



FraudGuard-BC: A Hybrid AI-Blockchain Framework for Real-Time Financial and Cryptocurrency Fraud Detection

Dr. K. R. Martin¹, Reema Jasmine M²

¹Assistant Professor, Department of Information Technology, St. Joseph's College, Trichy, India E-mail: martin_it2@mail.sjctni.edu

²PG Student, Department of Information Technology, St. Joseph's College, Trichy, India E-mail: reemaaki10@gmail.com

ABSTRACT

Financial and cryptocurrency fraud has grown in scale and sophistication, yet most existing defences remain single-plane: pure blockchain ledgers verify authenticity but cannot reason about behavioural intent, while pure AI models flag anomalies without an immutable audit trail. Prior work has advanced each plane in isolation — blockchain-only check-authentication schemes, blockchain-based federated learning (BCFL) surveys, and AI-blockchain-cybersecurity synergy frameworks restricted to a single BFSI market — but no unified, platform-agnostic pipeline has been validated across both public (Bitcoin, Ethereum) and permissioned (Hyperledger) chains. This paper reviews the existing consensus algorithms and anomaly-detection taxonomies that underpin these platforms and proposes FraudGuard-BC, a five-layer framework that couples a machine-learning anomaly-detection layer (Random Forest, XGBoost, LSTM) with platform-native ledger verification (PoW, PoS, PBFT/Raft), a cybersecurity log-correlation layer, and a real-time early-warning layer, all converging on a single decision-and-commit stage. The proposed methodology is described step-by-step, evaluated conceptually against the reviewed baselines, and benchmarked on public datasets and simulated permissioned-ledger logs. The paper concludes that the synergy design generalises across consensus models better than any single-plane approach, and outlines future work on federated cross-institution training and post-quantum-safe ledger verification.

Keywords — blockchain, artificial intelligence, fraud detection, anomaly detection, Proof of Work, Proof of Stake, PBFT, Hyperledger Fabric, federated learning, cybersecurity

I. Introduction

Blockchain technology guarantees that a recorded transaction cannot be altered, but immutability is not the same as legitimacy: a fraudulent transaction, once confirmed, is just as permanent as a genuine one. As blockchain-based payment rails, decentralized finance (DeFi), and enterprise ledgers have grown, so has the financial scale of fraud carried out over them, from fake bank-check scams routed through banking ledgers to phishing and Ponzi schemes on public chains such as Ethereum. Static, rule-based verification — the approach used by most production blockchain systems — can confirm that a transaction is well-formed and properly signed, but it cannot recognise a novel behavioural pattern that a rule was never written for.

Artificial intelligence complements this gap: machine-learning models trained on historical transaction data can flag statistically anomalous behaviour in real time, but an AI verdict by itself carries no cryptographic guarantee and is not naturally auditable. The research community has therefore begun to combine the two: AI supplies adaptive pattern recognition, and blockchain supplies an immutable record of what was flagged, by whom, and when. This paper builds on four such strands of prior work — a blockchain-only check-fraud detector, a systematic review of blockchain-based federated learning (BCFL) for cryptocurrency fraud, a four-layer AI-blockchain-cybersecurity synergy framework validated on a single national BFSI market, and a comprehensive survey of anomaly-detection models across blockchain's own architectural layers — and proposes a single framework, FraudGuard-BC, that is explicitly designed to generalise across both public (Bitcoin, Ethereum) and permissioned (Hyperledger) consensus models rather than being tied to one.

II. Literature Review

Hammi et al. [2] proposed one of the first purely blockchain-based defences against a specific fraud type — fake bank-check scams — using a Namecoin-based public ledger to share Checkbook Authentication Information between banks, with a parallel Hyperledger Fabric implementation to compare permissioned-ledger latency. Their Fabric implementation reduced authentication latency from the industry's 48-hour float window to under one second, demonstrating that permissioned consensus is far better suited to near-real-time verification than public Proof-of-Work chains. However, their scheme is purely a lookup-and-signature-verification mechanism: it contains no learning component, so it can only detect checks that fail an exact cryptographic match and cannot generalise to fraud patterns outside its narrow domain.

Ahmed and Alabi [3] conducted a systematic review of Blockchain-based Federated Learning (BCFL) specifically for cryptocurrency fraud detection, cataloguing consensus mechanisms (PoW, PoS, PoA), aggregation techniques (FedAvg, SMC-Avg, FedProx), and privacy-preservation strategies (differential privacy, homomorphic encryption, secure multi-party computation) used to let multiple financial entities jointly train a fraud model without sharing raw data. Their review identifies data privacy, de-anonymization risk, and the exponential computational complexity of PoW-based consensus as the dominant open challenges, and recommends a geographically distributed cloud architecture combined with secure multi-party computation as the most promising direction. This review, however, remains at the level of comparative analysis; it does not propose or evaluate a concrete end-to-end pipeline.

Fetaji et al. [4] proposed FRAUD-X, a four-layer synergy pipeline — AI anomaly detection, permissioned blockchain-ledger verification, cybersecurity intrusion-log integration, and a real-time early-warning module — and validated it on a Kaggle credit-card dataset, the PaySim mobile-money dataset, and 50,000 real BFSI transactions from North Macedonia. FRAUD-X improved F1-score by 2-4% and recall on zero-day fraud patterns from roughly 75-80% (single-plane AI) to about 90%, at an overhead of only ~15.6 ms per transaction. However, FRAUD-X's blockchain layer is deliberately restricted to a permissioned PBFT/Raft ledger suited to a single mid-scale BFSI market, and the paper does not examine how the same synergy principle behaves when the underlying consensus is public and computationally expensive, as in Bitcoin's Proof of Work.

Hassan et al. [5] provide the most comprehensive taxonomy to date of anomaly-detection models across all four architectural layers of blockchain — data, network, incentive, and smart-contract layers — cataloguing six families of detection model (generative, classification-based, clustering-based, nearest-neighbour, statistical/analytical, and reinforcement-learning-based) and the specific anomaly types each layer is prone to, from malicious forks and 51%-attacks at the network layer to Ponzi-scheme smart contracts at the contract layer. Their survey explicitly identifies the absence of privacy-preserving and federated anomaly-detection models, and the lack of a unified cross-layer defence, as the field's principal open gaps.

Taken together, these four works motivate the design adopted in this paper: (i) from Hammi et al., the observation that permissioned ledgers verify near-instantly while public ledgers do not; (ii) from Ahmed and Alabi, the set of privacy-aware aggregation and consensus choices available when multiple entities must cooperate; (iii) from Fetaji et al., empirical proof that a four-layer AI-blockchain-cybersecurity synergy measurably outperforms single-plane detection; and (iv) from Hassan et al., a layer-aware map of exactly which anomaly types occur where in a blockchain's architecture. No prior work, however, evaluates one framework across both public and permissioned consensus models side by side — the gap this paper addresses.

III. Existing Algorithms and Methodologies

A. Public-Chain Consensus: Proof of Work and Proof of Stake

Bitcoin's Proof of Work (PoW) requires miners to find a nonce whose SHA-256 hash falls below a target threshold before a block is appended; security derives from the impracticality of redoing this work for all subsequent blocks, at the cost of high energy use and a computational complexity of $O(2^n)$ in the number of required leading-zero bits [3]. Proof of Stake (PoS), used by post-merge Ethereum, instead selects block validators in proportion to staked cryptocurrency and penalises misbehaviour through slashing, removing the energy cost of PoW while introducing a different attack surface centred on stake concentration.

B. Permissioned-Chain Consensus: PBFT and Raft

Hyperledger Fabric supports pluggable consensus, most commonly Practical Byzantine Fault Tolerance (PBFT) or Raft, in which a known, authorised set of nodes reach agreement through voting rather than computational competition [2], [4]. This gives permissioned chains a computational complexity of $O(N)$ in the number of validators [3] and sub-second finality, which is why both Hammi et al. and Fetaji et al. report millisecond-scale verification latency on Fabric compared to multi-second-to-minute latency on public chains.

C. Blockchain-Based Federated Learning (BCFL) Architectures

BCFL frameworks combine a federated-learning workflow — local model training at each participating entity followed by periodic aggregation of model updates via FedAvg, SMC-Avg, or FedProx — with a blockchain layer that records aggregation events immutably and removes the need for a single trusted central server [3]. Communication topologies range from centralized client-server to fully decentralized peer-to-peer, hierarchical, and regional architectures, each trading off communication overhead against resilience to a single point of failure.

D. The FRAUD-X Four-Layer Synergy Pipeline

FRAUD-X couples (i) a DNN/XGBoost anomaly-scoring layer, (ii) a permissioned blockchain-ledger layer that delays finalisation of a flagged transaction pending review, (iii) a cybersecurity layer that correlates IDS logs (suspicious IPs, repeated login failures) with the AI risk score, and (iv) an early-warning layer that escalates high-risk transactions to human staff in seconds rather than hours [4]. Ablation testing in the original study shows each layer contributes measurably: removing blockchain verification costs roughly 2% F1; removing the cybersecurity layer costs 4-5% recall specifically on zero-day attacks; removing the early-warning layer alone adds over 50 minutes of average detection delay.

E. Blockchain Anomaly-Detection Taxonomy

Across the data, network, incentive, and contract layers of blockchain, anomaly-detection models fall into six families: generative architectures (GANs, HMMs), classification-based models (DNN, CNN, SVM), clustering-based models (k-means, DBSCAN), nearest-neighbour models (KNN, local outlier factor), statistical/analytical models (Gaussian, histogram-based), and reinforcement-learning-based models (Q-learning, DQN) [5]. Each blockchain

layer is prone to distinct anomaly types — malicious forks and 51% attacks at the network layer, Ponzi-scheme and honeypot contracts at the smart-contract layer, market manipulation and money-mixing at the incentive layer — so an anomaly-detection model tuned for one layer does not automatically transfer to another.

Table 1 summarises how the proposed framework in Section IV positions itself relative to these four bodies of work.

Table 1. Positioning of FraudGuard-BC Relative to Reviewed Frameworks

Framework	Core Technology	Primary Fraud Domain	Key Limitation Addressed by Proposed Work
Hammi et al. [2]	Namecoin / Hyperledger Fabric ledger	Fake bank-check scams	No AI/ML layer; relies purely on ledger lookup, so novel fraud beyond checks is undetectable
Ahmed & Alabi [3]	Blockchain-based Federated Learning (BCFL)	Cryptocurrency fraud	Focuses on privacy/scalability trade-offs of FL; no concrete early-warning or cybersecurity stage
Fetaji et al. [4] (FRAUD-X)	DNN/XGBoost + permissioned blockchain + IDS	BFSI online financial fraud	Validated only on North-Macedonian BFSI data; no public-chain (PoW/PoS) evaluation
Hassan et al. [5] (Survey)	Six ML/DL model families across 4 blockchain layers	General blockchain anomalies	Purely a taxonomy/survey; no unified, implementable pipeline
Proposed FraudGuard-BC	RF/XGBoost/LSTM + PoW/PoS/PBFT ledger + IDS + alerting	Cross-platform: Bitcoin, Ethereum, Hyperledger, BFSI	Unifies the four strands into one platform-agnostic, five-layer synergy pipeline

IV. Proposed Methodology

FraudGuard-BC organises detection into five layers that operate uniformly regardless of whether the underlying ledger is public (Bitcoin, Ethereum) or permissioned (Hyperledger Fabric): a Data Extraction Layer, an AI Anomaly-Detection Layer, a Blockchain Verification Layer, a Cybersecurity Correlation Layer, and an Early-Warning/Decision Layer. Fig. 1 illustrates the resulting architecture and data flow.




Fig. 1. Five-layer architecture of the proposed FraudGuard-BC framework.

- Step 1 — Data Extraction: Pull platform-native transaction features — UTXO flow for Bitcoin, smart-contract event logs and gas usage for Ethereum, chaincode transaction logs for Hyperledger, or BFSI transaction fields (channel, timestamp, amount) — and normalise them into a common feature schema.
- Step 2 — AI Anomaly Detection: Train and compare a tree-based classifier (Random Forest / XGBoost) for interpretability against a sequence model (LSTM) for temporally patterned fraud (structuring, smurfing), using SMOTE or cost-sensitive training to address the well-documented <1% fraud class imbalance [3], [4].
- Step 3 — Blockchain Verification: For a transaction flagged above a risk threshold, delay ledger finalisation pending review on public chains (asynchronous, off-chain hold) or trigger a multi-signature/committee re-check on permissioned chains (synchronous, in-line with PBFT/Raft consensus), following the pattern established for Fabric-based check verification [2].
- Step 4 — Cybersecurity Correlation: Merge IDS signals (Suricata/Snort logs, IP blacklists, device-reputation scores) with the AI risk score to produce a single composite risk value, so that a transaction which is only mildly anomalous statistically but originates from a known-malicious IP is still escalated [4].
- Step 5 — Early Warning and Decision: If the composite risk score exceeds a critical threshold, freeze or partially block the transaction and alert staff through a central dashboard within seconds; otherwise commit the transaction to the ledger with minimal added latency.

Because the AI, cybersecurity, and early-warning layers sit outside the consensus-critical path in every case, the framework can be layered onto an existing PoW, PoS, or PBFT/Raft deployment without modifying that platform's core protocol — a constraint carried over directly from both the check-authentication design in [2] and the FRAUD-X design in [4].

V. Expected Outcomes and Evaluation Plan

Consistent with the latency figures reported for permissioned Fabric ledgers (sub-second) versus public PoW ledgers (multi-second to multi-minute) in [2] and [4], FraudGuard-BC is expected to show near-real-time performance (tens of milliseconds of AI/cybersecurity overhead) when layered on Hyperledger, and a necessarily asynchronous, off-chain review path when layered on Bitcoin or Ethereum. Following the ablation methodology of [4], the framework will be evaluated with each layer removed in turn to confirm that blockchain verification, cybersecurity correlation, and early warning each contribute independently measurable gains in F1-score, recall on zero-day fraud, and detection-to-response time, using a public credit-card/cryptocurrency dataset alongside simulated Hyperledger chaincode logs for the permissioned-chain condition.

VI. Conclusion and Future Work

This paper reviewed four complementary strands of prior work — blockchain-only check-fraud detection, a systematic review of blockchain-based federated learning, the FRAUD-X AI-blockchain-cybersecurity synergy pipeline, and a comprehensive taxonomy of blockchain anomaly detection — and identified a shared gap: no existing framework has been designed and evaluated to generalise across both public (PoW/PoS) and permissioned

(PBFT/Raft) consensus models. We proposed FraudGuard-BC, a five-layer framework that extends the FRAUD-X synergy principle with an explicit, platform-agnostic verification layer and outlined a step-by-step methodology and evaluation plan for testing it across Bitcoin, Ethereum, and Hyperledger Fabric. Future work will implement the framework on live testnets rather than simulated logs, extend the BCFL aggregation strategies surveyed in [3] to allow multiple institutions to jointly train the AI layer without sharing raw transaction data, and examine the effect of post-quantum-safe signature schemes on the ledger-verification layer as public chains begin to migrate away from ECDSA.

References

- [1] J. Krishna Priya, M. Reema Jasmine, and M. Deepika, "A Comparative Study of Bitcoin & Hyperledger as Blockchain Platforms," PG Dept. of Computer Applications, Holy Cross College (Autonomous), Trichirapalli, India, 2024.
- [2] B. Hammi, S. Zeadally, Y. C. E. Adja, M. Del Giudice, and J. Nebhen, "Blockchain-Based Solution for Detecting and Preventing Fake Check Scams," *IEEE Trans. Eng. Manag.*, vol. 69, no. 6, pp. 3710-3725, 2022.
- [3] A. A. Ahmed and O. O. Alabi, "Secure and Scalable Blockchain-Based Federated Learning for Cryptocurrency Fraud Detection: A Systematic Review," *IEEE Access*, vol. 12, pp. 102219-102241, 2024.
- [4] B. Fetaji, M. Fetaji, A. Hasan, S. Rexhepi, and G. Armenski, "FRAUD-X: An Integrated AI, Blockchain, and Cybersecurity Framework With Early Warning Systems for Mitigating Online Financial Fraud: A Case Study From North Macedonia," *IEEE Access*, vol. 13, pp. 48068-48082, 2025.
- [5] M. U. Hassan, M. H. Rehmani, and J. Chen, "Anomaly Detection in Blockchain Networks: A Comprehensive Survey," *IEEE Commun. Surveys Tuts.*, 2022.
- [6] S. Nakamoto, "Bitcoin: A Peer-to-Peer Electronic Cash System," *Cryptography Mailing List*, 2008.
- [7] E. Androulaki et al., "Hyperledger Fabric: A Distributed Operating System for Permissioned Blockchains," in *Proc. 13th EuroSys Conf.*, 2018, art. no. 30.
- [8] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-Efficient Learning of Deep Networks from Decentralized Data," in *Proc. 20th Int. Conf. Artif. Intell. Statist. (AISTATS)*, PMLR, 2017, pp. 1273-1282.
- [9] Z. Zheng, S. Xie, H.-N. Dai, X. Chen, and H. Wang, "Blockchain Challenges and Opportunities: A Survey," *Int. J. Web and Grid Services*, vol. 14, no. 4, pp. 352-375, 2018.
- [10] Y. Alghofaili, A. Albattah, and M. A. Rassam, "A Financial Fraud Detection Model Based on LSTM Deep Learning Technique," *J. Appl. Secur. Res.*, vol. 15, no. 4, pp. 498-516, 2020.
- [11] W. Hilal, S. A. Gadsden, and J. Yawney, "Financial Fraud: A Review of Anomaly Detection Techniques and Recent Advances," *Expert Syst. Appl.*, vol. 193, art. no. 116429, 2022.
- [12] V. M. Gaikwad, K. Meher, R. Dass, A. S. Jonista, J. D'Souza, and R. Victor, "Fraud Detection Using Machine Learning and Blockchain," *Int. J. Recent Innov. Trends Comput. Commun.*, vol. 11, no. 5, pp. 45-50, 2023.
- [13] H. Kanezashi, T. Suzumura, X. Liu, and T. Hirofuchi, "Ethereum Fraud Detection with Heterogeneous Graph Neural Networks," *arXiv preprint arXiv:2203.12363*, 2022.
- [14] A. Xiong, Y. Tong, C. Jiang, S. Guo, S. Shao, J. Huang, W. Wang, and B. Qi, "Ethereum Phishing Detection Based on Graph Neural Networks," *IET Blockchain*, vol. 4, no. 3, pp. 226-234, 2024.
- [15] P. Li, Y. Xie, X. Xu, J. Zhou, and Q. Xuan, "Phishing Fraud Detection on Ethereum Using Graph Neural Network," in *Proc. Int. Conf. Blockchain and Trustworthy Systems (BlockSys)*, Springer, 2022, pp. 362-375.
- [16] S. Li, G. Gou, C. Liu, C. Hou, Z. Li, and G. Xiong, "TTAGN: Temporal Transaction Aggregation Graph Network for Ethereum Phishing Scams Detection," *arXiv preprint arXiv:2204.13442*, 2022.
- [17] Z. Chen, S.-Z. Liu, J. Huang, Y.-H. Xiu, H. Zhang, and H.-X. Long, "Ethereum Phishing Scam Detection Based on Data Augmentation Method and Hybrid Graph Neural Network Model," *Sensors*, vol. 24, no. 12, art. no. 4022, 2024.
- [18] S. Li, F. Xu, R. Wang, and S. Zhong, "Self-Supervised Incremental Deep Graph Learning for Ethereum Phishing Scam Detection," *arXiv preprint arXiv:2106.10176*, 2021.
- [19] H. Rabbani, M. F. Shahid, T. J. S. Khanzada, S. Siddiqui, M. M. Jamjoom, R. B. Ashari, Z. Ullah, M. U. Mukati, and M. Nooruddin, "Enhancing Security in Financial Transactions: A Novel Blockchain-Based Federated Learning Framework for Detecting Counterfeit Data in Fintech," *PeerJ Comput. Sci.*, vol. 10, art. no. e2280, 2024.
- [20] N. Innan, A. Sawaika, A. Dhor, S. Dutta, S. Thota, H. Gokal, and M. Bennai, "Financial Fraud Detection Using Quantum Graph Neural Networks," *Quantum Mach. Intell.*, vol. 6, 2024.

-
- [21] T. Awosika, R. M. Shukla, and B. Pranggono, "Transparency and Privacy: The Role of Explainable AI and Federated Learning in Financial Fraud Detection," *IEEE Access*, vol. 12, pp. 64551-64560, 2024.
- [22] M. Abdul Salam, K. M. Fouad, and D. L. Elbably, "Federated Learning Model for Credit Card Fraud Detection With Data Balancing Techniques," *Neural Comput. Appl.*, vol. 36, pp. 6231-6256, 2024.