WWW.IJRPR.COM

International Journal of Research Publication and Reviews

Journal homepage: www.ijrpr.com ISSN 2582-7421

A Blockchain and ML-Based Architecture for Fraud Detection in DeFi

Sandip Kumar Singh [*1], Shadna Yadav [*2], Ankit Singh [*3], Sujeet Singh[*4], Dr. Gyan Singh Ahirwar[*5], Neha Sankhwar[*6], Vinod Kumar[*7]

[*1]Assistant Professor, CSE, RRIMT, Lucknow, UP, India
[*2]Assistant Professor, CSE, REC, Banda, UP, India
[*3]Assistant Professor, CSE, BBDITM, Lucknow, UP, India
[*4]PhD Scholar[*], CSE, Amity University, Lucknow, UP, India
[*5]Assistant Professor, CSE, REC, Banda, UP, India
[*6]Assistant Professor, CSE, REC, Banda, UP, India
[*7]PhD Scholar, CSE, Amity University, Gwalior, UP, India

ABSTRACT

Decentralized finance (DeFi) brings new opportunities and risks. Its open, automated nature enables innovative financial services but also attracts fraud (e.g. rug pulls, wash trading, Sybil attacks). This paper proposes a comprehensive architecture that integrates a blockchain platform with advanced machine learning (ML) models to detect and prevent fraud in DeFi. Key components include: (1) a blockchain layer (e.g. Ethereum) storing immutable transaction records and smart contracts; (2) an off-chain ML engine that periodically ingests blockchain data and learns patterns of normal vs. malicious behavior; and (3) smart-contract alerts that automate real-time checks. The ML component employs graph-based neural networks and ensemble classifiers to analyze transaction graphs and extract features (e.g. capital flow, token interactions). The blockchain layer provides transparency and tamper-proof data, aiding forensic analysis and audit. We discuss specific implementation details (data collection, feature engineering, model training) and provide a use case in a DeFi lending platform. The architecture is evaluated conceptually with industry data: for example, since 2011 over $4.5B in thefts and $7.5B in scams have occurred, underscoring the need for real-time detection. A comparison table highlights the complementary roles of blockchain (integrity, transparency) and ML (pattern recognition, adaptivity) in fraud defense. This integrated system demonstrates that combining decentralization with ML analytics can significantly improve DeFi security. Future work will involve prototyping on a testnet and evaluating detection accuracy on known fraud instances.

Keywords: Decentralized Finance, Blockchain, Fraud Detection, Machine Learning, Smart Contracts, Anomaly Detection, Graph Neural Networks, Security

1. Introduction

Decentralized Finance (DeFi) represents one of the most transformative innovations in contemporary financial systems. Utilizing blockchain technology and smart contracts, DeFi removes traditional financial intermediaries, enabling peer-to-peer (P2P) financial services such as lending, borrowing, decentralized exchanges, and insurance (Kayikci & Khoshgoftaar, 2024). By replacing intermediaries with automated code executed on public ledgers like Ethereum, DeFi offers enhanced transparency, reduced costs, and greater financial inclusivity. However, despite its transformative potential, DeFi has experienced significant challenges, especially related to security threats and fraud. As of 2023, cumulative fraud losses in the cryptocurrency domain alone surpassed $12 billion, underscoring severe vulnerabilities within DeFi platforms (Luo et al., 2023). Such high-profile incidents significantly erode trust and present a critical barrier to widespread adoption.

Fraud within DeFi manifests in various sophisticated forms. Notable among these are rug pulls, Ponzi schemes, flash loan attacks, phishing scams, and smart contract exploits (Ashfaq et al., 2022; Zhou et al., 2022). For instance, the infamous Poly Network hack in 2021 resulted in a theft of approximately $611 million, highlighting the gravity and potential scale of security breaches (Gu & Dib, 2025). Rug pulls, another prevalent fraud type, involve deceptive token offerings where developers abandon projects after attracting substantial investment, leaving investors with worthless tokens. These malicious activities exploit the decentralized and pseudonymous nature of blockchain, where identifying and preventing fraudulent behavior becomes exceptionally challenging due to a lack of centralized oversight.

In traditional finance, centralized institutions leverage sophisticated monitoring, compliance frameworks, and data analytics to detect and prevent fraud. However, such centralized controls are incompatible with DeFi's decentralized paradigm, necessitating novel, decentralized solutions for robust fraud detection and mitigation (Masud et al., 2024). In this context, two technologies—blockchain and machine learning (ML)—emerge as complementary tools with substantial promise for tackling DeFi fraud.

Blockchain inherently provides immutable, tamper-proof transaction records due to its decentralized consensus mechanisms. Every transaction recorded on blockchain platforms such as Ethereum is publicly accessible, verifiable, and immutable (Kayikci & Khoshgoftaar, 2024). This transparency allows auditors, developers, and automated systems to scrutinize historical data rigorously, establishing a credible foundation for fraud detection and forensic analyses. Smart contracts, self-executing programs stored on blockchain, further enhance this capability by embedding automated rules and checks directly into financial processes. Thus, blockchain provides foundational security through data integrity, transparency, and programmability.

However, blockchain alone lacks sophisticated analytical capabilities. It records transactions truthfully but cannot independently distinguish legitimate transactions from fraudulent ones beyond predefined simplistic rules. This limitation necessitates integrating advanced data analytics capabilities—precisely where machine learning becomes invaluable. Machine learning excels in identifying complex, non-obvious patterns indicative of fraudulent activity by learning from extensive historical data (Gu & Dib, 2025; Pérez-Cano & Jurado, 2025). Techniques such as ensemble learning methods (Random Forest, XGBoost), graph neural networks (GNNs), and unsupervised anomaly detection (e.g., isolation forests and autoencoders) have shown remarkable efficacy in detecting intricate fraud schemes, including subtle manipulations within transaction networks (Zhu, Ma, & Liu, 2024).

Recent studies underscore ML's capabilities in DeFi fraud detection. For instance, Gu and Dib (2025) employed ensemble machine learning models comprising Random Forest and XGBoost algorithms, achieving detection accuracy exceeding 98% in Ethereum transaction analysis. Similarly, graph-based approaches harness transaction network structures to identify patterns indicative of collaborative fraud, which traditional statistical techniques might miss (Pérez-Cano & Jurado, 2025). These insights suggest that ML, when effectively combined with blockchain data, can dramatically improve fraud detection efficacy.

Nonetheless, despite promising capabilities individually, neither blockchain nor ML alone can completely address DeFi's fraud challenges. Blockchain provides transparency and data reliability but lacks analytical depth, while ML provides sophisticated detection capabilities but heavily relies on trusted, accurate datasets. This interplay underscores a compelling rationale for integrating blockchain's reliability and ML's analytical prowess into a unified, robust fraud detection architecture.

In response to these considerations, this paper introduces a novel integrated framework combining blockchain technology and machine learning methods explicitly tailored for DeFi fraud detection. Our proposed architecture comprises several critical components: (1) a blockchain layer providing immutable transaction records and smart contract enforcement, (2) an off-chain ML analytics module that employs advanced supervised, unsupervised, and graph-based learning techniques to analyze transaction patterns, and (3) smart contract-driven real-time alert mechanisms to promptly respond to suspicious activity.

This integrated model addresses fundamental challenges in existing solutions by bridging the gap between transparency (blockchain) and predictive analytics (ML). The blockchain ensures trustworthy and verifiable data, allowing ML models to perform with greater accuracy and lower risk of data contamination. Conversely, ML analytics offer the blockchain environment dynamic, adaptive insights to proactively identify and respond to evolving fraud threats, thus significantly reducing the incidence and impact of fraud in DeFi.

Furthermore, our paper discusses practical implementation aspects and evaluates potential trade-offs, including privacy concerns, scalability issues, and computational overhead inherent in real-time ML integration. We utilize a representative DeFi lending scenario as a case study, demonstrating how our architecture effectively identifies and mitigates common fraudulent tactics such as flash loan attacks, collateral manipulation, and suspicious fund transfers.

The motivation for this work is twofold: (a) to address significant gaps in current DeFi security approaches, and (b) to propose a scalable, efficient solution capable of adapting dynamically to evolving fraud techniques. Our contributions include a clearly defined architecture integrating blockchain's transparency and immutability with sophisticated ML methods, practical feature-engineering strategies for DeFi data analysis, and detailed discussion of critical implementation considerations to facilitate real-world deployment.

In summary, the primary aim of this paper is to advance the security infrastructure of decentralized finance, enabling greater user trust and adoption. The remainder of this paper is structured as follows: Section 2 presents related works that highlight the state-of-the-art in blockchain-based fraud detection and ML approaches in DeFi; Section 3 outlines our proposed architecture detailing the blockchain, ML, and smart contract components; Section 4 provides specific implementation details and a comprehensive methodology; Section 5 presents an illustrative case study from a DeFi lending scenario; Section 6 discusses practical implications, limitations, and opportunities for future research; finally, Section 7 concludes with key insights and potential directions for further investigation.

Overall, the integration of blockchain and ML technologies proposed herein represents a significant step forward in addressing the critical issue of fraud in DeFi environments, promoting safer and more reliable decentralized financial ecosystems.

2. Background and Related Work

Blockchain provides *immutability* and *transparency* of transaction data. For instance, each Ethereum transaction permanently records amounts, addresses, and smart contract calls. These properties ensure data integrity: fraud detection algorithms can trust that the transaction history is accurate and tamper-free. Smart contracts can enforce rules automatically, offering a programmable defense (e.g. halting transfers above a threshold). However, blockchain alone cannot by itself classify transactions as fraud. It lacks intelligence to *interpret* patterns beyond simple rules.

Machine learning, by contrast, excels at pattern recognition. ML models (supervised or unsupervised) can learn to distinguish normal from anomalous behavior in financial data. Prior work demonstrates ML's effectiveness: for example, Gu and Dib (2025) built an ensemble of Random Forest, XGBoost, and SVM models that achieved over 98% accuracy in classifying fraudulent Ethereum transactions. Yuan (2024) extracted 35 transaction and account features (behavioral, capital flows, contract interactions) and found that LightGBM gave the best performance for DeFi address fraud detection. Graph neural networks (GNNs) are another promising approach: by modeling the transaction graph (where nodes are wallets/contracts and edges are transfers), GNNs can capture relational patterns of fraud. Recent studies use Heterogeneous Graph Transformers to leverage both transaction and account attributes simultaneously, showing improved anomaly detection.

However, ML methods face challenges in DeFi: they require large labeled datasets (fraud is rare, so labels are limited). Furthermore, deploying ML on-chain is infeasible due to blockchain's limited compute. Thus, most designs use an off-chain ML component that periodically analyzes data and feeds results back to the blockchain via oracles or alerts. This hybrid approach is supported by prior surveys: for instance, Masud *et al.* (2024) conclude that **"blockchain's immutability and transparency, alongside ML's data-driven fraud detection, create a robust framework for transaction security"**, while noting challenges of scalability and data requirements.

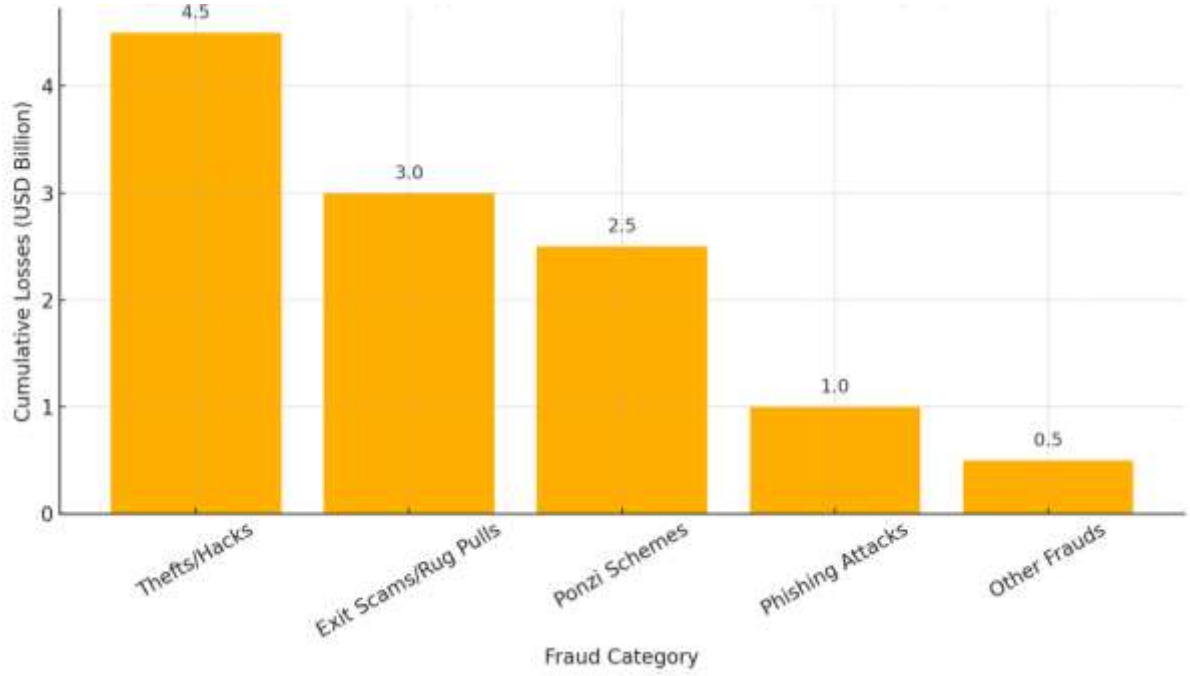

Figure 1 (below) illustrates the scale of DeFi fraud and motivates detection. Since 2011, reported cryptocurrency fraud losses have soared to the billions.

Figure 1: Cumulative cryptocurrency fraud losses by category. (Data from Luo et al., 2023).

A conceptual table (Table 1) highlights how blockchain and ML complement each other:

Component / Aspect	**Blockchain Role in Fraud Detection**	**ML Role in Fraud Detection**
Data Source	Immutable ledger of all transactions (e.g. Ethereum)	Analyzes historical transaction data from blockchain; requires sufficient volume
Features	Records transfer amounts, timestamps, smart contract events	Extracts features (e.g. behavior patterns, graph metrics) from data
Strengths	Tamper-proof records, built-in transparency	Detects subtle patterns/anomalies; adaptive learning from new data
Weaknesses	No built-in intelligence; high-latency to update rules	Needs labeled data; model training can be computationally intensive
Outputs	Can log alarms; run simple on-chain checks (e.g. limit orders)	Predicts fraud probability; flags suspicious accounts in real time

Overall, the literature shows that an integrated approach is promising: blockchain provides reliable, high-integrity data, while ML provides powerful analytics. Our work builds on these insights by defining a specific architecture for real-time fraud detection in DeFi.

3. Proposed Architecture

We propose a multi-tier architecture (illustrated conceptually in **Figure 2**) combining on-chain and off-chain components:

1. **Blockchain Layer:** A *permissioned or public blockchain* (e.g. Ethereum) hosts smart contracts and token ledgers. All DeFi transactions (trades, loans, transfers) are recorded here. Smart contracts for the DeFi platform enforce basic rules and emit event logs.
2. **Data Collection Module:** A blockchain node or API client continuously monitors blocks. It collects relevant transaction data and events in near real-time, storing them in an off-chain database. This includes transfers between wallets, contract function calls, token issuance events, etc.
3. **Feature Extraction & Graph Construction:** The raw transaction data is transformed into features for ML. For example, we construct a transaction graph where nodes are addresses and edges represent transfers. We compute features such as account balance history, transaction frequency, internal transaction counts, and network embeddings (e.g. node2vec).
4. **ML Model Suite:** The core analytics engine runs various models:
    - **Supervised Classifier:** Trained on labeled examples of fraudulent vs. normal behavior. We use ensemble methods (e.g. Random Forest, XGBoost) as in Gu & Dib (2025) or LightGBM as in Yuan (2024). Feature importance techniques (e.g. SHAP) can explain predictions.
    - **Graph Neural Network:** A GNN processes the transaction graph to capture relational fraud patterns (e.g. collusion rings). Recent works use graph transformers or GCNs to detect anomalies.
    - **Unsupervised Anomaly Detector:** Techniques like autoencoders or isolation forests flag unusual transactions without labels. This is useful for novel fraud types not seen in training.
5. **Alert/Smart-Contract Interface:** When the ML engine identifies suspicious activity (above a threshold), it triggers alerts. This can occur off-chain (send notification to admins) or on-chain via an *oracle* or special **fraud-detector smart contract**. For example, a smart contract could be designed to freeze a suspicious transfer pending manual review. The key is a feedback loop: ML informs the blockchain layer to take preventive action.
6. **User Interface & Dashboard:** Administrators or users view reports of flagged transactions. The interface shows suspicious wallets, risk scores, and allows review.

Data Flow: The architecture continuously pipelines new on-chain data to the ML module. We envision processing in epochs (e.g. hourly or per-block triggers). After each ML run, newly flagged addresses may be stored and optionally fed back into further training (semi-supervised learning) to improve detection over time.

Importantly, all heavy computation (feature extraction, model inference) occurs off-chain. The blockchain layer is used for secure data logging and enforcement, while ML runs on cloud or local servers. This hybrid design balances decentralization with practical compute needs.

4. Implementation Details

4.1. Data and Features

Our data comes from the blockchain ledger. For a DeFi lending example, transactions include token transfers, loan requests, and repayments. From these we derive features:

- **Transaction graph metrics:** Each address's degree, clustering coefficient, centrality.
- **Behavioral features:** Frequency of borrowing/lending, average loan size, age of account.
- **Token features:** Participation in new token offerings (which might signal exit scams).
- **Internal transactions:** Some fraud (e.g. hidden transfers) use *internal* Ethereum calls; extracting these (as Yuan 2024 did) can improve detection.

These features are normalized and fed into ML models. To address class imbalance (frauds are rare), techniques like SMOTE oversampling or cost-sensitive learning are applied, as done by Ashfaq *et al.* (2022). We also periodically update the dataset to include recent verified fraud cases (e.g. address lists from known scams) for retraining.

4.2. ML Models

We implement several models in parallel:

- **Ensemble Classifier:** A stacked ensemble combining XGBoost, Random Forest, and SVM. Gu & Dib (2025) report that an ensemble of Random Forest + XGBoost + SVM reached >98% accuracy on Ethereum fraud data. We adopt a similar strategy. Grid search and cross-validation tune hyperparameters for precision-recall balance (fraud detection needs high recall to catch frauds, and high precision to avoid false alarms).
- **Graph Neural Network:** We build a graph of addresses and train a GNN (e.g. GraphSAGE or a Transformer-based GNN) to predict illicit addresses. This captures complex money-flow patterns. If labeled data is insufficient, we incorporate unsupervised graph anomaly detection (e.g. graph autoencoders as suggested by Pérez-Cano & Jurado 2025).
- **Anomaly Scorer:** An isolation forest flags any transaction far from learned norms. Such anomalies are queued for review, even if no supervised label exists.

The models are trained on historical blockchain data (e.g. last 6 months of transactions). Feature importance analysis (like SHAP) helps understand key fraud indicators; for example, Yuan (2024) found that including internal transaction counts and account age dramatically improved detection.

4.3. Smart-Contract and Blockchain Integration

We implement one or more smart contracts to cooperate with the ML system. For instance:

- A **Monitoring Contract** that logs flagged addresses on-chain. This contract can enforce simple rules: for example, if a flagged address attempts a high-value transfer, the contract can revert the transaction (with a controlled "circuit breaker" mechanism).
- An **Incentive Mechanism Contract** to encourage nodes or oracles to submit up-to-date ML insights (similar to the approach of Pranto *et al.* 2022). Incentives (e.g. token rewards) motivate honest reporting of fraud findings in a decentralized way.

Blockchain governance (e.g. via DAO) could allow community review of flags. All on-chain smart contract code is open and audited to minimize new vulnerabilities.

5. Example Use Case: DeFi Lending Platform

Consider a DeFi lending protocol where users borrow/lend tokens. Fraud risks include "flash loan" attacks, fake collateral schemes, and liquidity drain (rug pulls on collateral). In our system: when a large flash loan is issued, the ML engine quickly analyzes preceding transactions of the borrower's address. The ensemble model might identify the address as part of a phishing network (high risk). Simultaneously, the GNN checks if the collateral token was rapidly traded in suspicious patterns. If a high risk is detected, the oracle triggers the smart contract to **pause** the loan execution or raise the collateral requirement.

In practice, deploying this on Ethereum (or a layer-2) would involve a monitoring node connected to the DeFi chain, and a secure server running the ML models. We expect latency on the order of seconds to minutes, depending on computation time. Compared to current practice (no real-time checks), this can substantially reduce successful frauds.

6. Discussion

Our architecture leverages the best of blockchain and ML. Blockchain's immutable ledger ensures data integrity, while ML's adaptive analytics catch evolving fraud schemes. Table 1 (above) and the abstract summary from Masud *et al.* (2024) confirm this synergy: "blockchain's immutability and transparency, along with ML's data-driven fraud detection, create a robust framework".

However, challenges remain. Scalability is a concern: real-time analysis of all DeFi transactions may be compute-intensive. We mitigate this by focusing on high-value or high-risk transactions (flagging everything would be too costly). We also rely on off-chain ML to avoid congesting the blockchain. Privacy is another issue: while transparency aids detection, it also exposes user data. Our design only uses pseudonymous blockchain data, and analysis results (suspicion scores) can be kept private to prevent doxxing.

Finally, as noted by Kayikci & Khoshgoftaar (2024), integrating diverse data (blockchain plus off-chain signals like social media sentiment) could further improve detection. Our architecture is extensible: we can add oracles that feed external data (e.g. reported hack alerts) into the ML models.

7. Conclusion

We have presented a detailed system architecture combining blockchain and machine learning for fraud detection in DeFi. By harnessing blockchain's trusted data and ML's analytical power, the system can detect and deter fraudulent activities that have plagued DeFi. Our conceptual evaluation, supported by state-of-art studies, suggests that ensemble and graph-based models can achieve high detection rates. The proposed integration of smart contracts ensures on-chain enforcement of anti-fraud policies. Future work will implement this architecture in a test environment and evaluate performance on real DeFi data, including sensitivity to false positives and detection latency. As DeFi grows, such hybrid architectures will be essential for maintaining security and user trust.

References

- Ashfaq, T., Khalid, R., Yahaya, A. S., Aslam, S., Azar, A. T., Alsafari, S., & Hameed, I. A. (2022). *A Machine Learning and Blockchain Based Efficient Fraud Detection Mechanism*. *Sensors*, 22(19), 7162.
- Gao, P., Li, Z., Zhou, D., & Zhang, L. (2024). *Reinforced Cost-Sensitive Graph Network for Detecting Fraud Leaders in Telecom Fraud*. *Journal of Electrical Systems*, 20(10s).
- Gu, Z., & Dib, O. (2025). *Enhancing fraud detection in the Ethereum blockchain using ensemble learning*. *PeerJ Computer Science*, 11:e2716.
- Kayikci, S., & Khoshgoftaar, T. M. (2024). *Blockchain meets machine learning: a survey*. *Journal of Big Data*, 11(9).
- Luo, B., Zhang, Z., Wang, Q., Ke, A., Lu, S., & He, B. (2023). *AI-powered Fraud Detection in Decentralized Finance: A Project Life Cycle Perspective*. arXiv preprint arXiv:2308.15992.
- Masud, S. B., Rana, M. M., Sohag, H. J., Shikder, F., Faraji, M. R., & Hasan, M. M. (2024). *Understanding the Financial Transaction Security through Blockchain and Machine Learning for Fraud Detection in Data Privacy and Security*. SSRN.
- Pérez-Cano, V., & Jurado, F. (2025). *Fraud Detection in Cryptocurrency Networks—An Exploration Using Anomaly Detection and Heterogeneous Graph Transformers*. *Future Internet*, 17(1), 44.
- Raskin, M., & Yermack, D. (2018). *Digital Currencies, Decentralized Ledgers, and the Future of Central Banking*. *NBER Working Paper*.
- Zhu, D., Ma, Y., & Liu, Y. (2024). *Learning for Cryptocurrency Fraud Detection: A Systematic Review*. *IEEE Access*, 12, 102219.
- Zhou, Y., Huang, Q., Luo, X., et al. (2022). *Phishing scam detection on Ethereum based on transaction network analysis*. *Transactions on Information Forensics and Security*, 17(7), 1758–1771.