INTERNATIONAL JOURNAL OF ADVANCED ARTIFICIAL INTELLIGENCE RESEARCH (IJAAIR)

eISSN: 3087-4297

Volume. 02, Issue. 11, pp. 32-38, November2025"

Integrated Real-Time Fraud Detection and Response: A Streaming Analytics Framework for Financial Transaction Security

Dr. Leila K. Moreno
School of Information Systems, University of Sydney, Australia

Article Received: 05/10/2025, Article Accepted: 06/11/2025, Article Published: 30/11/2025

ABSTRACT

This article develops a comprehensive and practically oriented research contribution that synthesizes streaming analytics, deep learning, and pragmatic systems engineering to create an integrated framework for real-time detection and response to financial transaction fraud. The work builds on contemporary technical reports and peer-reviewed research on streaming platforms (Kafka, Spark, Flink), real-time fraud detection models, anti-money-laundering (AML) machine learning, and operational best practices in financial services (Rajeshwari & Babu, 2016; Abakarim et al., 2018; Nicholls et al., 2021; Saxena & Gupta, 2017). The proposed Integrated Streaming Fraud Response Framework (ISFRF) is defined by four tightly coupled layers—Event Ingestion & Stream Processing, Low-Latency Triage & Authorization, Contextual Enrichment & Deep Analysis, and Forensic Evidence & Governance—and is accompanied by a suite of engineering patterns for feature computation in streaming contexts, drift-adaptive learning, adversarial resilience, privacy-preserving forensic commitments, and human-in-the-loop adjudication. Methodologically, the article employs a structured conceptual synthesis of the literature and derives implementation prescriptions, evaluation metrics adapted to streaming contexts, and a multi-stage empirical agenda including realistic sandbox pilots, adversarial red-team testing, and cross-institutional studies for network-level detection. The analysis emphasizes trade-offs—latency vs. complexity, explainability vs. predictive power, and privacy vs. auditability—and offers concrete design decisions (such as hybrid on-chain/off-chain commitments for audit integrity and tiered model pipelines) to reconcile these tensions. Finally, the work identifies critical gaps in extant knowledge—particularly the scarcity of longitudinal, production-scale evaluations and adversarial field studies—and proposes a prioritized research roadmap for practitioners and scholars engaged in building resilient, trustworthy transaction monitoring systems.

KEYWORDS

Real-time fraud detection; streaming analytics; Kafka; deep learning; anti-money laundering; forensic logging; adversarial robustness.

INTRODUCTION

Modern financial ecosystems are characterized by high-velocity digital transactions, diverse payment channels, and complex intermediation networks. This environment creates economic opportunity but simultaneously increases exposure to fraud, identity theft, money laundering, and other illicit financial activity (Federal Trade Commission, 2024; News, 2024). The scale and timeliness of contemporary fraud threats demand detection and response capabilities that operate at streaming speeds—able to inspect, score, and act on transactions in near real time—while retaining rigorous evidentiary trails for compliance and forensic inquiry (Rajeshwari & Babu, 2016; Martín Hernández, 2015). Advances in machine learning, particularly deep learning, provide powerful models for pattern recognition and anomaly detection but often introduce latency, interpretability, and governance challenges when deployed in authorization paths (Abakarim et al., 2018; Nicholls et al., 2021). Parallel maturation of streaming platforms (Kafka, Spark Streaming, Flink) has created the opportunity to operationalize low-latency analytics at scale, enabling architectures that can reconcile the need for speed with depth of analysis (Saxena & Gupta, 2017; Hivemind Technologies, 2024).

Policy and industry reports underscore the urgency: consumer losses from investment scams and impostor tactics have risen, and credit-card-related losses continue to create substantial recovery and reputational costs for issuers and merchants (Federal Trade Commission, 2024; Axios, 2024). National and regional statistics signal demographic and channel-specific trends—illustrating how millennials and online-first consumers experience distinct fraud vectors—suggesting that detection systems must be agile across use cases and customer segments (News, 2024). At the same time, anti-money-laundering detection research shows the efficacy of machine learning methods when combined with network and sequence analysis, yet highlights the difficulty of deploying such methods in streaming, high-throughput settings (Jullum et al., 2020; Cardoso et al., 2022).

This article is motivated by three interdependent goals: (1) to articulate a detailed, implementable framework for streaming, real-time fraud detection that balances operational constraints and detection quality; (2) to synthesize methodological and engineering patterns from contemporary literature into practical prescriptions for feature engineering, model life cycle management, and forensic logging; and (3) to identify a concrete empirical agenda—pilots, adversarial testing, and cross-institutional experiments—that addresses current evidence gaps. The resulting Integrated Streaming Fraud Response Framework (ISFRF) proposes explicit layering, dataflow templates, evaluation metrics tailored to streaming contexts, and governance practices that make real-time fraud detection viable in production financial systems.

METHODOLOGY

The research method is a structured conceptual synthesis and design-oriented systems analysis. Given the applied engineering and socio-technical nature of the problem, the methodology blends systematic literature assimilation with architectural design reasoning and operational constraint mapping. The process entails five interlinked steps:

Corpus Selection and Validation: A curated body of literature and authoritative technical sources was assembled, focusing on (a) real-time fraud detection methods and streaming analytics (Rajeshwari & Babu, 2016; Martín Hernández, 2015; Abakarim et al., 2018), (b) platforms for streaming data processing and engineering guidance (Saxena & Gupta, 2017; Hivemind Technologies, 2024), (c) AML and transaction-monitoring research emphasizing network analysis and deep learning (Jullum et al., 2020; Cardoso et al., 2022; Nicholls et al., 2021), and (d) policy and incident reports that frame the societal stakes (Federal Trade Commission, 2024; Axios, 2024). The corpus includes peer-reviewed articles, conference proceedings, industry reports, and recent credible technical blog posts and whitepapers to bridge the gap between academic insight and practitioner practice.

Thematic Extraction: Sources were read in detail with an eye to extracting recurring patterns, constraints, and solution motifs. Thematic categories included: latency budgets and authorization-path tolerances; feature types viable in streaming settings; model families and their online learning capacities; class-imbalance and drift-adaptation techniques; logging and evidence commitment patterns; and governance/operational practices for dispute resolution and regulatory reporting. Each theme was annotated with source-backed evidence.

Architectural Synthesis: The ISFRF architecture emerged through iterative mapping of themes to system design decisions—selecting platform primitives (e.g., Kafka partitioning semantics, stateful stream processors), modeling roles (fast linear models vs. deep graph/sequence models), and logging strategies (hybrid on-line/off-line with cryptographic commitments). Trade-off analyses were performed to explicitly identify where one design choice would necessitate compensating mechanisms elsewhere (e.g., asynchronous forensic committing to avoid latency penalties).

Operational Prescriptions: The synthesis produced concrete engineering patterns for feature computation (sliding-window aggregation, streaming graph sketches, online embedding updates), model lifecycle management (continuous monitoring, drift detection, periodic mini-batch retraining), evaluation metrics adapted to streaming contexts (time-to-detection distribution, cost-weighted utility), and operational governance (human-in-the-loop adjudication, retention and access controls for forensic logs).

Research Agenda Formulation: Based on gaps identified in the literature (limited production-scale evaluations, insufficient adversarial field studies, jurisdictional heterogeneity in evidentiary acceptance), a prioritized empirical agenda was developed recommending sandbox pilots, red-team adversarial experiments, cross-institutional prototype studies, and user-centered research on dispute-resolution acceptability.

This methodological approach is intentionally pragmatic: it privileges solutions that are implementable within existing banking infrastructure while remaining grounded in the empirical claims and technical limits documented in the referenced literature (Saxena & Gupta, 2017; Abakarim et al., 2018; Nicholls et al., 2021).

RESULTS

The principal outcome is the Integrated Streaming Fraud Response Framework (ISFRF). The following detailed exposition describes the architecture, component-level engineering patterns, modeling strategy, logging and forensic architecture, evaluation metrics, and deployment recommendations.

Architectural Overview: Four Coupled Layers

ISFRF is organized as four interdependent layers that map to operational responsibilities and latency constraints:

Event Ingestion & Stream Processing Layer: This foundational layer is responsible for reliable, ordered ingestion of transaction events and associated telemetry (device fingerprints, geolocation, merchant metadata). The layer exploits partitioning semantics (e.g., partition by account or card identifier) to preserve event locality and to support deterministic stateful processing (Saxena & Gupta, 2017). Exactly-once or effectively-once processing semantics are desirable to avoid duplication in stateful aggregates; engineering for idempotent transformations reduces the operational risk associated with consumer restarts.

Low-Latency Triage & Authorization Layer: Operations in the authorization path must meet tight latency budgets to preserve user experience and throughput in payment networks. ISFRF prescribes ultralightweight heuristics (velocity checks, merchant blacklists) and simple explainable models (regularized logistic regression, compact decision-tree ensembles) as the primary decision inputs at this layer (Rajeshwari & Babu, 2016). Actions include immediate declines for high-confidence fraud, step-up authentication (challenge prompts), or authorization with enhanced monitoring. The primary design constraint is that any added latency must be within acceptable authorization windows (sub-second ideally), and models must be thoroughly tested to maintain acceptably low false-positive rates.

Contextual Enrichment & Deep Analysis Layer: To improve detection fidelity, the architecture routes referenced events to enrichment services that augment transactions with history-driven aggregates, behavioral embeddings, device reputations, and cross-entity link indicators. Mid-path scoring engines leverage moderately complex models (boosted trees, compact neural nets) under relaxed latency constraints for soft policy enforcement and to triage events into human review workflows when ambiguity is high. Deep models—sequence models, graph neural networks, variational anomaly detectors—are executed in the background or on micro-batched windows to detect sophisticated patterns such as money-laundering sequences or cross-account collusion (Jullum et al., 2020; Cardoso et al., 2022).

Forensic Evidence & Governance Layer: For regulatory compliance and dispute resolution, systems must preserve immutable attestations of the decisions and supporting evidence. ISFRF recommends a hybrid approach: retain comprehensive, access-controlled raw logs in secure off-line storage with retention consistent with regulatory obligations, while persisting cryptographic digests or succinct evidence summaries to an append-only tamper-resistant service (which can be implemented via secure ledger services or hardened audit logs) to enable proof of integrity without exposing personal data on immutable public ledgers (Nguyen et al., 2020; Abbassi et al., 2023).

Key Engineering Patterns and Techniques

Partitioning and State Management: Partition transaction streams by account or card identifier to enable efficient, low-contention stateful aggregations for velocity and rolling-window features. Employ state stores that are backed by replicated, fault-tolerant storage to survive consumer crashes (Saxena & Gupta, 2017).

Windowing and Late-Arrival Handling: Use sliding windows with carefully chosen durations for velocity features (e.g., per-minute, hourly, 24-hour aggregates). Implement mechanisms for handling late-arriving events through watermarking and corrective replay to avoid spurious anomalies due to out-of-order data (Martín Hernández, 2015).

Online-Updateable Embeddings: Represent recent behavioral history with low-dimensional embeddings that can be updated incrementally to capture short-term drift without expensive reprocessing. These embeddings support fast similarity comparisons in the authorization path and are stabilized in the deep layer for richer pattern discovery (Nicholls et al., 2021).

Graph Sketches for Connectivity: Maintain approximate graph sketches or small neighborhood caches to detect emergent link patterns (shared devices, merchant clustering) that are indicative of collaborative fraud or laundering rings. Such sketches are memory-efficient approximations suitable for streaming settings (Zhang & Trubey, 2019; Cardoso et al., 2022).

Tiered Model Serving: Deploy a serving topology where the fastest models are collocated with the authorization microservices, while heavier models are served via dedicated model servers with autoscaling and asynchronous invocation. Use feature caching and nearline feature stores to minimize repeated computation (Saxena & Gupta, 2017).

Handling Class Imbalance and Concept Drift

Resampling and Cost-Sensitive Losses: Given the low prevalence of fraud, training must incorporate classbalancing techniques and cost-aware objective functions that weight false negatives appropriately relative to false positives (Abakarim et al., 2018). In streaming contexts, implement windowed resampling where recent fraud examples are emphasized to adapt to emergent schemes.

Online and Continual Learning: Use incremental learning algorithms or frequent mini-batch retraining using the latest labeled adjudications to maintain model relevancy. Monitor distributional changes in feature inputs and scores using statistical divergence monitoring and trigger retraining when drift exceeds tolerances (Nicholls et al., 2021; Jullum et al., 2020).

Human-in-the-Loop Adjudication and Active Learning: Integrate adjudicator feedback loops and active learning strategies that surface ambiguous cases for human labeling. This both improves model training data and embeds human judgment where automated systems risk high customer impact (Abakarim et al., 2018; Labanca et al., 2022).

Forensic Logging, Privacy, and Auditability

Hybrid Log Architecture: Retain raw, detailed logs in access-controlled object storage for forensic need and regulatory retention periods; concurrently persist cryptographic commitments (hash digests of log segments) to an immutable service to enable later verification that logged records were unaltered (Nguyen et al., 2020; Abbassi et al., 2023).

Privacy-Preserving Evidence Summaries: Where jurisdictional or policy constraints forbid long-term retention of personal data on immutable ledgers, store only evidence summaries or anonymized signal aggregates on tamper-resistant ledgers, enabling verification of integrity without exposing PII (Remmide et al., 2024).

Chain-of-Custody and Access Governance: Implement role-based access controls, detailed audit trails, and strict separation of duties to ensure that forensic logs are accessed only for legitimate investigatory or regulatory purposes (The Business Research Company, 2025).

Evaluation Metrics Tailored to Streaming Fraud Contexts

Time-to-Detection Distribution: Measure the latency from transaction event to first high-confidence flag and examine the distribution to ensure that the system meets operational SLAs for in-flight mitigation (Rajeshwari & Babu, 2016).

Cost-Weighted Utility: Compute a utility function that balances money saved from prevented fraud against costs of manual review, customer friction from false positives, and regulatory penalties; such evaluation reflects real-world operational trade-offs (Nicholls et al., 2021).

Drift Sensitivity and Stability Metrics: Quantify the rate at which score distributions change and the system's sensitivity to retraining cadence; use such metrics to set automated retraining and threshold recalibration policies (Jullum et al., 2020).

Adversarial Robustness Tests: Evaluate models under simulated adversarial manipulations—poisoning, mimicry, low-and-slow behavior—to understand worst-case performance and to harden defenses (Nicholls et al., 2021; Abbassi et al., 2023).

Deployment Recommendations and Operational Phasing

Incremental Rollout: Start with a shadow deployment that scores transactions without impacting authorizations to validate detection performance and false positive behavior under production traffic. Progressively enable low-risk automated mitigations (e.g., step-up authentication) before high-impact automated declines.

Monitoring and Alerting: Deploy comprehensive observability—latency, throughput, error rates, and model drift indicators—to ensure rapid detection of operational anomalies.

Cross-Institution Collaboration: For network-level fraud detection (e.g., laundering rings), explore privacy-preserving collaboration mechanisms—secure multiparty computation, federated learning, or anonymized graph sharing—balanced against legal constraints and competitive concerns (Cardoso et al., 2022; The Business Research Company, 2025).

DISCUSSION

The ISFRF synthesizes a wide range of technical techniques and organizational practices into a coherent pattern for operationalizing real-time fraud detection in modern banking and payments environments. This discussion elaborates theoretical implications, addresses counter-arguments, highlights limitations, and maps a research agenda.

Reconciling Latency with Model Sophistication

One of the most significant tensions in fraud detection systems is that of latency versus analytical depth. Deep sequence models and graph-based approaches capture temporal and relational patterns that are highly indicative of complex fraud but often require windowed aggregation and heavier computation (Zhang & Trubey, 2019; Cardoso et al., 2022). ISFRF reconciles this by decoupling the authorization decision from forensic analysis: fast, explainable models are used in the control loop, while heavy models run asynchronously on enriched windows to refine detection and improve future models. This architectural decision reflects operational reality—immediate user experience constraints trump perfect detection in the authorization moment (Saxena & Gupta, 2017).

Explainability and Accountability vs. Predictive Performance

Black-box models achieve strong predictive performance but hinder explainability, which is critical for customer disputes and regulatory scrutiny (Nicholls et al., 2021). The ISFRF proposes a tiered epistemic approach: require explainable rationales for any high-impact automated action and maintain deeper, possibly opaque models to support advisory requests and forensic synthesis where human adjudication is available. Moreover, the architecture embeds techniques to extract local explanations (feature attributions) even from complex models to support regulatory and customer-facing narratives.

Adversarial Dynamics and Defensive Posture

Financial fraud is an adversarial domain. Attackers adapt to detection strategies, performing low-and-slow operations or mimicking legitimate behavior to evade detection (Nicholls et al., 2021). ISFRF emphasizes continuous adversarial testing and ensemble diversity as defensive mechanisms. Furthermore, logging forensic commitments increases the cost of evidence tampering and supports post-incident attribution—both deterrence and accountability functions (Abbassi et al., 2023).

Privacy, Legal Regimes, and Cross-Border Evidence Management

The need to maintain forensic integrity collides with privacy rules and cross-border legal heterogeneity. Immutable logs or ledgers containing personal transaction data can create compliance risk. The hybrid logging and cryptographic commitment approach offers a practical compromise—enable verification of unaltered logs without exposing PII on immutable public platforms (Nguyen et al., 2020; Remmide et al., 2024). Nonetheless, the legal efficacy of cryptographic commitments as admissible evidence varies by jurisdiction and needs empirical and legal research to clarify admissibility standards.

Organizational and Cost Considerations

Implementing ISFRF requires cross-functional capabilities—real-time data engineering, MLops, security and compliance operations, and human adjudication teams. For smaller institutions, the required investment may be prohibitive; consortium models, managed services, and open standards can lower entry barriers but introduce dependency and governance trade-offs. The design therefore recommends phased adoption and exploration of collaborative models subject to robust privacy and governance safeguards (The Business Research Company, 2025).

Limitations and Gaps in the Evidence Base

The literature provides robust conceptual techniques but lacks extensive longitudinal, production-scale evaluations. Many published models are validated on historical datasets that do not fully capture adversarial adaptation, drift dynamics, or the human operational costs of false positives (Abakarim et al., 2018; Nicholls et al., 2021). In addition, there is limited cross-institutional research on network-level detection that respects legal privacy constraints. These empirical gaps motivate the research agenda described next.

Research Agenda and Priority Experiments

Sandbox Pilots with Live Traffic Shadowing: Deploy ISFRF components in shadow mode to capture production traffic behavior, allowing measurement of false positive impacts and time-to-detect distributions without customer impact.

Adversarial Red-Teaming: Conduct systematic attacks—poisoning, mimicry, low-and-slow—to quantify model degradation and to iterate defense strategies.

Cross-Institution Privacy-Preserving Collaboration Trials: Implement prototypes for federated graph analytics or privacy-preserving watchlist sharing to evaluate network-level detection effectiveness.

Legal-Technical Studies on Evidentiary Acceptance: Collaborate with legal scholars and regulators to evaluate the admissibility of cryptographic commitments, standards for tamper-evident logs, and acceptable retention policies.

Human Factors and Customer Experience Research: Study user responses to step-up friction and false-positive remediation to optimize communication, minimize churn, and ensure fair treatment in automated workflows.

CONCLUSION

This article presents the Integrated Streaming Fraud Response Framework (ISFRF), a design-oriented synthesis that maps streaming analytics primitives, tiered modeling strategies, hybrid forensic logging, and governance practices into a coherent blueprint for real-time transaction fraud detection and response. The architecture is pragmatic: it explicitly separates latency-sensitive authorization decisions from deeper asynchronous analysis; it prescribes hybrid forensic commitments to balance auditability and privacy; and it embeds human adjudication and active learning to maintain model relevance and fairness. While the ISFRF is informed by substantial research evidence, critical empirical work remains: longitudinal production deployments, adversarial resilience testing, and cross-jurisdictional legal analyses are necessary to validate and refine the framework.

Financial institutions and researchers should pursue coordinated pilot programs, share standardized evaluation metrics, and engage regulators early to design evidence-based retention and audit policies. As fraudsters adapt, the defensive posture must evolve: integrating streaming analytics with robust governance, privacy-aware forensics, and continuous adversarial testing offers the most promising path to resilient transaction security in the digital age.

REFERENCES

1. Federal Trade Commission. (2024). Investment Scams 2024. Retrieved from https://www.ftc.gov
2. Axios. (2024). Impostor Scam Trends 2024. Retrieved from https://www.axios.com
3. Axios. (2024). Fraud Losses by State in 2024. Retrieved from https://www.axios.com
4. News. (2024). Millennials and Credit Card Fraud in Australia 2024. Retrieved from https://www.news.com.au
5. Xu, J., Yang, T., Zhuang, S., Li, H., & Lu, W. (2024). AI-Based Financial Transaction Monitoring and Fraud Prevention with Behaviour Prediction. Applied and Computational Engineering, 67, 76-82.
6. Gartner Hype. (2024). Gartner Hype Cycle for Artificial Intelligence: A Comprehensive Guide. Retrieved from https://aicoach.co.za/2024-hype-cycle-for-artificial-intelligence/
7. Hebbar, K. S. (2025). AI-DRIVEN REAL-TIME FRAUD DETECTION USING KAFKA STREAMS IN FINTECH. International Journal of Applied Mathematics, 38(6s), 770-782.
8. Glassbox. (2023). Customer acquisition in banking: 10 proven strategies you should implement right now. Retrieved from https://www.glassbox.com/blog/customeracquisition-in-banking/
9. Kaushik Sathupadi et al. (2025). BankNet: Real-Time Big Data Analytics for Secure Internet Banking. Big Data and Cognitive Computing.
10. Kai Waehner. (2023). The State of Data Streaming for Financial Services. Retrieved from https://www.kai-waehner.de/blog/2023/04/04/the-state-of-data-streaming-for-financial-services-in-2023/
11. Noussair Fikri et al. (2019). An Adaptive and Real-Time Based Architecture for Financial Data Integration. Journal of Big Data, 6(1).
12. Tundis, A., Nemalikanti, S., & Mühlhäuser, M. (2021). Fighting organized crime by automatically detecting money laundering-related financial transactions. Proceedings of the 16th International Conference on Availability, Reliability and Security.
13. Jullum, M., Løland, A., Huseby, R. B., Ånonsen, G., & Lorentzen, J. (2020). Detecting money laundering transactions with machine learning. Journal of Money Laundering Control, 23(1), 173–186.
14. Zhang, Y., & Trubey, P. (2019). Machine learning and sampling scheme: an empirical study of money laundering detection. Computational Economics, 54, 1043–1063.
15. Labanca, D., Primerano, L., Markland-Montgomery, M., Polino, M., Carminati, M., & Zanero, S. (2022). Amaretto: an active learning framework for money laundering detection. IEEE Access, 10, 41720–41739.
16. García-Bedoya, O., Granados, O., & Cardozo Burgos, J. (2021). AI against money laundering networks: the Colombian case. Journal of Money Laundering Control, 24(1), 49–62.
17. Segovia-Vargas, M.-J., et al. (2021). Money laundering and terrorism financing detection using neural networks and an abnormality indicator. Expert Systems with Applications, 169, 114470.
18. Guevara, J., García-Bedoya, O., & Granados, O. (2020). Machine learning methodologies against money laundering in non-banking correspondents. In Applied Informatics: Third International Conference, ICAI 2020.
19. Larik, A. S., & Haider, S. (2011). Clustering based anomalous transaction reporting. Procedia Computer Science, 3, 606–610.
20. Alexandre, C., & Balsa, J. (2015). Client profiling for an anti-money laundering system. arXiv:1510.00878.
21. Sahin, Y., & Duman, E. (2011). Detecting credit card fraud by decision trees and support vector machines.
22. Tang, J., & Yin, J. (2005). Developing an intelligent data discriminating system of anti-money laundering based on SVM.
23. Cardoso, M., Saleiro, P., & Bizarro, P. (2022). LaundroGraph: self-supervised graph representation learning for anti-money laundering.

24. Khan, W., & Haroon, M. (2022). An efficient framework for anomaly detection in attributed social networks. International Journal of Information Technology, 14(6), 3069–3076.

25. Kashika, P., & Venkatapur, R. B. (2022). Automatic tracking of objects using improvised YOLOv3 algorithm and alarm human activities in case of anomalies. International Journal of Information Technology, 14(6), 2885–2891.

26. Iliyasu, A. S., & Deng, H. (2022). N-gan: a novel anomaly-based network intrusion detection with generative adversarial networks. International Journal of Information Technology, 14(7), 3365–3375.

27. Gómez, J. A., Arévalo, J., Paredes, R., & Nin, J. (2018). End-to-end neural network architecture for fraud scoring in card payments. Pattern Recognition Letters, 105, 175–181.

28. Remmide, M. A., Boumahdi, F., Ilhem, B., & Boustia, N. (2024). A privacy-preserving approach for detecting smishing attacks using federated deep learning. International Journal of Information Technology.

29. Abbassi, H., Abdellah, B., Mendili, S., & Youssef, G. (2023). End-to-end real-time architecture for fraud detection in online digital transactions. International Journal of Advanced Computer Science and Applications.

30. Alkhalili, M., Qutqut, M. H., & Almasalha, F. (2021). Investigation of applying machine learning for watch-list filtering in anti-money laundering. IEEE Access.

31. The Business Research Company. (Jan. 2025). Financial Services Market Definition. The Business Research Company Insight.