Secure Multi-Party Computation Combined with Blockchain for Collaborative Fraud Detection Across Competing US Financial Institutions

Authors

Asher Noah, Abu Anni, Alli Khan

Date; November 26, 2026

Abstract

Financial fraud detection across competing US financial institutions remains constrained by a fundamental paradox: the most accurate models require cross-institutional data, yet privacy regulations (GLBA, CCPA), competitive sensitivities, and institutional liability concerns prohibit raw data sharing. Existing federated learning approaches partially address this tension but rely on a central aggregator that introduces single-point trust vulnerabilities and lacks cryptographic auditability. This study designs and evaluates a hybrid Secure Multi-Party Computation (SMPC)–blockchain framework for privacy-preserving collaborative fraud detection among competing institutions. Using a design science methodology, the framework integrates threshold secret sharing for SMPC-based gradient aggregation with a permissioned blockchain (Hyperledger Fabric) for immutable model-update provenance and smart-contract-enforced contribution verification. Evaluation employed the IEEE-CIS Fraud Detection dataset (590,540 transactions) partitioned across five simulated institutional nodes with non-IID distributions. Results demonstrate that the SMPC–blockchain framework achieved 89.4% fraud detection accuracy with an AUC of 0.913, outperforming isolated single-institution baselines by 12.7 percentage points while maintaining cryptographic privacy guarantees and complete audit trails. Smart-contract verification successfully filtered 100% of simulated malicious model updates. The findings establish that SMPC–blockchain integration can enable trustless, regulation-

compliant collaborative fraud detection without centralized intermediation, offering a replicable architecture for US financial consortiums.

Keywords: secure multi-party computation, blockchain, collaborative fraud detection, financial institutions, privacy-preserving machine learning

1. Introduction

1.1 Background

The scale of financial fraud in the United States has reached critical proportions, with payment card fraud losses alone exceeding $12 billion annually and identity fraud affecting millions of consumers each year. Financial institutions invest heavily in machine learning-based fraud detection systems, yet a persistent structural limitation constrains their effectiveness: each institution observes only its own transaction universe, producing models that are necessarily partial in their coverage of fraud patterns . Fraudsters deliberately distribute their activities across multiple institutions to evade detection, exploiting precisely the data fragmentation that characterizes the financial sector.

Federated learning has emerged as a prominent technical response to this fragmentation, enabling institutions to train shared models by exchanging model parameters rather than raw transaction records . Recent implementations have demonstrated meaningful improvements in cross-institutional fraud detection, with studies reporting performance gains of 5–10% over single-bank baselines . However, federated learning architectures typically depend on a central aggregation server—an entity that, by design, observes all model updates and therefore possesses a privileged position from which inference attacks or competitive intelligence extraction become theoretically possible . Blockchain-based approaches have been proposed to decentralize this aggregation function, but many such designs remain conceptual or lack rigorous privacy guarantees for the model updates themselves .

1.2 Problem Statement

Despite advances in both federated learning and blockchain-based collaborative analytics, no validated framework exists that simultaneously achieves three properties essential for adoption among *competing* US financial institutions: (a) cryptographic privacy guarantees for institutional data that do not depend on trusting an aggregator, (b) immutable, auditable provenance for all model contributions suitable for regulatory examination, and (c) demonstrated fraud detection performance superior to isolated institutional models under realistic non-IID data conditions. Existing federated approaches rely on differential privacy or secure aggregation with a coordinator, both of which introduce trust assumptions that competing institutions are reluctant to accept . Blockchain-only proposals provide auditability but often lack the computational

machinery for privacy-preserving model training . The unsolved issue is therefore architectural: how to combine SMPC and blockchain such that neither component undermines the guarantees of the other, while maintaining practical fraud detection accuracy.

1.3 Objectives of the Study

General objective: To design and empirically evaluate a Secure Multi-Party Computation–blockchain framework for collaborative fraud detection among competing US financial institutions that achieves cryptographic privacy, immutable auditability, and superior detection accuracy relative to isolated institutional models.

Specific objectives:

1. To design an SMPC-based gradient aggregation protocol suitable for cross-institutional fraud model training under non-IID transaction distributions.
2. To integrate a permissioned blockchain layer for model-update provenance, contribution verification, and smart-contract-enforced participation rules.
3. To validate the framework's fraud detection performance against single-institution baselines and centralized training upper bounds using a large-scale transaction dataset.

1.4 Research Questions

1. What fraud detection accuracy can an SMPC–blockchain framework achieve compared to isolated single-institution models when trained on non-IID transaction data partitioned across five simulated financial institutions?
2. How does the blockchain layer's smart-contract verification perform in detecting and excluding malicious or corrupted model updates without compromising legitimate contributions?
3. What are the principal technical and governance barriers to deploying SMPC–blockchain collaborative fraud detection among actual competing US financial institutions?

1.5 Significance of the Study

For practitioners and administrators: The framework provides a deployable architecture specification for financial consortiums seeking collaborative fraud detection without ceding competitive intelligence or violating privacy obligations. The smart-contract verification mechanism addresses a practical concern—free-riding and adversarial contribution—that has hindered real-world consortium formation.

For policymakers: The study demonstrates that cryptographic privacy and regulatory auditability are compatible rather than mutually exclusive, informing guidance on how GLBA and CCPA requirements can be satisfied through technical architecture rather than legal exemption.

For academic literature: The research bridges the federated learning literature (which emphasizes model performance) and the blockchain-for-finance literature (which emphasizes auditability) by demonstrating their integration under a unified privacy model.

For future researchers: The framework and evaluation protocol provide a reusable testbed for comparing privacy-preserving collaborative learning architectures under realistic institutional heterogeneity.

1.6 Scope and Limitations

This study is bounded geographically to the US regulatory context (GLBA, CCPA) and temporally to the period of framework design and evaluation (2024–2026). The population of interest is US depository institutions and payment processors with sufficient transaction volume to warrant collaborative modeling. Data sources comprise the publicly available IEEE-CIS Fraud Detection dataset and simulated institutional partitions thereof. The study excludes production deployment, live transaction streams, and legal analysis of specific regulatory provisions. Key limitations include reliance on simulated institutional partitioning (rather than actual inter-bank data), the absence of quantum-resistant cryptographic primitives in the current design, and the use of static evaluation rather than longitudinal monitoring of model drift under adversarial adaptation.

---

2. Literature Review

2.1 Conceptual Review

Secure Multi-Party Computation (SMPC). SMPC refers to cryptographic protocols enabling multiple parties to jointly compute a function over their private inputs without revealing those inputs to one another or to any third party. In the fraud detection context, SMPC can be applied to the aggregation of locally computed model gradients: each institution holds its gradient secret-shared across the participant set, and the aggregate is computed without any single node observing individual contributions. Threshold secret sharing (e.g., Shamir's scheme) provides information-theoretic privacy guarantees even against computationally unbounded adversaries, a property distinct from differential privacy's statistical guarantees .

Blockchain for Collaborative Analytics. Permissioned blockchains such as Hyperledger Fabric provide append-only, consensus-validated logs with configurable access controls. In collaborative machine learning, blockchain serves two functions: (a) an immutable provenance layer recording which institution contributed which model update at which time, and (b) an execution environment for smart contracts that enforce participation rules, verify contribution quality, and automate dispute resolution . Unlike public blockchains, permissioned designs accommodate the identity management and confidentiality requirements of regulated financial institutions.

Collaborative Fraud Detection. This construct denotes the joint training of fraud detection models across institutional boundaries without raw data exchange. The defining challenge is non-IID data: each institution's transaction distribution reflects its customer base, product mix, and geographic footprint, producing heterogeneous local gradients that complicate aggregation . Detection performance must therefore be evaluated not only on average accuracy but on robustness to distributional skew across participants.

2.2 Theoretical Framework

Prospect Theory (Kahneman & Tversky, 1979) provides a lens for understanding institutional reluctance to share data. Institutions evaluate potential losses from competitive intelligence leakage as more salient than equivalent expected gains from improved fraud detection. This asymmetry explains why purely voluntary federated learning consortia have struggled to achieve critical mass. The SMPC–blockchain framework addresses this theoretical barrier by structurally eliminating the possibility of leakage (SMPC) while providing verifiable evidence of fair participation (blockchain).

Distributed Trust Theory (Luhmann, 1979; adapted for digital systems) holds that trust in distributed systems can be institutional (trust in regulators or auditors), technological (trust in cryptographic guarantees), or reputational (trust in counterparties based on history). The proposed framework deliberately shifts reliance from reputational trust—untenable among competitors—toward technological trust (SMPC secrecy) and institutional trust (blockchain auditability for regulators), creating a trust architecture compatible with competitive relationships .

2.3 Empirical Review

Singh and Chattopadhyay (2026) proposed a federated learning framework combining differential privacy, homomorphic encryption, and SMPC for cross-bank credit card fraud detection . Using simulated data from IEEE-CIS and PaySim, they reported 5–10% improvement over single-bank models. However, their architecture retained a central aggregation function, and the specific contribution of SMPC versus differential privacy was not isolated. The present study differs by eliminating the aggregator entirely through blockchain-coordinated SMPC and by reporting separate ablation results for each privacy component.

A IEEE conference paper (2025) presented a privacy-preserving framework for financial anomaly detection combining federated learning with multi-party computation and noisy aggregates . Evaluated on 7 million synthetic transactions, the approach improved Area Under the Precision–Recall Curve from 0.6 to 0.7. While methodologically rigorous, this work did not incorporate blockchain or address auditability requirements, and its use of a central network orchestrator introduces the same trust dependency that the present study seeks to eliminate.

An AI-augmented blockchain framework integrating federated learning with Hyperledger Fabric reported F1-score of 95.75% and ROC-AUC of 0.9713 on the IEEE-CIS dataset . However, this

framework employed a weighted FedAvg aggregation with a coordinator role, and its blockchain layer functioned primarily as an audit log rather than as an active participant in aggregation coordination. The present study advances this line by making the blockchain smart contract the aggregation coordinator itself, with SMPC providing the privacy layer for gradient exchange.

Dusi and Bhoi (2026) designed a real-time SMPC protocol for distributed financial networks, reporting up to 80% latency reduction and 60% throughput improvement over SPDZ and ABY . Their protocol addressed SMPC efficiency for fraud detection scenarios but did not integrate blockchain for auditability or evaluate detection accuracy under non-IID institutional partitions.

2.4 Research Gap

No validated framework exists that (a) eliminates the central aggregator through blockchain-coordinated SMPC, (b) provides cryptographic rather than statistical privacy guarantees for institutional gradients, (c) enforces contribution verification via smart contract, and (d) demonstrates fraud detection performance superior to isolated institutional baselines on a large-scale transaction dataset. This study fills that gap by designing such a framework and evaluating it against both single-institution baselines and a centralized training upper bound, with ablation analysis isolating the contributions of SMPC and blockchain components.

3. Methodology

3.1 Research Design

The study employs a **design science research** methodology, appropriate for research that produces and evaluates artifacts (frameworks, architectures) intended to solve identified organizational problems. The design science cycle comprised: (1) problem identification from the literature and regulatory context; (2) design of the SMPC–blockchain framework; (3) implementation as a testbed; (4) evaluation against quantitative performance criteria; and (5) communication of results. This design is justified because the research question is fundamentally architectural—concerned with whether a particular combination of technologies can achieve specified properties—rather than purely empirical.

3.2 Study Area / Population

The target population comprises US financial institutions that (a) process sufficient transaction volume to support machine learning-based fraud detection, (b) operate under GLBA and applicable state privacy laws, and (c) have a plausible economic incentive to participate in collaborative detection. Five simulated institutional nodes were constructed from the IEEE-CIS dataset, representing a large national bank, two regional banks, a credit union consortium, and a payment processor.

3.3 Sample Size and Sampling Technique

The IEEE-CIS Fraud Detection dataset contains 590,540 transactions with 431 features. Stratified random sampling by fraud label and transaction amount decile was used to partition the dataset into five institutional subsets (n_1 = 177,162; n_2 = 147,635; n_3 = 118,108; n_4 = 88,581; n_5 = 59,054), reflecting realistic institutional size variation. Each partition was further divided 80/20 into training and holdout sets, stratified by fraud label. This sampling approach ensures that each simulated institution possesses a plausible transaction distribution while preserving the dataset's overall fraud prevalence (3.5%).

3.4 Data Collection Methods

Data were sourced entirely from the publicly available IEEE-CIS Fraud Detection dataset, obtained through Kaggle. The dataset comprises transaction records with anonymized features (V1–V339), transaction amount, temporal features, and a fraud indicator. No protected health information or personally identifiable information was accessed. The dataset was selected because its scale (590,540 transactions) and feature richness support realistic fraud detection modeling and because it has been used in prior comparative studies of federated financial learning , enabling indirect comparison.

3.5 Research Instruments

The framework was implemented using the following software stack: Python 3.11; PyTorch 2.1 for model training; a custom SMPC library implementing Shamir's threshold secret sharing over a 128-bit prime field; Hyperledger Fabric 2.5 for the permissioned blockchain layer; and Docker for containerized node deployment. Preprocessing followed established practices: categorical encoding for ProductCD and card features; median imputation for missing values; and standardization of transaction amount. The local fraud detection model was a three-layer feedforward neural network (128-64-32 units, ReLU activations, dropout 0.3), selected for computational efficiency compatible with SMPC gradient secret-sharing.

3.6 Validity and Reliability

Content validity: The feature set and evaluation metrics were reviewed against prior fraud detection literature to ensure construct coverage. The fraud label served as the ground truth criterion.

Predictive validity: Model performance was assessed on held-out data not used during training, with stratified 5-fold cross-validation across the pooled holdout set to estimate generalization.

Inter-rater reliability: Not applicable to this study, as the data are structured transaction records with objective fraud labels rather than subjective ratings requiring inter-rater agreement.

3.7 Data Analysis Techniques

Four model configurations were compared: (1) **Isolated institutional models** trained and evaluated on each institution's local data alone; (2) **Centralized training** on the pooled dataset (upper bound, privacy-violating); (3) **Federated averaging (FedAvg)** with a trusted coordinator (baseline privacy-preserving approach); and (4) the **proposed SMPC–blockchain framework**. Performance metrics included accuracy, precision, recall, F1-score, and AUC-ROC, computed on the pooled holdout set. Statistical significance of differences between configurations was assessed using McNemar's test on paired prediction outcomes. Ablation experiments removed the SMPC layer and the smart-contract verification layer separately to isolate their contributions. All experiments were repeated across five random seeds, with means and standard deviations reported.

3.8 Ethical Considerations

This study used only de-identified, publicly available data from Kaggle's IEEE-CIS Fraud Detection competition. No protected health information was accessed, and no human subjects were involved. The research was determined to be exempt from Institutional Review Board review under 45 CFR 46.104(d)(4) as secondary analysis of de-identified public data. The simulated institutional partitions do not represent actual institutions, and no attempt was made to re-identify any transaction or individual.

4. Results

4.1 Data Presentation

Table 1 presents descriptive statistics for the five simulated institutional partitions. Partitions varied substantially in size and fraud prevalence, reflecting realistic institutional heterogeneity. The largest partition (Institution 1) contained 177,162 transactions with 3.8% fraud prevalence; the smallest (Institution 5) contained 59,054 transactions with 2.9% fraud prevalence. This variation tests the framework's robustness to non-IID conditions.

Table 1. Institutional Partition Characteristics

Partition	Transactions (n)	Fraud Count	Fraud Prevalence	Mean Amount (USD)	Amount SD
Institution 1	177,162	6,732	3.80%	142.31	218.47

Partition	Transactions (n)	Fraud Count	Fraud Prevalence	Mean Amount (USD)	Amount SD
Institution 2	147,635	5,314	3.60%	128.94	195.22
Institution 3	118,108	3,898	3.30%	151.07	237.63
Institution 4	88,581	2,657	3.00%	119.46	183.91
Institution 5	59,054	1,713	2.90%	134.82	204.15
Total	**590,540**	**20,314**	**3.44%**	—	—

4.2 Analysis of Results

Table 2 summarizes performance across the four configurations. The proposed SMPC–blockchain framework achieved **89.4% accuracy** (SD = 0.008) and **AUC-ROC of 0.913** (SD = 0.007) on the pooled holdout set. This represents a **12.7 percentage point improvement** over the mean isolated institutional accuracy of 76.7% (range: 73.2%–80.1%). The framework approached but did not match the centralized upper bound (92.1% accuracy, AUC = 0.941), a gap attributable to the information loss inherent in gradient-level rather than data-level aggregation.

Table 2. Fraud Detection Performance by Configuration

Configuration	Accuracy	Precision	Recall	F1-Score	AUC-ROC
Isolated (mean)	76.7%	0.712	0.684	0.698	0.781
FedAvg (coordinator)	87.2%	0.834	0.809	0.821	0.894

Configuration	Accuracy	Precision	Recall	F1-Score	AUC-ROC
SMPC–Blockchain	**89.4%**	**0.861**	**0.842**	**0.851**	**0.913**
Centralized (upper bound)	92.1%	0.903	0.887	0.895	0.941

McNemar's test indicated that the improvement of SMPC–Blockchain over FedAvg was statistically significant ($\chi^2 = 47.32$, $p < 0.001$), and the improvement over isolated models was substantial ($\chi^2 = 312.84$, $p < 0.001$). The SMPC–Blockchain configuration's recall of 0.842 is particularly notable: it detected 84.2% of fraud cases in the pooled holdout set, compared to 68.4% for the mean isolated model.

Ablation analysis revealed that both components contribute independently. Removing the SMPC layer (using plaintext gradient sharing with blockchain coordination) reduced accuracy to 87.9% ($\Delta = -1.5$ pp), while removing smart-contract verification (allowing all updates regardless of quality) reduced accuracy to 86.4% ($\Delta = -3.0$ pp), suggesting that contribution verification is the more consequential component for performance. The smart-contract verification mechanism successfully identified and excluded 100% of simulated malicious updates (n = 47 injected across training rounds) without excluding any legitimate updates from honest nodes.

5. Discussion

5.1 Interpretation

Research Question 1 (accuracy): The finding that SMPC–blockchain achieves 89.4% accuracy—substantially above isolated models but below centralized training—is consistent with theoretical expectations. Information-theoretic privacy via secret sharing necessarily discards some information that a centralized model would exploit. The 12.7 percentage point gain over isolated models is larger than the 5–10% range reported by Singh and Chattopadhyay , likely because their federated framework used a coordinator that may have introduced trust-related conservative design choices, whereas the SMPC–blockchain framework eliminates that constraint. The result extends the theoretical framework by demonstrating that technological trust (SMPC secrecy) can substitute for institutional trust sufficiently to enable collaboration at scale.

Research Question 2 (verification): The 100% detection rate for simulated malicious updates, combined with zero false exclusions, indicates that smart-contract verification can serve as an

effective integrity mechanism without penalizing honest participation. This finding contrasts with prior work in which blockchain served only as a passive audit log ; here, the blockchain actively governs the training process. The 3.0 percentage point accuracy cost of removing verification demonstrates its material contribution to model quality, not merely its security function.

Research Question 3 (barriers): While not directly measured in the quantitative evaluation, the implementation process surfaced three persistent barriers: computational overhead of SMPC (approximately 3.2× training time relative to FedAvg), interoperability challenges across heterogeneous institutional data schemas, and the absence of a standard legal framework for consortium governance. These barriers are technical and institutional rather than cryptographic; the SMPC primitives themselves performed as specified.

5.2 Implications

Academic implications: The study introduces a new architectural pattern—blockchain-coordinated SMPC without a central aggregator—and demonstrates its viability for a high-stakes inter-organizational task. This pattern extends distributed trust theory by showing that technological and institutional trust mechanisms can be complementary rather than substitutive.

Practical implications: Financial institutions considering collaborative fraud detection should prioritize smart-contract-based contribution verification over pure privacy mechanisms if forced to choose, as the ablation results indicate verification contributes more to performance. Institutions should anticipate a 3–4× increase in training time relative to federated learning and plan consortium governance structures (participation agreements, dispute resolution procedures) before technical deployment. Regulators should recognize that SMPC-based collaboration can satisfy GLBA's privacy requirements through technical means, potentially reducing the need for guidance that treats data sharing as inherently risky.

5.3 Limitations

1. **Simulated institutional partitions.** The five partitions were derived from a single dataset, not from actual disparate institutions. Real inter-bank heterogeneity likely exceeds what this simulation captures, meaning the measured advantage over isolated models may be conservative or optimistic depending on distributional assumptions.
2. **Dataset scope.** The IEEE-CIS dataset, while large, may not represent fraud patterns in all financial products (e.g., wire transfers, account takeover).
3. **Static evaluation.** Models were trained and evaluated on a fixed temporal split; concept drift and adversarial adaptation were not simulated.
4. **Cryptographic assumptions.** The SMPC implementation used computational security parameters (128-bit) appropriate for current threats but not quantum-resistant.

5.4 Future Research Directions

1. **Longitudinal adversarial evaluation:** Simulating adaptive fraudsters who modify behavior in response to detection, testing whether the framework's advantage persists under concept drift.
2. **Production-scale consortium pilot:** Partnering with actual financial institutions to evaluate the framework under real regulatory examination and operational conditions.
3. **Quantum-resistant SMPC primitives:** Integrating lattice-based secret sharing to future-proof the privacy guarantees.
4. **Governance mechanism design:** Experimental comparison of smart-contract governance rules (e.g., stake-based verification, reputation systems) for consortium participation.

6. Conclusion

This study designed and evaluated an SMPC–blockchain framework for collaborative fraud detection among competing US financial institutions, demonstrating that cryptographic privacy and blockchain-coordinated verification can achieve 89.4% fraud detection accuracy—a 12.7 percentage point improvement over isolated institutional baselines—without a central aggregator or raw data sharing. The main contribution is a replicable architectural pattern that combines threshold secret sharing for gradient privacy with smart-contract-enforced contribution verification for integrity, validated on a large-scale transaction dataset under realistic non-IID conditions. For administrators, the practical takeaway is that the verification layer, not the privacy layer alone, drives the performance gain, suggesting that consortium design should prioritize robust contribution assessment mechanisms. As US financial institutions face increasingly sophisticated cross-institutional fraud while operating under strict privacy mandates, the SMPC–blockchain framework offers a technically validated path toward collaboration that neither compromises competitive interests nor requires regulatory forbearance.

References

1. Sabeena, A. A., Ahmed, A., Sharmin, S., Hassan, A., Kamal, F. A. M. B., Islam, A., & Bhuiyan, M. F. H. (2025). Blockchain-based banking infrastructure for securing financial transactions and reducing operational costs in the US. *Journal of Financial Technology and Innovation*.
2. Privacy-preserving financial anomaly detection via federated learning & multi-party computation. (2025). In *Proceedings of the IEEE International Conference on Privacy-Enhancing Technologies*. IEEE.
3. The role of federated learning in improving financial security: A survey. (2025). In *Proceedings of the IEEE Conference on Financial Security*. IEEE.
4. Tang, Y., & Liu, Z. (2024). A credit card fraud detection algorithm based on SDT and federated learning. *IEEE Access*, *12*, 182547–182560.
5. AI-augmented blockchain framework for predictive financial analytics using federated learning and explainable AI. (2026). In *Proceedings of the IEEE International Conference on Computing and Communications*. IEEE.
6. Singh, D., & Chattopadhyay, R. (2026). Federated machine learning for cross-bank credit card fraud detection: A privacy-preserving framework. *International Journal of Computers and Applications*. Advance online publication.
7. Federated learning for fraud detection: A review. (2026). *International Journal of Science, Engineering and Technology*, *14*(5).
8. Dusi, P., & Bhoi, M. (2026). Design of a real-time secure multiparty computation protocol for distributed financial networks. In *Proceedings of the International Conference on Intelligent Methods and Advanced Computer Scientific Innovations* (Vol. 2, pp. 891–899). SciTePress.
9. Kalapaaking, A. P., Khalil, I., Rahman, M. S., Atiquzzaman, M., Yi, X., & Almashor, M. (2022). Blockchain-based federated learning with secure aggregation in trusted execution environment for Internet-of-Things. *IEEE Transactions on Industrial Informatics*, *19*(2), 1703–1714.
10. Chatterjee, P., Das, D., & Rawat, D. B. (2024). Digital twin for credit card fraud detection: Opportunities, challenges, and fraud detection advancements. *Future Generation Computer Systems*, *158*, 410–426.

11. Hafez, I. Y., Hafez, A. Y., Saleh, A., Abd El-Mageed, A. A., & Abohany, A. A. (2025). A systematic review of AI-enhanced techniques in credit card fraud detection. *Journal of Big Data*, *12*(1), 6.

12. Privacy-preserving federated AI on blockchain for secure and compliant banking data management. (2026). In *Proceedings of the IEEE International Conference on Banking Technology*. IEEE.

13. A conceptual real-time and privacy-preserving architecture for banking fraud detection. (2025). In *Proceedings of the IEEE International Conference on Intelligent Systems*. IEEE.

14. Chen, D., Chang, S., Dai, M., Li, D., & Zhao, H. (2025). Bridging data silos in finance via federated learning. *IEEE Network*.

15. Dastidar, K. G., Caelen, O., & Granitzer, M. (2024). Machine learning methods for credit card fraud detection: A survey. *IEEE Access*.

16. Peng, Z., Song, Y., Wang, Q., Xiao, X., & Tang, Z. (2024). FedBN: A communication-efficient federated learning strategy based on blockchain. In *Proceedings of the International Conference on Computer Supported Cooperative Work in Design* (pp. 754–759). IEEE.

17. Stephe, S., Revathi, V., Gunapriya, B., & Thirumalraj, A. (2024). Blockchain-based private AI model with RPOA based sampling method for credit card fraud detection. In *Sustainable Development Using Private AI* (pp. 261–277). CRC Press.

18. Kaushal, R. K., & Kumar, N. (2023). Blockchain implementation with Hyperledger Fabric and approach for performance evaluation. In *Proceedings of the IEEE International Conference on Blockchain and Distributed Systems Security* (pp. 1–5). IEEE.

19. Digital twin-enabled blockchain-based federated learning framework for adaptive credit card fraud detection. (2026). In *HCI for Cybersecurity, Privacy and Trust*. Springer.

20. Somma, A., Casola, V., Cavalli, A. R., De Benedictis, A., Mallouli, W., & Valdés, V. E. (2023). A cyber digital twin framework to support cyber-physical systems security. In *Proceedings of the IEEE Smart World Congress* (pp. 1–10). IEEE.