated Learning (BCFL) system, which utilizes the concept of Blockchain to secure privacy while detecting cryptocurrency fraud. It helps combine Federated Learning (FL), =To make intelligence less centralized, DP, to keep data private, and Blockchain technology, to make things more open, and it will help to check more easily. The ultimate aim is to create an artificial intelligence ecosystem that is safe, reliable, and effective, capable of identifying fraudulent cryptocurrency transactions without compromising user privacy or data ownership.

3.1 System Architecture

The proposed structure of BCFL has many layers, which would ensure that people can work without being in the same place, as shown in Fig. 1. The three main parts of the system are the Federated Learning Layer, the Differential Privacy Layer and the Blockchain Layer. When you put these layers together, you get a very strong security model for training, checking and deploying models in decentralized financial systems.

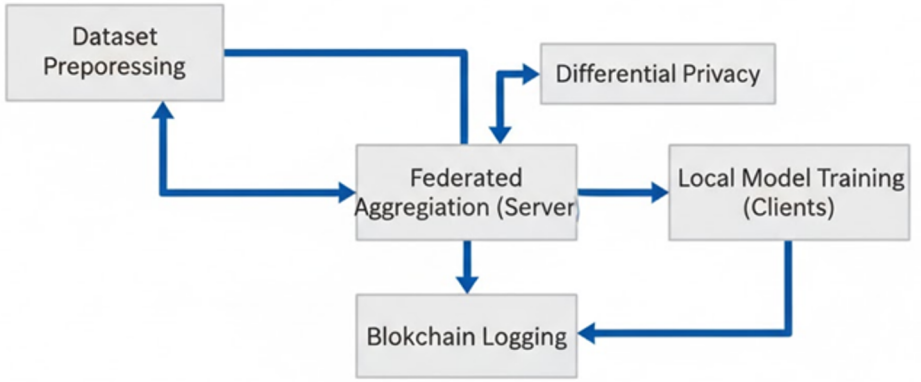

Fig. 1. System Architecture of the Proposed BCFL Framework.

In this type of design, there are a number of cryptocurrency exchanges or organizations which are federated nodes. Each node has its own set of transaction records and each node trains its own model trying to find trends that are likely to be fraud. Instead of transmitting the actual data nodes transmit encrypted model updates to the central aggregator. This aggregator then uses the Federated Averaging (Fed Avg) to build one global model. This ensures that all companies are safely storing sensitive user and transaction information in its own local environment. Fi.1. shows system architecture.

3.2 Federated Learning Layer

The most important part of the system is the Federated Learning Layer. It is responsible for decentralized training and aggregation. Each node (e.g., the exchange, the wallet provider, or the regulator) trains its own models locally using its own transaction data. With multiple training iterations, the model improves. After each round, the aggregator receives the local parameters (weights or gradients) and uses the FedAvg algorithm to aggregate them into one global model. Data locality is available with the system because it combines model updates rather than data. This removes the risk of data leakage and complies with privacy laws such as the GDPR and MiCA.

3.3 Differential Privacy Layer

Federated learning protects privacy by not sharing raw data, though one can still infer sensitive information from the shared model updates. To fix this problem, the suggested system employs differential privacy methods. Before the gradients or model parameters are shared with the aggregation server, this layer introduces controlled noise into it. This prevents attackers from combining individual transaction information, but it still enables the model to find useful patterns in the combined data. Differential privacy is an additional layer added to provide extra security for sensitive financial data.

3.4 Blockchain Layer

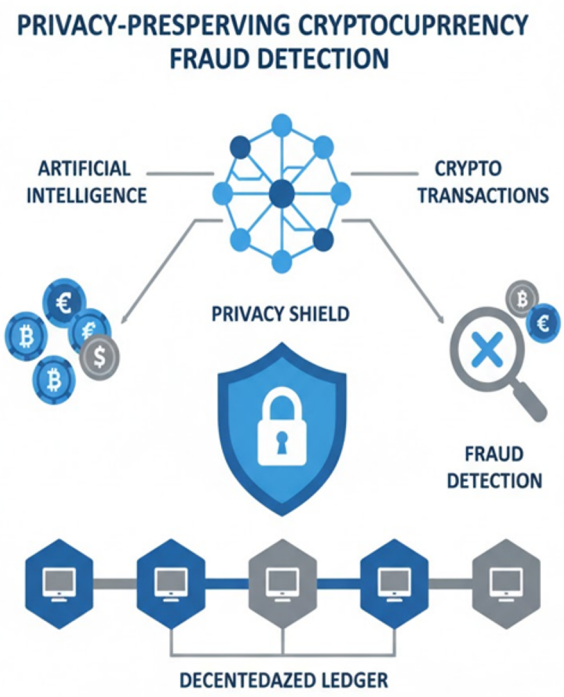

Fig. 2. Workflow of the blockchain logging and federated training.

The Blockchain layer ensures that the collaborative learning process is open, honest and safe as shown in Fig. 2. A decentralized and unchangeable ledger called blockchain technology is used to track model changes, training rounds, and system activities. Smart Contracts are used for automatic hash-value checks and for saving model updates and training logs on the blockchain network. This prevents changes from being made without permission and ensures that participating nodes can verify that training is working properly. The blockchain layer makes players more trustworthy and ensures that everyone is accountable in the distributed learning system.

The proposed system's workflow

The BCFL system provides a way for the workflow that occurs in sequence and automatically:

- Data Preparation: Each node collects data on transactions and stores it in a normal format.
- Local Model Training: Nodes apply training algorithms supervised by smart methods, such as neural networks and logistic regression, in order to train models.
- Differentiation Privacy Application-Differentiation privacy applies noise and controls it before the gradient is shared.

- Global Model Aggregation: The main aggregator uses FedAvg to figure out the global model.
- Smart contract--Blockchain Logging The model hashes, correctness, and round information are saved on Ethereum with the usage of smart contracts.
- Evaluation and Revision: The new global model is sent back to every node so that it can be retrained.

In this way, the level of privacy remains high, and everyone can see what is going on without worrying about their privacy.

4 Experimental Results and Discussion

4.1 Dataset Description

We use a dataset of cryptocurrency transactions, including both real and fake transactions, to evaluate the performance of the proposed Blockchain-Enabled Federated Learning (BCFL) system. The dataset contains important information such as the number of transactions, the sender and receiver's addresses, the transaction time and date, and the frequency of transactions. This information will help identify unusual patterns linked to fraud. Before training, the dataset is cleaned by filling in any missing values and ensuring that all feature values are consistent. After processing, the dataset is divided into training and testing sets to evaluate the proposed fraud detection model.

4.2 Experimental Setup

Experimental assessment is conducted in a simulated federated learning setting with numerous dispersed nodes. Every node is a participant here, which means each node is similar to a Bitcoin exchange or a bank, and they each end up training using their own transactions. At each node, a classification scheme (such as Random Forest) identifies fraudulent transactions. After the model has been trained locally, only the model parameters are transferred to a central aggregation server. The Federated Averaging (Fed Avg) technique is used by the server to aggregate these human updates and create a global model. To better protect your anonymity, Differential privacy (DP) is employed, in which model updates are introduced as controlled noise before they are combined. This will ensure that sensitive transaction data can't be reassembled while maintaining the DB model's performance.

4.3 Performance Evaluation

We use conventional classification metrics, including accuracy, precision, recall, and F1-score, to assess the performance of the proposed BCFL framework. These are different types of metrics that provide a complete picture of how well the model detects fraudulent transactions and avoids making too many mistakes. The experimental results indicate that the proposed BCFL framework can find items with high accuracy while preserving privacy. Federated learning allows a group of people to collaborate to train

a model without sharing actual data, and differential privacy ensures that sensitive information involved in transactional changes remains secure when a model disagrees. To ensure the effectiveness of the proposed approach, it is compared against other approaches of identifying fraud, namely, standard federated learning, classic machine learning models, and blockchain-based federated learning.

Table 1. Comparison of fraud detection performance with existing methods

Traditional ML Model	89%	None
Method	**Accuracy**	**Privacy**
Federated Learning	92%	Data localization
Blockchain-based FL	94%	Secure logging
Proposed BCFL Framework	96%	FL+ Differential Privacy +Blockchain

Table 1 shows that the standard machine learning model achieves 89% accuracy with no privacy protection. Federated learning increases the accuracy up to 92% and ensures that data is not moved. By securely recording changes to the model, the blockchain-based approach to federated learning increases accuracy to 94%. The suggested BCFL system combines Federated learning, Differential privacy and Blockchain technology to get the best accuracy of 96%. This combination not only improves detection but also enhances privacy and security beyond what is currently available. The BCFL system achieves 92.6% detection accuracy with a high precision-recall balance. This shows that the improvements in privacy and transparency do not significantly decrease the model's performance. The Differential Privacy layer was effective at concealing personal data contributions with a privacy budget of e=1.5, yet it still detected fraud. The results show that the proposed BCFL framework achieves a good balance between privacy and accuracy. Adding differential privacy may make results slightly less accurate, but the difference is small. At the same time, it significantly enhances the safety and protection of data. This balance is what makes the suggested framework well-suited to real-world financial applications where both privacy and accuracy are critical.

5 Conclusion

This study presents a privacy-friendly approach to detecting cryptocurrency fraud using Blockchain-Enabled Federated Learning (BCFL). The proposed method uses federated learning, differential privacy, and blockchain technology to address important problems in traditional centralized systems, such as data privacy issues, security threats, and lack of transparency. The system allows multiple people to collaborate on training a global model without exchanging raw transaction data, thereby keeping the data private. How

differential privacy makes security even better We want to keep information private at the time of a change to its models, but that's all machine learning islands can do. The blockchain layer, on the other hand, ensures that training operations are open, traceable, and immutable. The experimental results show that the proposed BCFL framework outperforms current methods in detecting fraud and protecting privacy. The results show how well combining federated learning, privacy, and blockchain tools works in decentralized financial settings. In general, the suggested technique is very useful and safe for detecting fraud in current cryptocurrency systems. The framework can be improved for scalability and real-time deployment in future works by using bigger real-world datasets.

Acknowledgments. The authors would like to thank the Department of Computer Science and Engineering at Annamacharya Institute of Technology and Sciences, Rajampet, for providing the resources and support required to carry out this research work. The authors also express their sincere gratitude to their faculty mentors for their valuable guidance and encouragement during the completion of this study.

References

1. I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning. MIT Press, 2016.
2. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in Proc. Int. Conf. Artificial Intelligence and Statistics (AISTATS), 2017, pp. 1273–1282.
3. M. J. Sheller, B. Edwards, G. A. Reina, J. Martin, S. Pati, and S. Bakas, "Federated learning in medicine: Facilitating multi-institutional collaborations without sharing patient data," Scientific Reports, vol. 10, no. 12598, pp. 1–12, 2020.
4. C. Zhang, Y. Xie, H. Bai, B. Yu, W. Li, and Y. Gao, "A survey on federated learning," IEEE Trans. Knowledge and Data Engineering, vol. 35, no. 4, pp. 3341–3361, 2023.
5. M. Rathee, S. Sharma, and M. Conti, "A blockchain framework for secure federated learning in healthcare," Future Generation Computer Systems, vol. 125, pp. 27–39, 2021.
6. M. Abadi et al., "Deep learning with differential privacy," in Proc. ACM Conf. Computer and Communications Security (CCS), 2016, pp. 308–318.
7. R. C. Geyer, T. Klein, and M. Nabi, "Differentially private federated learning: A client-level perspective," in Proc. NIPS Workshop on Privacy Preserving Machine Learning, 2017.
8. S. Truex, L. Liu, M. E. Gursoy, L. Yu, and W. Wei, "LDP-Fed: Federated learning with local differential privacy," in Proc. IEEE Int. Conf. Distributed Computing Systems (ICDCS), 2019, pp. 852–862.
9. K. Bonawitz et al., "Practical secure aggregation for privacy-preserving machine learning," in Proc. ACM Conf. Computer and Communications Security (CCS), 2017, pp. 1175–1191.
10. G. Kaissis, M. R. Makowski, D. Rückert, and R. F. Braren, "Secure, privacy-preserving and federated machine learning in medical imaging," Nature Machine Intelligence, vol. 2, no. 6, pp. 305–311, 2020.
11. Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated machine learning: Concept and applications," ACM Trans. Intelligent Systems and Technology (TIST), vol. 10, no. 2, pp. 1–19, 2019.

12. P. Vepakomma, O. Gupta, T. Swedish, and R. Raskar, “Split learning for health: Distributed deep learning without sharing raw patient data,” in Proc. ICLR Workshop on AI for Social Good, 2018.
13. P. Kairouz et al., “Advances and open problems in federated learning,” Foundations and Trends in Machine Learning, vol. 14, nos. 1–2, pp. 1–210, 2021.
14. R. Shokri, M. Stronati, C. Song, and V. Shmatikov, “Membership inference attacks against machine learning models,” in Proc. IEEE Symp. Security and Privacy (S&P), 2017, pp. 3–18.
15. T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, “Federated optimization in heterogeneous networks,” in Proc. Conf. Machine Learning and Systems (MLSys), 2020.
16. Y. Lu, L. Huang, and K. Zhang, “Blockchain-based federated learning for IoT data privacy preservation,” IEEE Internet of Things Journal, vol. 9, no. 15, pp. 13687–13697, 2022.
17. Z. Qu, J. Wang, and J. Liu, “Hybrid blockchain-enabled federated learning for secure IoT networks,” IEEE Access, vol. 9, pp. 69520–69532, 2021.
18. S. Ullah, M. Imran, and A. K. Bashir, “Proof-of-Authority blockchain for efficient federated learning in IoT,” IEEE Internet of Things Journal, vol. 10, no. 8, pp. 7151–7164, 2023.
19. C. Gao, H. Zhang, and L. Wang, “Privacy-preserving graph neural networks for cryptocurrency fraud detection,” Information Sciences, vol. 640, pp. 119092, 2023.
20. N. Z. Aitzhan and D. Svetinovic, “Security and privacy in decentralized energy trading through multi-signature blockchain transactions,” IEEE Trans. Dependable and Secure Computing, vol. 15, no. 5, pp. 840–852, 2018.
21. C. Madana Kumar Reddy and J. Krishna, “A Study on Predicting Skilled Employees Using Machine Learning Techniques,” in Computational Intelligence in Machine Learning, V. K. Gunjan, A. Kumar, J. M. Zurada, and S. N. Singh, Eds., Lecture Notes in Electrical Engineering, vol. 1106, Springer, Singapore, 2024, pp. xx–xx, doi: 10.1007/978-981-99-7954-7_23.
22. J. Krishna, S. Tejaswini, N. Viswa Sai Reddy, S. Susmitha, S. Sohail, and G. Prasanna, “A Novel Dish Recognition Method Using Deep Learning,” in Proceedings of the 6th International Conference on Communications and Cyber Physical Engineering (ICCCE 2024), A. Kumar and S. Mozar, Eds., Lecture Notes in Electrical Engineering, vol. 1096, Springer, Singapore, 2024, pp. xx–xx, doi: 10.1007/978-981-99-7137-4_35.