International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(10s) 2026 pp. 482-490
https://doi.org/10.70917/ijcisim-2026-3630
Article

Predictive RegTech: Real-Time AML Compliance in Instant Payment Networks using Graph Neural Networks (GNNs)

Abhinav Reddy Jutur

The State University of New York at New Paltz, USA

Abstract: This article presents a novel, first-of-its-kind predictive RegTech solution to address this challenge using machine learning methods. The rapid global adoption of Real-Time Payment Systems (RTPS) has created a significant "velocity gap" in regulatory compliance. While expanding financial accessibility, these systems introduce new vulnerabilities into existing AML frameworks. Static, rule-based systems and batch processing architectures cannot effectively counter money laundering in sub-second transaction environments. This limitation enables sophisticated activities such as digital layering and smurfing that move illicit financial flows across networks faster than regulatory systems can react. The core of our solution involves the use of graph neural networks (GNNs). This approach enables real-time, pre-settlement risk assessment, preventing illicit transactions before execution. Unlike traditional AML systems that evaluate transactions in isolation, this framework analyzes the entire transaction network to detect coordinated illicit behavior in real time. GNNs capture complex structures such as loops, funnels, and bridges that indicate illicit activity. To support efficient implementation, the framework integrates Event-Driven Architecture (EDA). The proposed architecture introduces the concept of the Zero-Knowledge Proof (ZKP) protocol layer in order to make risk-sharing possible in a secure manner across multiple institutions. This allows the banks to cooperate with each other in order to combat financial crimes while maintaining their data sovereignty. With predictive graph analytics, event-driven integration, and private cooperation, the proposed architecture enables proactive compliance in real-time payment environments, including real-time payment systems such as FedNow, against high-speed financial crime.

Keywords: RegTech, Anti-Money Laundering (AML), Graph Neural Networks (GNNs), Real-Time Payment Systems (RTPS), Financial Crime, Zero-Knowledge Proofs

1. Introduction & The Velocity Challenge

1.1 The Crisis of Speed in Compliance

The global financial ecosystem has undergone significant transformation through technologies such as the Real-Time Payment System (RTPS). Technological advances such as those made in India's UPI system and the USA's FedNow system have ensured that real-time transactions take place, and, therefore, any delay in payment settlement is avoided [11]. However, RTPS, apart from making things convenient, poses threats to financial security. There is an emerging threat that takes advantage of RTPS to facilitate rapid transactions involving different accounts. This is referred to as digital layering, where the origins of money become obscured as the money is transferred rapidly through a series of financial transactions. This creates a critical compliance failure, where illicit activity operates at the same speed as legitimate transactions, rendering traditional controls insufficient for real-time enforcement. This growing disparity between transaction speed and compliance capability introduces systemic risk, requiring advanced solutions to protect critical financial infrastructure operating under real-time constraints.

Aspect	Details	Implications
RTPS Adoption	Platforms like UPI (India) and FedNow (USA) enable instant payments	Expands financial inclusion and efficiency

Velocity Gap	Transactions occur in sub-second windows	Compliance systems struggle to match speed
Exploitation	Criminals use digital layering to obscure fund origins	Makes illicit flows harder to trace
Compliance Challenge	Criminal activity occurs at the same pace as legitimate commerce	Requires predictive, real-time detection methods

Table 1: Crisis of Speed in Compliance [11]

1.2 The Problem: Limitations of Legacy AML

Rule-based mechanisms dominate traditional AML processes. These mechanisms identify suspicious activities such as threshold breaches or interactions with sanctioned entities. While effective in slower environments, this model presents significant challenges in RTPS settings. The first issue is the high prevalence of false positives. Rule-based systems often generate alerts for transactions that are not illicit, overwhelming compliance teams with noise. A majority of alerts do not correspond to genuine threats, creating what is commonly described as a compliance bottleneck [1]. These inefficiencies consume resources that could otherwise support targeted investigations, delaying the identification of genuine threats. A further limitation lies in the inability to capture the relational nature of financial crime. These systems evaluate transactions in isolation, preventing detection of cyclical patterns across multiple accounts. As a result, patterns that emerge only through aggregated transaction relationships remain undetected. Financial crime is inherently relational, involving multiple interconnected accounts, entities, and recipients. This challenge is amplified by the rise of instant payment systems, where existing compliance mechanisms remain insufficient. The continued expansion of RTPS is expected to intensify this gap in the absence of more advanced detection strategies.

Limitation	Description	Consequence
Rule-Based Thresholds	Flags transactions exceeding fixed values or involving sanctioned entities	Effective in slower systems but inadequate for RTPS
False Positives	A significant majority of alerts are non-illicit	Creates compliance bottlenecks and wastes investigative resources
Lack of Relational Context	Transactions analyzed in isolation	Fails to detect cyclical flows across multiple accounts
Operational Inefficiency	Compliance teams overwhelmed by noise	Delays detection of genuine threats

Table 2: Limitations of Legacy AML [1, 6]

1.3 Scope of Study

This article presents a paradigm shift from reactive rule-based systems to graph-based predictive intelligence. The framework utilizes graph neural networks (GNNs) to model the relational topology of payment networks. Through graph representation of the financial ecosystem, GNNs reveal structures such as cycles, funnels, and bridges that remain undetected in linear models [2], [7]. This enables pre-transaction intervention, shifting anti-money laundering systems from reactive detection to predictive prevention in high-speed financial environments.

The scope is defined across three components. First, GNNs are applied directly to transactions to generate a risk score prior to completion. This enables pre-transaction fraud prevention rather than post-transaction detection [10]. Second, GNNs are integrated with event-driven architecture to ensure inference operates within RTPS latency constraints [3]. Finally, privacy-enhancing technologies such as zero-knowledge proofs support interorganizational collaboration without compromising data security [4]. By addressing technical, operational, and regulatory challenges, this framework enables predictive compliance in real-time payment environments. The integration of graph intelligence in AML systems represents a significant advancement in compliance capabilities [10]. This capability is critical for U.S. real-time payment infrastructure, including FedNow, where transaction finality requires risk evaluation prior to settlement. where transactions are irreversible and require pre-settlement risk evaluation [12].

Strengthening detection at this stage improves the resilience of financial infrastructure against high-speed fraud. This transition from entity-level monitoring to network-level intelligence represents a fundamental shift in AML system design.

2. Literature Review- From Tabular to Graph Data

2.1 The Geometry of Financial Crime

Financial data exhibits strong relational properties. Each transaction represents a relationship between entities such as accounts, vendors, and devices. In graph representation, a transaction is modeled as an edge connecting two nodes. Traditional machine learning methods such as Random Forest and XGBoost transform relational data into tabular formats for classification tasks. As a result, traditional models fail to capture coordinated behavior distributed across multiple entities. Illicit operations often remain undetected when transactions are evaluated in isolation. Examples include layering, where funds circulate across multiple accounts, and smurfing, where numerous accounts deposit funds into a central collector. Such structures are effectively identified using graph representations, where transactions are modeled as edges rather than tabular entries. Graph-based solutions are better suited to detect such patterns, as they preserve network topology [1], [6]. Relational analysis enables detection of deviations that remain unrecognized in non-relational models. This capability is particularly important in RTPS environments, where rapid activity requires timely detection. Unlike traditional systems that evaluate transactions in isolation, this framework analyzes the entire transaction network to identify coordinated patterns of illicit activity.

2.2 Evolution of Graph Neural Networks (GNNs)

Graph Neural Networks (GNNs) represent a significant advancement in learning from non-Euclidean data. Unlike classical neural networks operating on grid-structured data, GNNs directly model graph structures as input [2], [7]. A key feature of GNNs is message passing, where each node aggregates information from its neighbors to update its representation. This concept has several implications for analyzing financial transactions of bank accounts. A node is evaluated based not only on its own activity but also on the behavior of connected nodes. As a result, otherwise normal activity may be flagged as risky when linked to high-risk entities. GraphSAGE enables inductive graph learning, allowing generalization to unseen nodes [8]. Unlike earlier techniques requiring full graph storage, GraphSAGE generates embeddings for new nodes by sampling local neighborhoods. This capability is essential as new nodes continuously enter and exit RTPS networks. Mule accounts can be identified even without prior history through contextual neighborhood information. GNN-based systems demonstrate improved anomaly detection, reduced false positives, and scalability in financial networks [7], [10]. These capabilities support the use of GNNs in predictive compliance systems.

Feature	Detailed Explanation	AML Application	Strategic Benefit
Message Passing	Each node aggregates information from its neighbors, updating its representation iteratively. This allows the model to capture both local and global context within the graph.	Accounts are evaluated not only by their own activity but also by the behavior of connected accounts. Suspicious associations raise risk scores even if direct activity appears normal.	Enables relational intelligence, uncovering hidden laundering chains and improving detection accuracy.
Inductive Learning (GraphSAGE)	GraphSAGE generates embeddings for unseen nodes by sampling their local neighborhood, rather than requiring the entire graph in memory.	Detects mule accounts that appear and disappear quickly in RTPS environments, even without prior transaction history.	Provides scalability and adaptability in dynamic payment networks where new accounts are constantly created.
Relational Intelligence	GNNs preserve network topology, identifying subgraph structures such	Detects laundering techniques like layering (cycles), smurfing (funnels), and	Reduces false positives by focusing on suspicious relational patterns

	as cycles, funnels, and bridges.	cross-border intermediaries (bridges).	rather than isolated thresholds.
Scalability	GNNs can operate on large, heterogeneous, and temporal graphs, leveraging distributed computation and sampling strategies.	Supports RTPS environments processing thousands of transactions per second without performance degradation.	Ensures compliance systems remain resilient under high throughput, meeting latency requirements of <500ms.

Table 3: Graph Neural Networks for AML [2, 7, 8, 10]

2.3 RegTech and the RTPS Infrastructure

Regulatory technology (RegTech) has traditionally focused on post-transaction auditing. Log monitoring occurs after transaction completion to identify potential fraud. While this method supports compliance, it is insufficient in RTPS environments where suspicious activity may already be completed. This requires a shift from retrospective detection to proactive prevention of fraudulent activity. Event-driven architecture enables this transition by supporting real-time GNN inference and pre-transaction risk scoring [3], [9]. Integration of AI within financial services remains necessary, particularly under evolving regulatory constraints [11]. The combination of GNNs and event-driven architecture provides a scalable and technically feasible solution. Privacy-preserving protocols such as zero-knowledge proofs enable inter-institutional cooperation while maintaining data sovereignty [4]. Therefore, the RTPS protocol will call for the change in approach. Compliance must evolve from rule-based enforcement to intelligence-driven systems capable of real-time relational analysis to protect critical financial infrastructure.

3. Methodology-GNN Framework for AML

3.1 Graph Construction

In order to accurately model the relationships between entities that constitute a financial ecosystem, the payment network is formulated as a heterogeneous temporal graph. The use of this form of modeling allows us to incorporate various types of nodes, edges, and attributes into our model. For example, nodes can be used to denote accounts, merchants, and IP addresses, as well as device identifiers. The edges can depict relationships such as transactions, common login devices, or common beneficiaries, while attributes add context-specific knowledge to our graph such as transaction frequency, speed, and geographical variance. The inclusion of temporal attributes will allow the model to recognize the dynamics that are associated with time frames. Many financial crimes occur during a particular period in time due to rapid transfer that helps conceal the source of money. The inclusion of heterogeneity is necessary as financial criminals can use more than just accounts for their activities. In such cases, other entities such as merchants or devices can play a role that is not seen when analyzing tabular data [1][2], [6]. Therefore, a graph approach enables a predictive intelligence approach because it can enable a predictive intelligence capability by capturing relational structure within the data. This capability supports pre-transaction risk evaluation, which is essential for real-time fraud prevention. This representation enables continuous updating of node and edge states, allowing the system to reflect evolving transaction behavior in real time. Such dynamic modeling is essential for detecting high-speed laundering patterns in RTPS environments [23].

3.2 The GNN Architecture (GraphSAGE)

The architecture chosen for this framework is GraphSAGE, an algorithm designed for inductive representation learning on large graphs [8]. Unlike earlier models that required the entire graph to be loaded into memory, GraphSAGE generates embeddings for new nodes by sampling their local neighborhood. This capability is critical in RTPS environments, where new accounts are created and discarded rapidly. Mule accounts, often used for laundering, can be detected even if they have no prior history, because their neighborhood context reveals suspicious patterns. This inductive capability supports detection in previously unseen transaction scenarios, which is critical in environments where adversarial behavior continuously evolves.

GraphSAGE performs according to the scheme of sampling and aggregation. Every node gathers information about itself, as well as from its neighbor nodes, and then performs aggregation using various operations, such as averaging or pooling. As a result, the inductive nature of GraphSAGE helps build local and global context. For

applications to AML, this implies that the risk characteristics of an account not only depend on the transactions performed by it but also on its relationships with other accounts [7], [10]. The inductive character of GraphSAGE contributes to its scalability. Real-time payment processing systems operate with thousands of transactions every second, thus making the need for generating representations for new nodes crucial. This architecture enables scalable, real-time inference, making it suitable for deployment in high-throughput payment systems processing thousands of transactions per second [18].

3.3 Anomaly Detection Logic

The anomaly detection logic is designed to identify subgraph structures indicative of financial crime. These include cycles, funnels, and bridges. Cycles happen where funds circulate across multiple accounts before returning to the origin. This occurs quite often in cases of layering, whereby the funds are moved between various bank accounts to confuse their source. It is essential to detect cycles through tracking multi-hop connections, which can be accomplished by utilizing GNNs [1], [7]. The funnel consists of many small accounts being funneled into one large collector node. This happens mostly in cases of smurfing, which is the breaking down of illicit funds into small chunks to evade threshold limits. Funnel identification can be achieved via neighborhood analysis of collector nodes [9], [10].

Bridges are nodes that function as the only mediator in high-risk areas. It is important to detect such nodes since their presence links regions that would not be otherwise related and thus allows for money to be transferred from one zone into another. The task of finding bridges involves considering structural as well as temporal characteristics, since these types of nodes tend to be active for short periods of time [2], [4]. The rationale behind the anomaly detection algorithm lies in combining structural and temporal characteristics. This way, by teaching the model to distinguish between such subgraphs, we can provide risk assessments of transactions prior to their execution.This structure-aware detection enables identification of coordinated laundering strategies that remain undetectable in isolated transaction analysis. The resulting risk scoring supports pre-transaction intervention in real-time payment systems [13].

4. Result: Real-Time Integration via Event-Driven Architectures

4.1 The In-Stream Inference Pipeline

Real-time payment systems demand compliance checks that operate within extremely tight latency windows, often less than 500 milliseconds. Conventional compliance solutions based on batch processing or post-transaction auditing are insufficient. For real-time detection of fraud, the GNN must be integrated into the transaction process flow itself. Event-driven architecture enables this integration. This architecture captures transactions through brokers such as Apache Kafka, where each transaction is processed as an event. Processing is initiated at transaction time rather than after settlement. This ensures that high-risk transactions are intercepted before settlement, preventing irreversible financial loss. The transaction flow consists of three stages: event capture, feature hydration, and GNN inference. Features are retrieved through graph-based feature stores. Transaction history is retrieved within milliseconds, enabling relational awareness for the model [1], [6]. Upon hydration of features, the GNN executes inference on the transaction. The transaction gets evaluated based on the aggregation of data from adjacent nodes, looking for cycles, funnels, or bridges that could signal money laundering or smurfing operations. The output is a probability score representing the likelihood of the transaction being unlawful. In case the probability is above a predetermined threshold, the system triggers actions such as holding the transaction or notifying compliance systems [3], [9], [10].This in-stream processing enables pre-settlement intervention, which is essential for preventing fraud in high-speed payment environments.This capability is critical for systems such as FedNow, where real-time settlement requires immediate fraud prevention to protect national financial infrastructure [11], [12].

4.2 Technical Workflow

The event-driven pipeline consists of four sequential stages:

Event Triggering: A payment request is received as a pending event in the system. This ensures that the compliance check takes place before any money release.

Feature Hydration: Feature attributes, including transaction and relationship history, are retrieved from the graph database. This process takes only milliseconds for the GNN model to possess sufficient context information to infer accurately [2], [7].

Inference of GNN: It includes inferring whether the transaction is fraudulent through relational intelligence. This is achieved with the help of the GNN algorithm, where the transaction is ranked depending on its probability of being fraudulent according to any anomaly in the subgraph structure among the neighboring nodes [8].

Outputting Risk Scores: The broker executes the hold event if the risk score is above the defined threshold value. Otherwise, it allows for the payment to be executed if it falls below the threshold [9], [10]. This event-driven integration ensures a balance between efficiency and robustness in its design.

Stage	Function	Compliance Benefit
Event Capture	Payment request published as pending event	Ensures compliance checks before fund release
Feature Hydration	A graph feature store provides historical and relational context	Enables relational awareness within milliseconds
GNN Inference	The model calculates probability of illicit activity	Detects anomalies in real time
Actionable Output	The broker triggers a hold event if the risk score exceeds the threshold.	Prevents disbursement and alerts compliance officers
Scalability	Distributed graph databases and sharding	Handles 10,000+ TPS without degradation

Table 4: Event-Driven Integration Pipeline [3, 9, 10, 11]

4.3 Scalability Analysis

Scalability is a critical consideration in real-time compliance. RTPS networks process tens of thousands of transactions per second, requiring compliance systems to operate at this scale without degradation. Distributed graph databases such as Amazon Neptune and Neo4j provide the required infrastructure. These systems allow GNNs to query multi-hop relationships across large graphs without performance degradation under heavy loads [3], [9]. Performance is sustained through distributed graph storage and parallel computation as transaction volumes increase. Sharding distributes graph data across multiple servers, enabling parallel processing. These features prevent any bottlenecks and ensure efficiency of the compliance systems under high transaction volumes. Event-driven architecture enables dynamic allocation of computing resources. Computing resources are adjusted based on transaction volume to maintain efficiency. With the help of hardware acceleration, institutions can ensure efficiency under RTPS requirements [10], [11]. Thus, the implementation of GNNs in event-driven pipelines offers a scalable solution to the problem of velocity gaps in compliance processes [13]. The inflow of financial transactions requires proactive detection and prevention of any illegal actions in order to address the issues in real-time mode [14], [15].

5. Discussion: Privacy-Preserving Risk Sharing (ZKPs)

5.1 The Compliance Paradox: Privacy vs. Transparency

Financial institutions confront an irony when dealing with money laundering problems. For example, it is necessary to provide transparent mechanisms to be able to identify criminal activity that frequently involves many banks or jurisdictions and payment routes. At the same time, legal frameworks like GDPR and CCPA do not allow exchanging confidential information even beyond national boundaries. In fact, there is a compliance paradox as institutions have to cooperate while being limited by legislation in exchanging information for that purpose [4, 11]. Money laundering is inherently an inter-institutional problem since criminals are exploiting the space between different banks to transfer their finances from one account to another and thereby conceal their identity. If there are no ways of collaborating, then each organization will have access only to the part of this laundering chain that will lead to its successful completion [15].

5.2 Implementing Zero-Knowledge Proofs (ZKPs)

Zero-knowledge proofs provide a solution to this paradox. A ZKP allows one party to prove knowledge of a fact without revealing the underlying data. Applied to financial compliance, ZKPs enable institutions to share risk signals without exposing customer identities. The proposed framework introduces a ZKP-enabled federated risk layer.

When Bank A detects a high-risk pattern using its GNN, it does not disclose the customer's name or transaction details. Rather, the proof of high risk is put on a shared ledger for which a cryptographically verifiable proof exists [2], [3]. Thus, other institutions have proof that this piece of information is true without any need to see the actual data. In case the actor wants to create another account with Bank B, the bank checks the proof of high risk in the shared ledger. Bank B thus learns about the proof of high risk and hence that the actor has already been identified as a potential threat somewhere else but still does not need to look at the private information owned by Bank A. This method ensures proper collaboration between the institutions while keeping them within the boundaries of legal regulations [16].

Component	Description	Strategic Impact
Compliance Paradox	Institutions must collaborate to detect laundering, but privacy laws (GDPR/CCPA) restrict data sharing	Creates blind spots across banks and jurisdictions
ZKP Mechanism	Proof of high risk published to a shared ledger without exposing customer identities	Enables secure collaboration while maintaining compliance
Federated Risk Layer	Banks verify cryptographic proofs instead of raw data	Builds trust and ensures legal adherence
False Positive Reduction	Shared signals combined with GNN intelligence improve accuracy	Investigators focus on high-probability threats, reducing wasted effort
Strategic Implications	Regulators gain confidence, institutions lower compliance costs, customers face fewer delays	Supports secure cross-border cooperation and resilience

Table 5: Privacy-Preserving Risk Sharing via ZKPs [4, 11]

5.3 Reducing False Positives

False positives are a persistent challenge in AML systems. Moreover, rule-based models often trigger alerts for genuine transactions, causing compliance departments to be inundated with notifications and wasting time on dealing with actual threats. The proposed framework eliminates this problem by leveraging relational intelligence via GNN and privacy-preserving institutional cooperation through ZKP.

By taking into account the relational context, GNNs minimize the number of false alerts by determining whether a transaction is truly high-value or just another laundered asset [1], [6], [7]. In addition, ZKP enables different institutions to share information about potential risks, allowing for more accurate predictions. Thus, if a certain actor was marked as suspicious by several banks, the likelihood of illicit activity grows tremendously. Otherwise, if an actor received an alert at only one bank, the shared database helps avoid generating false positives [17], [18], [19]. In conclusion, due to the implementation of GNNs, the system becomes more intelligent and efficient when predicting suspicious behavior. At the same time, through the use of ZKPs, institutional cooperation ensures that no unnecessary information is leaked to third parties, and the false positive rate is minimized [20].

5.4 Strategic Implications

The adoption of ZKPs within AML systems has some strategic significance for both the regulators and organizations. Regulators will be confident that the organizations are working well together without breaking any privacy regulations. Organizations will incur fewer costs in terms of compliance, as false positives decrease the number of investigations required. For customers, the process becomes simpler because their legitimate transactions are unlikely to trigger any unwarranted flags. In addition, the federated risk layer provides a basis for international coordination. It is common practice for money laundering activities to occur in various jurisdictions. This makes coordination among different compliance frameworks difficult. With the use of ZKPs, there can be international cooperation while maintaining data sovereignty [9], [10], [22], [23].

6. Conclusion & Regulatory Recommendations

6.1 Summary of Findings

The proposed framework represents a novel advancement in predictive RegTech for real-time anti-money laundering. Through the integration of graph neural networks and event-driven architecture technologies, banks will be able to create a real-time defense system against fraud. Unlike previous systems based on fixed thresholds and auditing methods, predictive graph intelligence technology will make it possible for them to detect the laundering activities at an early stage. By implementing privacy protection measures, this will allow banks to work together without infringing on their rights. This capability is essential for safeguarding high-speed payment ecosystems such as FedNow from emerging forms of financial crime.

6.2 Recommendations for Financial Regulators

Regulatory authorities can have a key role in realizing this revolution. For one, they should enforce the use of graph-ready data formats, making sure that the data collected during transactions allows for relational analysis. Secondly, the creation of regulatory sandboxes will enable banks to conduct experiments using GNN-based AML solutions. Lastly, zero-knowledge proof systems ought to be introduced for the purpose of incentivizing international collaboration. These systems will assist in preventing fragmentation, facilitating widespread usage of GNNs, and enhancing international payment systems.

6.3 Final Remarks

Combating financial crime in real-time payment ecosystems requires novel, intelligence-driven solutions. While there have been many significant changes in recent years, the adoption of graph neural networks represents an absolute necessity in order to ensure the resilience of critical financial infrastructure against high-velocity fraud. This advancement strengthens the resilience of critical U.S. financial infrastructure against high-velocity financial crime.

References

1. Shaziya Islam et al., “Detecting Fraudulent Transactions for Different Patterns in Financial Networks Using Layer Weighted GCN,” Human-Centric Intelligent Systems, vol. 5, pp. 181–195, Apr. 2025. https://link.springer.com/article/10.1007/s44230-025-00097-3
2. Sarah Mohammed Alshehri, et al., “Systematic Review of Graph Neural Network for Malicious Attack Detection,” Information, vol. 16, no. 6, p. 470, Jun. 2025. https://www.mdpi.com/2078-2489/16/6/470?utm_source=chatgpt.com
3. Manojkumar Reddy Peddamallu, “Real-Time Fraud Detection Using Graph Neural Networks and Federated Learning,” International Journal of Emerging Research in Engineering and Technology, Pearl Blue Research Group, vol. 6, no. 3, pp. 133–135, Sep. 2025. https://ijeret.org/index.php/ijeret/article/view/304/287
4. Batool Lakzaei, Mostafa Haghir Chehreghani, and Alireza Bagheri, “Disinformation Detection Using Graph Neural Networks: A Survey,” Artificial Intelligence Review, vol. 57, article no. 52, Feb. 2024. https://link.springer.com/article/10.1007/s10462-024-10702-9?utm_source=chatgpt.com
5. Linh Nguyen, Marcel Boersma, and Erman Acar, “Detecting Fraud in Financial Networks: A Semi-supervised GNN Approach with Granger-Causal Explanations,” Explainable Artificial Intelligence Conference, pp. 330–353, Oct. 2025. https://link.springer.com/chapter/10.1007/978-3-032-08330-2_16
6. Ahmad Asiri and K. Somasundaram, “Graph Convolution Network for Fraud Detection in Bitcoin Transactions,” Scientific Reports, vol. 15, article no. 11076, Apr. 2025. https://www.nature.com/articles/s41598-025-95672-w
7. William L. Hamilton, Rex Ying, and Jure Leskovec, “Inductive Representation Learning on Large Graphs,” arXiv preprint arXiv:1706.02216, Sep. 2018. https://arxiv.org/abs/1706.02216
8. R. Renuga Devi, Joseph Emerson Raja, and Yeo Boon Chin, “Reinforcement Learning with Graph Neural Network (RL-GNN) Fusion for Real-Time Financial Fraud Detection: A Context-Aware Community Mining Approach,” Scientific Reports, vol. 15, article no. 42953, Dec. 2025. https://www.nature.com/articles/s41598-025-25200-3
9. Vinh Truong Hoang et al., “Detecting Anomalies in FinTech: A Graph Neural Network and Feature Selection Perspective,” Computers, Materials and Continua, vol. 86, no. 1, pp. 1–40, Dec. 2025. https://www.sciencedirect.com/org/science/article/pii/S1546221825010768
10. Darko B. Vuković, Senanu Dekpo-Adza, and Stefana Matović, “AI Integration in Financial Services: A Systematic Review of Trends and Regulatory Challenges,” Humanities and Social Sciences Communications, vol. 12, article no. 562, Apr. 2025. https://www.nature.com/articles/s41599-025-04850-8
11. Nana Dwemoh Benneh, “Digital Banking And Fintech Integration: Implications For Financial Sector Efficiency,” Journal of International Crisis and Risk Communication Research, vol. 7, no. S9, pp. 3594–3603, Sep. 2024. https://doi.org/10.63278/jicrcr.vi.3770
12. S. Y. Reddy, M. K. Babu, and R. Unnikrishnan, “Decoupling Cloud Security (DCS): A Framework for Data Sovereignty and Cross-Border Cloud Compliance,” Journal of Information Systems Engineering and Management, vol. 10, no. 4, 2025. https://doi.org/10.52783/jisem.v10i4.9317
13. Nana Dwemoh Benneh, “Fintech Adoption and Digital Banking Transformation: A Structural Analysis of Financial Institutions,” International Journal of Computational and Experimental Science and Engineering, vol. 11, no. 4, 2025. https://doi.org/10.22399/ijcesen.5178

14. Raghu Gollapudi, “Autonomous Multi-Zone Replication for Zero-Loss Settlement Systems,” International Journal of Computational and Experimental Science and Engineering, vol. 12, no. 1, Jan. 2026. https://doi.org/10.22399/ijcesen.4817
15. F. A. Quintero, “Reducing Production Time Without Compromising Quality: Optimization Strategies in High-End VFX Simulations,” Sarcouncil Journal of Engineering and Computer Sciences, vol. 3, no. 8, pp. 1–8, 2024.
16. F. A. Clottey, “Strategic Leadership and Cross-Functional Alignment: Impacts on Procurement Efficiency, Financial Oversight, and Business Performance,” Journal of Information Systems Engineering and Management, vol. 10, no. 62s, pp. 1435–1445, Nov. 2025. https://jisem-journal.com/index.php/journal/article/view/14726
17. D. Bajaj, V. Shah, and S. Kanaparthi, “A Practical Framework for Migrating Large Manual Test Suites to Automation Pipelines: An Industrial Case Study,” Journal of Information Systems Engineering and Management, vol. 9, no. 3, pp. 1–8, 2024.
18. F. A. Clottey, “Optimizing Business Growth Through Strategic Leadership: Evidence from Team Development, Supply Chain Management, and Operational Efficiency,” IPHO–Journal of Advance Research in Business Management and Accounting, vol. 2, no. 11, pp. 32–39, 2024. https://doi.org/10.5281/zenodo.19603152
19. P. Sunil Kumar, “Index-State Uncertainty for Trustworthy Enterprise Retrieval-Augmented Generation (RAG),” Journal of Information Systems Engineering and Management, vol. 10, no. 36s, pp. 1258–1271, 2025. https://doi.org/10.52783/jisem.v10i36s.15043
20. Nana Dwemoh Benneh, “Financing National Infrastructure: The Role of Sovereign Capital Mobilization Strategies,” International Journal of Computational and Experimental Science and Engineering, vol. 8, no. 3, 2022. https://doi.org/10.22399/ijcesen.5175
21. Y. Suthari and P. Mohan, “Cloud-Driven Machine Learning-Based Framework for Measuring Customer Experience in Digital Touch-Points,” in Proceedings of the 2025 IEEE 11th International Conference on Smart Instrumentation, Measurement and Applications (ICSIMA), IEEE, pp. 328–333, 2025. https://doi.org/10.1109/ICSIMA66552.2025.11233566
22. Fauzia A. Clottey, “Project Oversight and Client Relationship Management as Drivers of Sustainable Business Growth under Strategic Leadership Frameworks,” International Journal of Computational and Experimental Science and Engineering, vol. 9, no. 4, 2023. https://doi.org/10.22399/ijcesen.5146