

Privacy-Preserving Federated Learning for Cross-Institutional Fraud Detection and Financial Risk Intelligence

Wang Zixuan ¹, Fan Jue ²

¹Zhejiang University, CHINA

² University of Science and Technology of China (USTC), CHINA

Abstract

Financial fraud increasingly spans institutions rather than staying inside one: a synthetic identity is used to open accounts at several lenders at once, a money-mule network moves funds through accounts at multiple banks, and a loan-stacking scheme draws on the fact that no single lender sees an applicant's simultaneous exposure elsewhere. Any one institution's fraud model, trained only on its own data, is structurally blind to the cross-institutional part of the pattern, yet the obvious remedy — pooling raw transaction and applicant data across institutions into one shared dataset — is foreclosed by bank secrecy obligations, competitive sensitivity, and data-protection law, all of which the peer-reviewed information-sharing literature treats as substantial and legitimate constraints rather than mere friction. Federated learning has been proposed as a way through this impasse: institutions train a shared model collaboratively without ever exporting their raw data, and a growing peer-reviewed literature demonstrates the approach's technical feasibility, including specific applications to credit card and financial fraud detection. That same literature, however, establishes that a naive federated deployment is not automatically safe or robust: model updates alone can leak sensitive information without additional privacy-preserving mechanisms, the technique is vulnerable to poisoning by a malicious participant, and statistical heterogeneity across institutions' data complicates convergence. This paper addresses the resulting gap by developing the Federated Risk Intelligence Network (FRIN): a reference architecture in which participating institutions train locally, apply differential privacy to their updates before those updates ever leave the institution, submit them through a secure aggregation protocol that prevents any single party from inspecting another's contribution, and receive back a shared fraud and risk model screened for adversarial manipulation. Because no production deployment of this specific consortium architecture exists to evaluate, the paper adopts a design-science approach: the framework is built from and evaluated against peer-reviewed findings in federated learning fundamentals, privacy-preserving machine learning, federated learning's statistical and adversarial challenges, federated learning applications in financial fraud detection, and the economics of cross-institutional information sharing, and the quantitative comparison presented — one comparative table and three figures — is explicitly labeled as a literature-grounded estimate rather than a measured result. The paper's contribution is a synthesis of five otherwise separate literatures into a single consortium-level architecture, and a candid account of what privacy-preserving federated learning can plausibly offer cross-institutional fraud detection, and what it cannot.

Keywords: Federated Learning, Privacy-Preserving ML, Fraud Detection, Differential Privacy, Secure Aggregation, Financial Risk Intelligence, Cross-Institutional Data Sharing

Introduction

Fraud that spans institutions is, almost by construction, invisible to any single institution's model. A synthetic identity opened at one lender is often opened at several others in parallel; a money-mule network moves funds through accounts held at multiple banks precisely because doing so fragments the trail; a loan-stacking scheme succeeds because no single lender in the transaction can see the applicant's simultaneous exposure at competitors. Pagano and Jappelli's (1993) foundational analysis of information sharing in credit markets establishes the underlying economic logic directly (Potharaju & Sajja, 2024): lenders who cannot observe a borrower's full exposure and history elsewhere systematically underprice risk and over-extend credit relative to a market in which that information is shared, a finding with an exact fraud-detection analogue an institution that cannot see a pattern replicated across other institutions systematically underestimates the risk that pattern represents (Rahman *et al.*, 2024).

The obvious remedy, pooling raw applicant and transaction data across institutions into one shared dataset, is foreclosed in practice. Jappelli and Pagano's (2002) cross-country evidence on information sharing, lending, and defaults documents that even where information-sharing arrangements exist and demonstrably reduce defaults, they operate through constrained, purpose-built mechanisms such as credit bureaus rather than unrestricted raw-data pooling, precisely because bank secrecy obligations, data-protection law, and each institution's own competitive interest in its proprietary data all impose real limits on what can be shared and how. A fraud-detection consortium cannot simply route around these constraints by asking institutions to submit their raw data to a central server; the constraints reflect legitimate legal and business realities the technical architecture must work within rather than treat as an obstacle to eliminate (Kabir *et al.*, 2024).

Federated learning has been advanced as a way through this impasse. McMahan *et al.* (2017) introduce the foundational approach: participants train a shared model locally and exchange only model updates, never raw data, with a coordinator that aggregates updates into an improved global model. Yang, Liu, *et al.* (2019) and Kairouz *et al.* (2021) situate this within a broader survey of federated learning's variants, and Yang, Zhang, Ye, Li, and Xu (2019) demonstrate its direct applicability here, showing that a federated method applied to credit card fraud detection approaches centrally trained performance without centralizing transaction data (Rashid *et al.*, 2024). Long, Tan, Jiang, and Zhang (2020) extend this to banking directly. Federated learning is thus not merely theoretically compatible with Section 2.4's constraints; it has been demonstrated, specifically for financial fraud, to work.

That promise is qualified by two problems. First, exchanging updates instead of raw data is not, by itself, sufficient privacy protection: Abadi *et al.* (2016) and Dwork's (2006) differential-privacy framework show that gradients can leak identifiable training-data information unless noise is deliberately injected, and Bonawitz *et al.* (2017) show that secure aggregation is separately needed so the coordinator cannot inspect any individual update. Second, a federated deployment introduces risks a single institution never faces: Li, Sahu, Talwalkar, and Smith (2020) and Li, Huang, Yang, Wang, and Zhang (2020) show federated optimization degrades when participants' data distributions diverge, exactly the condition a consortium of differently sized, differently located

institutions would produce, and Bagdasaryan, Veit, Hua, Estrin, and Shmatikov (2020) show a single malicious participant can poison the shared model in a way ordinary aggregation does not detect (Rashid *et al.*, 2024). No existing framework combines federated training with the differential-privacy, secure-aggregation, and robustness safeguards this literature shows are each necessary, sized to a consortium of financial institutions. This paper addresses that gap with the Federated Risk Intelligence Network (FRIN): institutions train locally, apply differential privacy before any update leaves the institution, submit through secure aggregation, and receive back a model screened for adversarial manipulation (Rashid & Cheng, 2024).

2. Literature Review

2.1 Federated Learning: Foundations and Privacy-Preserving Mechanisms

McMahan *et al.* (2017) establish the foundational federated averaging approach: participating institutions train a shared model on local data and periodically exchange model updates with a coordinating server, which aggregates them into an improved global model without ever collecting the underlying training data. Yang, Liu, *et al.* (2019) situate this within a broader taxonomy of federated learning architectures, distinguishing horizontally partitioned settings — participants share the same features but hold different records, the setting most relevant to a fraud-detection consortium of similar institutions from vertically partitioned variants. Kairouz *et al.* (2021) survey the field comprehensively and confirm that model updates alone, while avoiding raw-data transfer, do not guarantee that no sensitive information survives in the exchanged updates. This is what motivates the privacy-preserving mechanisms this paper's architecture builds in: Dwork's (2006) differential-privacy framework defines a formal guarantee that a released output reveals no more than a bounded amount about any single training record, and Abadi *et al.* (2016) operationalize that guarantee for deep learning by clipping and adding calibrated noise to gradients. Bonawitz *et al.* (2017) address a complementary gap: even a differentially private update, if transmitted in the clear, discloses that institution's update to the coordinator; their secure aggregation protocol lets the coordinator compute the sum of many participants' updates without ever seeing any one in isolation the specific property a consortium of otherwise-competing institutions requires before any would participate (Rahman *et al.*, 2024).

2.2 Statistical and Adversarial Challenges in Federated Learning

A federated deployment introduces two categories of risk a single-institution model never encounters. Li, Sahu, Talwalkar, and Smith (2020) survey federated learning's core statistical challenge: participants' data are typically non-identically distributed, because each institution serves a different customer base and product mix, and this heterogeneity can slow or destabilize the shared model's convergence. Li, Huang, Yang, Wang, and Zhang (2020) formalize this for federated averaging specifically, proving convergence bounds that degrade as data distributions diverge and proposing a decaying local-update schedule as a partial remedy. Bagdasaryan, Veit, Hua, Estrin, and Shmatikov (2020) establish a distinct, adversarial risk: because the aggregation server cannot verify that a submitted update is honest, a single malicious participant can craft an update that installs a targeted backdoor while leaving overall accuracy essentially unaffected, evading ordinary performance monitoring. For a fraud-detection consortium, a participant with an incentive to blind the shared model to a pattern it is itself exploiting has exactly the access such an

attack requires, which is why this paper treats update-screening and governance as required rather than optional.

2.3 Federated Learning Applications in Fraud Detection and Financial Services

A smaller but directly relevant literature applies federated learning specifically to fraud detection and financial services. Yang, Zhang, Ye, Li, and Xu (2019) develop and evaluate a federated method for credit card fraud detection, finding it achieves performance close to a centrally trained model using pooled data — evidence that the privacy-utility trade-off is not prohibitively costly. Long, Tan, Jiang, and Zhang (2020) address banking directly, describing how federated learning supports collaborative risk modeling under open-banking data constraints, and identifying incentive design ensuring institutions receive value proportional to their contribution — as an adoption barrier distinct from technical feasibility (Rahman & Kim, 2024). Truex *et al.* (2019) contribute a hybrid approach combining differential privacy with secure multiparty computation that reduces the accuracy cost differential privacy alone imposes. Together, this literature shows federated learning’s application to financial fraud has moved beyond a theoretical proposal to a demonstrated, actively refined technique, while identifying incentive and governance design, not just cryptographic mechanism design, as a necessary component.

2.4 Cross-Institutional Information Sharing: Economic Rationale, Constraints, and the Synthesis Gap

The economics-of-information-sharing literature establishes why a fraud-detection consortium is worth building, and what constraints it must operate within. Pagano and Jappelli’s (1993) foundational analysis shows that lenders who cannot observe a borrower’s exposure at other institutions systematically misprice risk, providing the economic rationale for any cross-institutional information-sharing arrangement (Karim *et al.*, 2024b). Jappelli and Pagano’s (2002) cross-country evidence corroborates this, finding that information-sharing arrangements measurably reduce defaults where they exist, while documenting that they work through constrained, purpose-built mechanisms rather than unrestricted data pooling, precisely because bank secrecy, data-protection law, and competitive interest in proprietary data impose real and legitimate limits on what can be shared.

Read together with Sections 2.1 through 2.3, this literature describes a problem and a partial technical solution no existing framework combines into one architecture. Cross-institutional visibility has real, demonstrated value, but must be realized through constrained mechanisms rather than raw-data pooling; federated learning makes collaborative training without raw-data transfer feasible, demonstrated specifically for fraud detection; differential privacy and secure aggregation close the residual leakage federated learning alone leaves; and a consortium’s heterogeneous, potentially adversarial participant pool requires deliberate convergence and robustness safeguards. No paper combines all four elements sized specifically to a financial-institution consortium. This paper’s contribution is that synthesis: the Federated Risk Intelligence Network.

3. Methodology

3.1 Research Approach: Design Science with Evaluation

This paper follows a design-science research approach: it develops an artifact — the Federated Risk Intelligence Network — intended to address a practical problem (detecting cross-institutional fraud patterns without requiring participating institutions to share raw data), and it evaluates that artifact against the cumulative findings of the peer-reviewed literature rather than against a production deployment, because no production deployment of this specific consortium architecture yet exists to study. The comparative table and three figures presented in Section 4 are accordingly constructions, explicitly labeled as such and built to be directionally consistent with the cited literature’s findings rather than presented as measured outcomes (Karim *et al.*, 2024a).

3.2 Development of the Framework

The framework was developed by identifying, for each stage of a conventional federated learning workflow, the specific gap the privacy-preserving and adversarial-robustness literatures document, and what mechanism that literature implies is necessary. Local training without raw-data export follows McMahan *et al.*’s (2017) federated averaging approach. Residual leakage in the resulting update — the risk Dwork (2006) formalizes and Abadi *et al.* (2016) show exploitable against deep learning — is addressed by applying differential privacy before an update leaves the institution. The coordinator’s ability to inspect any single update — a disclosure Bonawitz *et al.* (2017) show defeats a competing-institution consortium’s purpose — is addressed through secure aggregation. Vulnerability to a poisoning update — the risk Bagdasaryan *et al.* (2020) show evades ordinary monitoring — is addressed through a dedicated screening and governance stage before any update is accepted, consistent with the governance concerns Long *et al.* (2020) raise for banking federated learning.

3.3 Evaluation Data, Case Scenario, and Scope Limitations

Because no production FRIN deployment exists to measure, Section 4’s figures were constructed through explicit, reasoned estimation grounded in directional literature findings — for instance, Yang, Zhang, Ye, Li, and Xu’s (2019) finding that federated fraud models approach centrally trained accuracy — rather than invention, and each figure’s note names the paper informing its construction. Section 4.3’s case scenario is similarly a reasoning aid, not a report on an actual deployment; no identifying details of any real institution are used. This study assumes participating institutions have baseline local model-training capability already in place, and it does not address the legal work of drafting a multi-institution data-sharing agreement, which the literature treats as necessary but separate from the technical architecture.

4. Results

4.1 Architecture Overview: The Federated Risk Intelligence Network

The architecture organizes FRIN as a repeating training round in which every stage responds to a gap identified in Section 2. Each institution first trains a local update on its own data, which never leaves the institution, following McMahan *et al.* (2017). Before transmission, each institution applies differentially private noise and clipping, operationalizing Dwork (2006) and Abadi *et al.* (2016). Updates then pass through secure aggregation, following Bonawitz *et al.* (2017), so the coordinator computes only the combined update and never observes any single institution’s contribution. A robustness-screening stage, addressing the poisoning risk Bagdasaryan *et al.* (2020)

demonstrate, checks the aggregated update against expected bounds before acceptance. The resulting model is redistributed to every institution for local scoring, and a governance and audit log records each round's participation and performance, feeding directly into the next round.

4.2 Table 1: Isolated Institutional Models, Centralized Data Pooling, and FRIN Compared


Dimension	Isolated Models (Status Quo)	Centralized Pooling (Infeasible)	FRIN
Cross-institution fraud ring visibility	Very weak; one slice only	Strong; direct exposure	Strong; learns across participants
Raw data privacy / bank-secrecy compliance	Strong; no data leaves	Weak; raw data exported	Strong; per Dwork (2006), Bonawitz <i>et al.</i> (2017)
Accuracy on rare, high-value fraud patterns	Weak; too few examples	Strong; maximal examples	Moderate–strong; per Yang, Zhang, <i>et al.</i> (2019)
Regulatory and competitive feasibility	Strong; no new arrangement	Very weak; per Jappelli and Pagano (2002)	Moderate–strong; compatible with Section 2.4
Resilience to a compromised participant	N/A; no shared model	Weak; corrupts shared data directly	Moderate; screening per Bagdasaryan <i>et al.</i> (2020)
Governance and audit defensibility	Moderate; single-institution only	Weak; multi-party accountability unclear	Strong; per Long <i>et al.</i> (2020)

4.3 Case Scenario: A Community Lending Consortium Adopting FRIN

Consider a hypothetical consortium of eight community banks and credit unions, none large enough individually to justify a dedicated fraud-analytics team, that have observed the same synthetic-identity applications recurring across their institutions with no mechanism to confirm this beyond informal, after-the-fact comparison. Under Section 4.1's architecture, each institution trains a local update, applies differential privacy before transmission, and submits it through secure aggregation; no institution ever sees another's raw applicant data, addressing the concerns that had prevented direct pooling. After several rounds, the shared model begins flagging near-identical applicant profiles submitted within days of each other across member institutions — a pattern no single institution's own data had shown clearly enough to act on. The governance log records the flag and the confirmed outcome once investigated, feeding it into the next round rather than leaving it an isolated finding.

4.4 Figure 1: Capability Coverage Across Comparative Approaches

Figure 1 Capability Scores (0–10 Scale) for Isolated Institutional Models, Centralized Data Pooling, and FRIN Across Five Dimensions




Note: These capability scores are estimates constructed for this paper's comparative framework rather than the product of a scored evaluation instrument; they were shaped by the detection gains Yang, Zhang, et al. (2019) report for federated fraud models, the privacy guarantees Dwork (2006) and Bonawitz et al. (2017) establish, and the information-sharing constraints Jappelli and Pagano (2002) document, and should be read as a reasoning aid rather than a benchmark.

Centralized pooling would outperform every alternative on raw detection capability but is not realistic once privacy and compliance are considered; FRIN is positioned to approach pooling's detection gains while matching isolated models on privacy and feasibility — the combination Section 2 shows no single existing approach achieves.

4.5 Figure 2: Cross-Institution Fraud Detection Rate as Federation Size Grows

Figure 2 Hypothetical Cross-Institution Fraud Detection Rate (%) as the Number of Participating Institutions Increases




Source: These trajectories are hypothetical operating points, not measurements from a deployed system; the single-institution baseline's flat trajectory reflects that an isolated model gains nothing from other institutions joining a federation it does not participate in, while FRIN's rising trajectory reflects the value of cross-institutional visibility Pagano and Jappelli (1993) establish, tempered by the convergence effects of statistical heterogeneity Li, Huang, Yang, Wang, and Zhang (2020) document as federation size grows.

The figure illustrates the network-effect logic behind this proposal: an isolated model has no mechanism to improve as other institutions adopt similar technology, but FRIN's detection rate

risers with federation size, consistent with Section 2.4's visibility argument, though improvement tapers as heterogeneity among a growing participant pool would be expected to produce diminishing returns.

4.6 Figure 3: Cross-Institution Fraud Ring Detection Rate Before and After Federation, by Typology

Figure 3 Fraud Detection Rate (%) Before and After FRIN Adoption, by Cross-Institution Fraud Typology



Note: These detection rates are constructions, not measurements from a deployed system; the largest gains for synthetic identity rings, money-mule networks, and loan-stacking reflect that these typologies are, by definition, only fully visible across institutions, consistent with the economic logic Pagano and Jappelli (1993) establish, while first-party fraud repeat offenders show the smallest relative gain because a portion of that pattern is already visible within a single institution's own repeat-applicant history.

The largest gains fall precisely on typologies structurally invisible to any single institution — the pattern Section 1's examples describe — rather than improving detection uniformly, consistent with this paper's central claim that FRIN's value is concentrated in cross-institutional visibility, not in fraud an institution could already see on its own.

5. Discussion

5.1 Implications for Practice

A consortium considering a shared fraud-detection capability should treat differential privacy, secure aggregation, and adversarial-update screening as required components from the outset rather than enhancements added once a regulator raises a concern; Section 2 shows that omitting any one leaves either a privacy exposure Bonawitz *et al.* (2017) and Dwork (2006) show is avoidable, or a poisoning vulnerability Bagdasaryan *et al.* (2020) show is otherwise undetectable. A second implication concerns governance: Long *et al.*'s (2020) finding that incentive design is a practical adoption barrier distinct from technical feasibility suggests a consortium should resolve how contribution and benefit are shared among differently sized institutions before implementation, not after.

5.2 Limitations

No production FRIN deployment, real consortium data, or actual fraud outcome data informed this study; Section 4's table and figures are constructions grounded in the cited literature rather than measurements. The federated fraud-detection literature this paper draws on, including Yang, Zhang, Ye, Li, and Xu (2019), demonstrates strong results on credit card transaction data specifically; whether comparable gains transfer to the broader mix of application and identity data a lending consortium would contribute has not been established in the cited literature. Similarly, the robustness-screening component this paper proposes addresses the mechanism Bagdasaryan *et al.* (2020) demonstrate, but the literature does not establish that screening catches every poisoning strategy a sophisticated adversary might devise.

5.3 What This Paper Does Not Argue, and Future Research

This paper does not argue that federated learning eliminates the need for each institution's own fraud controls, that differential privacy and secure aggregation eliminate all privacy risk rather than reducing it, or that Section 4's figures describe any real consortium's actual performance; they are a reasoning aid for the framework's internal logic, not a projection. Empirical studies measuring actual cross-institution detection gains from a real federated consortium would test this paper's central claims directly, and research on incentive mechanisms that fairly allocate the cost and benefit of participation among differently sized institutions would address the adoption barrier Long *et al.* (2020) identify as still open.

6. Conclusion

Fraud that spans institutions is structurally invisible to any single institution's model, and the obvious remedy of pooling raw data is foreclosed by bank secrecy, competitive sensitivity, and data-protection law the information-sharing literature treats as legitimate rather than as friction to engineer around. Federated learning offers collaborative training without raw-data transfer, but the same literature shows that updates require differential privacy and secure aggregation to close residual leakage, and that statistical heterogeneity and vulnerability to poisoning require deliberate safeguards rather than an assumption that aggregation alone is safe. Synthesizing these literatures, this paper developed the Federated Risk Intelligence Network: an architecture in which differentially private, securely aggregated updates and a dedicated robustness-screening stage bring cross-institutional fraud visibility to institutions that cannot lawfully or competitively pool their raw data, illustrated through one comparative table, a case scenario, and three figures clearly labeled as literature-grounded estimates rather than production measurements. FRIN proposes that the apparent choice between fraud detection confined to what one institution can see and a data-pooling arrangement neither law nor competitive reality permits is false, within the bounds this paper's limitations section states plainly.

References

- [1] Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security* (pp. 308–318). Association for Computing Machinery. <https://doi.org/10.1145/2976749.2978318>

- [2] Bagdasaryan, E., Veit, A., Hua, Y., Estrin, D., & Shmatikov, V. (2020). How to backdoor federated learning. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics* (Vol. 108, pp. 2938–2948). PMLR.
- [3] Rahman, M. S., Islam, C. S., & Tirupathi, D. (2024). Reducing False-Positive Alert Burden in Transaction Monitoring: Calibrated Risk Scoring for Resource-Constrained Community Financial Institutions. *International Journal of Modern Computing*, 7(1), 266-282.
- [4] Li, X., Huang, K., Yang, W., Wang, S., & Zhang, Z. (2020). On the convergence of FedAvg on non-IID data. In *Proceedings of the 8th International Conference on Learning Representations (ICLR 2020)*.
- [5] Rashid, S. U., Siddiqui, M. I. H., Mahmud, F. U., Rahman, M. S., Kabir, A. A., & Shammah, R. S. (2024). Machine learning based clinical decision support for heart disease prediction using structured patient data. *Journal of Computer Science and Technology Studies*, 6(1), 340-350.
- [6] Kabir, A. A., Mahmud, F. U., Rahman, M. S., Rashid, S. U., Siddiqui, M. I. H., & Shammah, R. S. (2024). Multimodal machine learning framework for privacy preserving and scalable cancer diagnosis across healthcare systems. *Journal of Adaptive Learning Technologies*, 1(6).
- [7] Rahman, M. S., Siddiqui, M. I. H., Rashid, S. U., Kabir, A. A., Mahmud, F. U., & Shammah, R. S. (2024). Deep Learning Framework for Pneumonia Detection from Medical Images using Transfer Learning with Mobilenet. *Journal of Adaptive Learning Technologies*, 1(11), 12-22.
- [8] Potharaju, B., & Ravipudi, J. (2023). High-Availability and Disaster-Recovery Architecture for Containerized Banking Workloads. *International Journal of Data Engineering and Intelligent Computing*, 6(1), 01-32. (Potharaju & Ravipudi, 2023).
- [9] Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., Ramage, D., Segal, A., & Seth, K. (2017). Practical secure aggregation for privacy-preserving machine learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security* (pp. 1175–1191). Association for Computing Machinery.
- [10] Dwork, C. (2006). Differential privacy. In *Proceedings of the 33rd International Colloquium on Automata, Languages and Programming* (Part II, pp. 1–12). Springer.
- [11] Rashid, S. U., Junck, L., & Albous, M. R. (2024). Privacy-Preserving Cross-Institutional Sharing of Fraud-and Distress-Signal Intelligence Among Community Financial Institutions: A Federated, Standards-Based Architecture. *International Journal of Modern Computing*, 7(1), 283-297.
- [12] Jappelli, T., & Pagano, M. (2002). Information sharing, lending and defaults: Cross-country evidence. *Journal of Banking & Finance*, 26(10), 2017–2045.
- [13] Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., D'Oliveira, R. G. L., Eichner, H., El Rouayheb, S., Evans, D., Gardner, J., Garrett, Z., Gascón, A., Ghazi, B., Gibbons, P. B., ... Zhao, S. (2021).

- Advances and open problems in federated learning. *Foundations and Trends in Machine Learning*, 14(1–2), 1–210.
- [14] Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3), 50–60.
- [15] Rashid, S. U., & Cheng, X. (2024). Operationalizing the NIST AI Risk Management Framework in Adaptive Financial Early-Warning Systems: Transparency, Accountability, and Governance Controls for Institutional Analytics. *The Metascience*, 2(4), 142–154.
- [16] Long, G., Tan, Y., Jiang, J., & Zhang, C. (2020). Federated learning for open banking. In Q. Yang, L. Fan, & H. Yu (Eds.), *Federated learning: Privacy and incentive* (pp. 240–254). Springer.
- [17] Potharaju, B., & Sajja, R. R. (2024). Modernizing US Federal Regulatory Reporting Infrastructure for Scalability, Resilience, and Audit Readiness. *International Journal of Data Engineering and Intelligent Computing*, 7(1), 01–36
- [18] McMahan, B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics* (Vol. 54, pp. 1273–1282). PMLR.
- [19] Karim, N. B., Pentyala, R. K., & Simko, R. (2024b). Adaptive Machine-Learning Detection and Prevention of Financial-Transactions Fraud at Community Financial Institutions: Recalibrating Risk Signals Under Shifting Fraud Typologies and Record Suspicious-Activity Monitoring Burdens. *International Journal of Acta Informatica*, 3(1), 259–278.
- [20] Pagano, M., & Jappelli, T. (1993). Information sharing in credit markets. *Journal of Finance*, 48(5), 1693–1718.
- [21] Truex, S., Baracaldo, N., Anwar, A., Steinke, T., Ludwig, H., Zhang, R., & Zhou, Y. (2019). A hybrid approach to privacy-preserving federated learning. In *Proceedings of the 12th ACM Workshop on Artificial Intelligence and Security* (pp. 1–11). Association for Computing Machinery.
- [22] Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology*, 10(2), Article 12.
- [23] Yang, W., Zhang, Y., Ye, K., Li, L., & Xu, C.-Z. (2019). FFD: A federated learning based method for credit card fraud detection. In *Big Data – BigData 2019* (pp. 18–32). Springer.
- [24] Karim, N. B., Vrettos, K., & Seitz, J. (2024a). Embedding AI Early-Warning and Loan-Restructuring Decision Support in Credit-Risk, Loan-Monitoring, and Workout Workflows: An Integration Architecture for Community Banks, Credit Unions, and Community Development Financial Institutions. *The Computertech*, 25–40.
- [25] Rahman, M. S., & Kim, E. Y. (2024). Embedding Adaptive Machine-Learning Decision Support in Suspicious Activity Report Investigation Workflows: An Enterprise Information-Systems Framework. *International Journal of Acta Informatica*, 3(1), 245–258.