Journal of Computer Science and Technology Studies
ISSN: 2709-104X
DOI: 10.32996/jcsts
Journal Homepage: www.al-kindipublisher.com/index.php/jcsts

JCSTS
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

| RESEARCH ARTICLE

Graph-Based Anomaly Detection in Trading and Payment Networks: A Cloud-Native Approach

Leela Krishna Yenigalla
Independent Researcher, USA
Corresponding Author: Leela Krishna Yenigalla, **E-mail**: leelakrishnayen@gmail.com

| ABSTRACT

Graph-based anomaly detection leverages network theory to transform financial crime prevention by representing transactions and entities as interconnected structures rather than isolated events. The integration of graph analytics with cloud-native architectures enables financial institutions to identify sophisticated criminal activities that deliberately distribute operations across multiple accounts to evade traditional detection methods. By modeling the entire ecosystem of relationships between entities and their transactions, these systems reveal suspicious patterns through both structural anomalies and behavioral deviations, dramatically improving detection accuracy while reducing false positives. Cloud-native implementations provide the scalability, performance, and resilience required to process massive transaction volumes in real-time across global financial networks. This architectural approach fundamentally changes how financial institutions conceptualize and combat financial crime, moving from reactive investigation toward proactive prevention through earlier pattern recognition and contextual understanding of suspicious activities.

| KEYWORDS

Financial Crime Detection, Graph Neural Networks, Cloud-native Architecture, Transaction Monitoring, Network-based Anomaly Detection

| ARTICLE INFORMATION

ACCEPTED: 01 August 2025 **PUBLISHED:** 8 September 2025 **DOI:** 10.32996/jcsts.2025.7.9.46

I. Introduction

Financial institutions face significant challenges in detecting sophisticated fraud within complex trading and payment ecosystems. As transaction volumes surge across digital channels, conventional detection systems struggle to identify coordinated criminal activities disguised within legitimate financial flows. Compliance teams must navigate a delicate balance between thorough monitoring and operational efficiency, particularly as financial crime techniques grow increasingly sophisticated. Recent industry analysis indicates that fraud investigation teams devote a disproportionate amount of resources to false positives, while truly suspicious activities often remain undetected until after significant financial damage has occurred [1].

Traditional rule-based detection systems have formed the backbone of financial crime prevention for decades. These systems rely on predefined thresholds and static parameters to flag potentially suspicious transactions. While effective against rudimentary fraud attempts, rule-based approaches demonstrate critical limitations when confronting modern financial crime networks. By examining transactions in isolation rather than considering broader relationship patterns, these systems frequently miss sophisticated schemes distributed across multiple accounts and entities. The rigidity of rule-based detection creates predictable boundaries that experienced criminals deliberately circumvent through carefully structured operations. Additionally,

maintaining effective rule sets requires continuous manual updates, creating substantial operational overhead while still leaving institutions vulnerable to emerging criminal methodologies [1].

Graph-based approaches transform financial crime detection by representing the entire ecosystem as an interconnected network where entities become nodes and transactions form the connecting edges. This perspective aligns naturally with how financial crimes actually operate – through networks of related accounts and coordinated transaction patterns designed to appear innocuous when viewed individually. By analyzing relationship structures and transaction patterns simultaneously, graph-based systems can identify suspicious clusters and anomalous connections that remain invisible to traditional methods. Research demonstrates that network-based anomaly detection significantly improves the identification of fraudulent activities by evaluating entity relationships and behavioral patterns across the entire financial ecosystem [2].

The convergence of graph analytics with cloud-native architectures creates powerful new capabilities for financial crime detection. Cloud platforms provide the elastic computing resources necessary to process massive financial networks in near real-time, enabling continuous monitoring of transaction patterns and entity relationships at unprecedented scale. Distributed processing capabilities inherent in cloud environments align perfectly with the computational demands of graph analysis, allowing institutions to maintain comprehensive visibility across entire transaction networks without performance degradation during peak processing periods [2].

Graph-based anomaly detection, implemented through cloud-native technologies, offers unprecedented capabilities for identifying sophisticated financial crimes by comprehensively modeling the network of relationships between entities and transactions. This integrated approach enables financial institutions to detect evolving criminal methodologies with greater precision and contextual awareness than previously possible with traditional systems.

II. Theoretical Foundations of Graph-Based Financial Crime Detection

Graph theory provides the mathematical foundation for modern financial crime detection systems, offering a natural framework for modeling complex relationships within financial networks. Financial transactions and entities are represented as a graph G = (V, E), where vertices (V) represent entities such as accounts, individuals, or organizations, while edges (E) represent transactions or relationships between them. This representation captures the multi-dimensional nature of financial activities, where connections between entities often reveal patterns invisible when analyzing individual transactions in isolation. The inherent ability of graph structures to model relationships makes them particularly suitable for uncovering the interconnected nature of sophisticated financial crimes that deliberately distribute activities across multiple accounts to avoid detection [3].

In financial network analysis, nodes and edges serve as fundamental building blocks representing financial entities and their interactions. Nodes typically symbolize accounts, customers, or institutions, each characterized by attributes including transaction history, behavioral patterns, and demographic information. Edges capture financial flows between entities, carrying properties such as amount, timestamp, frequency, and transaction type. This dual representation allows detection systems to simultaneously analyze both network structure and entity attributes when identifying suspicious patterns. Graph-based detection leverages this rich representation to recognize anomalies through both unusual connection structures and atypical transaction characteristics, providing investigators with a comprehensive view of potential criminal activities [3].

Graph properties and metrics offer powerful analytical tools for uncovering suspicious patterns within financial networks. Centrality measures identify influential nodes that may function as coordinators or money laundering hubs, while clustering coefficients reveal unusually dense connection patterns indicative of collusion. Path analysis examines fund flows, identifying circular transaction chains characteristic of layering techniques. Temporal analysis of graph evolution captures suspicious structural changes over time, such as sudden formation of new connections or dramatic shifts in transaction patterns. Research demonstrates that different financial crime types exhibit distinctive topological signatures, with fraud networks typically showing specific structural characteristics that differentiate them from legitimate transaction networks [4].

Graph Neural Networks (GNNs) represent a significant advancement in financial crime detection, combining graph expressiveness with neural network learning capabilities. These models learn representations by aggregating information from neighboring nodes through multiple layers, capturing both local transaction patterns and global network structures. This neighborhood aggregation process enables GNNs to incorporate information from multiple hops away, mirroring how financial criminals distribute activities across intermediaries. Systematic literature reviews indicate that GNN-based approaches consistently outperform traditional methods in detecting coordinated fraud schemes that operate across multiple seemingly unrelated accounts [4].

Vector embeddings transform discrete financial networks into continuous vector representations that preserve relational properties while enabling efficient analysis. These embeddings map entities into high-dimensional space where proximity reflects

similarity in behavior or risk profiles. Advanced embedding techniques capture both structural roles and attribute-based characteristics, creating rich representations for pattern recognition. Once embedded, traditional machine learning methods can identify clusters or outliers indicating suspicious activities, bridging the gap between complex network structures and computational efficiency [3].

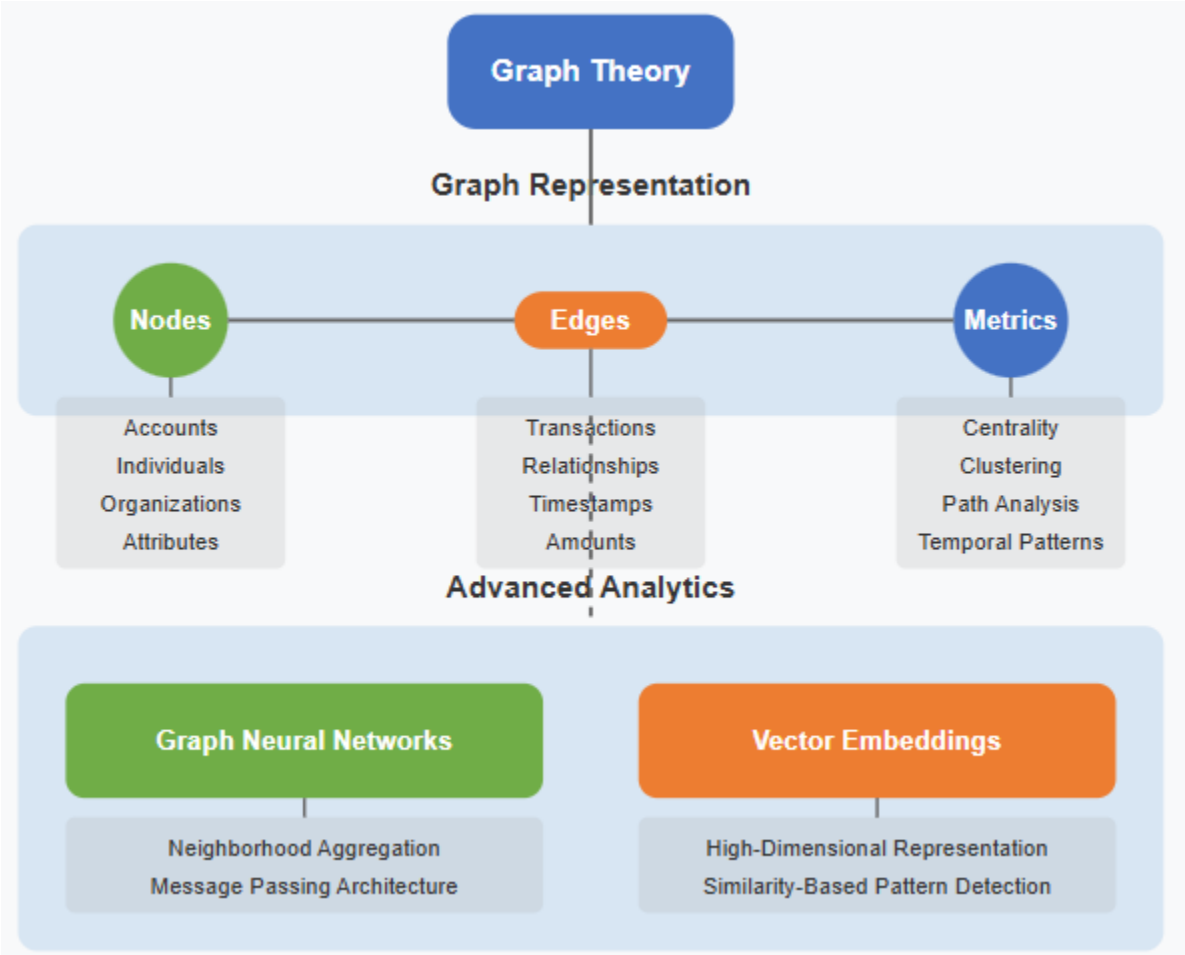

Fig 1: Theoretical Foundations of Graph-Based Financial Crime Detection [3, 4]

III. Cloud-Native Architecture for Graph-Based Anomaly Detection

Scalable graph database technologies form the foundation of modern financial crime detection systems, providing specialized storage and query capabilities optimized for relationship-based analysis. Cloud-native graph databases employ purpose-built data structures and indexing mechanisms that treat relationships as first-class citizens rather than derived attributes, dramatically accelerating the traversal operations essential for identifying suspicious transaction patterns. These databases maintain consistent performance even as financial networks grow to encompass billions of nodes and edges through sophisticated partitioning schemes that distribute data across commodity hardware while preserving locality for common query patterns. The distributed architecture enables horizontal scaling to accommodate growing transaction volumes without sacrificing query performance, a critical requirement for financial institutions monitoring global payment networks [5]. This scalability allows detection systems to maintain comprehensive visibility across entire transaction networks while providing interactive query performance for fraud investigators exploring suspicious patterns.

Real-time event processing pipelines transform raw transaction streams into continuously updated graph structures, enabling detection of suspicious patterns as they emerge rather than after funds have been transferred. These pipelines implement multi-stage processing workflows that validate, enrich, and transform transaction events before updating the graph database and evaluating detection rules. Cloud-native event processing leverages distributed stream processing frameworks that parallelize computation across hundreds or thousands of nodes, ensuring consistent throughput even during peak transaction periods. The event-driven architecture patterns enable seamless integration of multiple data sources, combining transaction data with

external risk signals and contextual information to create rich graph representations [6]. This streaming approach closes the detection time gap that sophisticated financial criminals exploit, enabling real-time intervention rather than post-facto investigation.

Containerization and orchestration technologies revolutionize the deployment and management of graph analytics workloads, enabling financial institutions to implement detection algorithms with unprecedented agility. By packaging algorithms and dependencies in standardized containers, organizations ensure consistent execution across diverse environments while eliminating deployment friction. Container orchestration platforms automatically distribute analytics workloads across compute clusters based on resource requirements and priority levels, ensuring critical detection processes receive necessary computing power during high-demand periods. Advanced orchestration capabilities such as auto-scaling and self-healing minimize operational overhead while maximizing system reliability [5]. This approach dramatically accelerates the deployment of new detection capabilities, enabling rapid response to emerging financial crime patterns.

Serverless computing models address the inherent variability in financial crime detection workloads by automatically scaling resources based on current demand without requiring explicit infrastructure provisioning. This approach is particularly valuable for computationally intensive graph algorithms that execute sporadically, such as global pattern analysis or investigation of suspicious subgraphs. The event-driven nature of serverless architectures aligns perfectly with the asynchronous workflow of financial crime detection, where initial alerts trigger progressively deeper analysis based on risk assessment [6]. This approach creates perfect alignment between computing costs and actual detection requirements, significantly reducing infrastructure expenses compared to traditionally provisioned environments.

Component	**Purpose**	**Benefit**
Graph Databases	Store and query relationship data	Fast traversal and scalability
Event Pipelines	Stream and process transactions in real time	Immediate anomaly detection
Containerization	Package and deploy detection algorithms	Consistent and agile deployments
Serverless Computing	Auto-scale resources on demand	Cost-effective workload handling
Multi-Region Setup	Distribute systems across geographies	Resilient and jurisdiction-compliant

Table 1: Cloud-Native Components for Graph-Based Anomaly Detection [5, 6]

Multi-region deployment architectures address the challenges of global financial networks that span geographic and jurisdictional boundaries. By distributing graph database instances across regions while maintaining logical connectivity, financial institutions can comply with data sovereignty requirements while still detecting cross-border transaction patterns that would remain invisible in siloed systems. This approach also provides critical resilience against regional outages, ensuring continuous monitoring capabilities even when individual regions experience disruption [6].

IV. Methodology and Implementation

Data ingestion and preprocessing establish the foundation for effective graph-based anomaly detection systems in financial environments. The process begins with collecting diverse data streams from multiple sources, including core banking platforms, payment networks, and customer information systems. Each system typically employs different formats, identifiers, and temporal frameworks, creating significant integration challenges. Preprocessing pipelines standardize these heterogeneous inputs through transformation workflows that normalize transaction types, align timestamps, and resolve entity references. Entity resolution represents a particularly critical challenge, as customers may appear under different identifiers across systems. Advanced frameworks employ probabilistic matching algorithms that consider multiple attributes to establish entity identity even with inconsistent data. Graph construction algorithms transform these standardized inputs into a coherent network structure, creating nodes for entities and edges for relationships [7]. The resulting graph serves as a unified representation of the financial ecosystem, capturing both direct and indirect relationships that might indicate suspicious activities.

Graph enrichment enhances detection capabilities by incorporating contextual information beyond basic transaction data. This process augments the graph structure with additional attributes that characterize entities and relationships based on historical patterns and domain knowledge. Entity attributes typically include account longevity, transaction patterns, risk classifications, and demographic information. Edge attributes capture transaction characteristics such as amount, frequency, channel, and temporal patterns. Advanced enrichment pipelines calculate derived features that capture behavioral patterns over time, such as

changes in transaction velocity or shifts in counterparty relationships. Graph structural features provide another critical dimension, with metrics such as centrality scores and community memberships offering insights into an entity's position within the broader financial network [7]. This rich representation enables detection systems to distinguish between normal variations in financial behavior and genuinely suspicious activities.

Graph neural network training addresses unique challenges in financial crime detection, particularly the extreme class imbalance where legitimate transactions vastly outnumber fraudulent ones. Effective training strategies implement techniques such as stratified sampling, synthetic minority oversampling, and cost-sensitive learning objectives. GNN architectures for financial applications typically employ message-passing frameworks that allow nodes to aggregate information from their neighborhood, capturing both direct relationships and broader network patterns. Advanced architectures incorporate attention mechanisms that focus on the most relevant connections when aggregating information [7]. Continuous learning frameworks periodically retrain models on recent data while preserving knowledge of established patterns, maintaining detection effectiveness as the financial ecosystem evolves.

Anomaly scoring and investigation workflows transform model outputs into actionable intelligence. Modern approaches implement multi-faceted scoring systems combining predictions from multiple detection methods, including supervised classification, unsupervised anomaly detection, and rule-based pattern matching. Dynamic thresholding approaches adjust alert criteria based on risk context and investigation capacity. Investigation platforms provide interactive graph visualization interfaces that allow analysts to explore entity relationships and transaction patterns, significantly improving alert disposition decisions [8]. Case management workflows maintain comprehensive audit trails of investigation steps, evidence gathered, and disposition decisions. The integration of graph-based visualization with structured investigation workflows enhances both efficiency and effectiveness, enabling compliance teams to process higher alert volumes while improving identification of genuinely suspicious activities.

Stage	Purpose	Benefit
Data Ingestion & Preprocessing	Standardize and resolve multi-source data	Unified and accurate graph construction
Graph Enrichment	Add contextual and behavioral attributes	Improved detection accuracy and pattern clarity
GNN Training	Learn from graph structures with imbalanced data	Detect complex and rare fraud patterns
Anomaly Scoring	Assign risk scores using multiple methods	Prioritized and context-aware alerting
Investigation Workflows	Visualize and manage suspicious patterns	Efficient and auditable fraud investigations

Table 2: Methodology and Implementation in Graph-Based Anomaly Detection [7, 8]

V. Case Studies and Performance Analysis

Graph-based anomaly detection has proven remarkably effective in identifying insider trading rings through the analysis of unusual transaction patterns that remain invisible to traditional monitoring approaches. Unlike conventional methods that examine individual accounts in isolation, graph-based techniques analyze the complex web of relationships between traders, revealing coordinated activities that appear innocuous when viewed separately. These systems construct temporal transaction graphs that capture not just the flow of funds but also the timing patterns that characterize information-based trading. By examining structural characteristics including unusual clustering coefficients, atypical centrality distributions, and distinctive temporal sequences, detection systems can identify coordinated trading behavior despite deliberate attempts to disguise connections between participants [9]. This methodology successfully identifies complex insider trading networks where information flows through multiple intermediaries before resulting in actual trades, creating sufficient distance between information sources and trading activity to evade traditional surveillance.

Money laundering operations spanning multiple jurisdictions present particularly challenging detection problems that graph-based approaches are uniquely positioned to address. These criminal operations deliberately structure activities across geographic and institutional boundaries, exploiting the fragmentation of monitoring systems. Graph-based detection overcomes these limitations by constructing comprehensive networks that integrate transaction data across boundaries, revealing complete fund paths despite deliberate fragmentation. Graph sampling techniques make this analysis computationally feasible even for

massive financial networks by focusing resources on suspicious substructures within the broader network [10]. This holistic perspective enables the detection of sophisticated money laundering techniques such as fan-in/fan-out structures, parallel transaction chains, and round-trip transactions that traditional systems examining individual transactions cannot recognize.

Performance metrics demonstrate that graph-based detection delivers substantial improvements over traditional approaches across multiple dimensions, including accuracy, false positive rates, and detection speed. By capturing the structural context in which transactions occur rather than analyzing them as isolated events, graph-based methods significantly reduce the ambiguity that leads to false positives. This contextual awareness is particularly valuable in reducing false positives for high-risk but legitimate activities such as remittance businesses serving developing economies or charitable organizations operating in conflict zones [9]. The structural patterns revealed through graph analysis provide powerful signals that complement traditional detection features based on transaction amounts and frequencies. Most significantly, graph-based approaches can identify suspicious patterns much earlier in their development, often after just a few transactions that establish distinctive structural signatures.

Scalability analysis confirms that modern graph-based detection systems can maintain performance under extreme transaction volumes. Importance sampling techniques reduce computational requirements by focusing on the most informative substructures within massive graphs, enabling analysis of networks with billions of nodes without exhaustive processing [10]. Distributed processing architectures further enhance scalability by partitioning graph analysis across multiple computing nodes. Memory-efficient representations reduce resource requirements by compressing graph structures while maintaining query performance. These scalability enhancements ensure that graph-based detection can handle the transaction volumes of even the largest global financial institutions without sacrificing effectiveness or timeliness.

Comparative analysis reveals fundamental advantages of graph-based approaches across diverse financial crime scenarios. While rule-based systems excel at detecting known patterns with clearly defined characteristics, they struggle with coordinated activities distributed across multiple accounts. Graph-based methods address these limitations by modeling relationships between entities as a core element of the detection process [9].

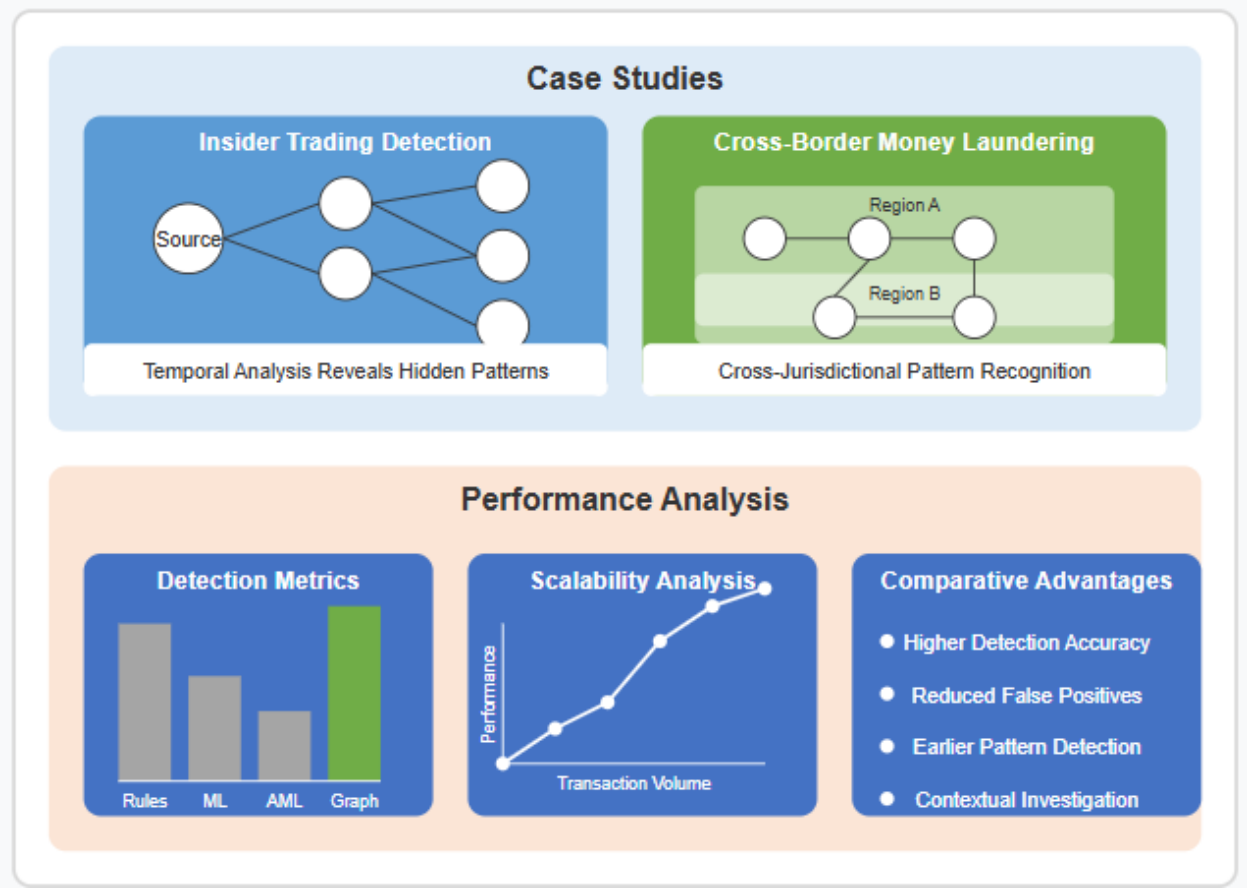

Fig 2: Case Studies and Performance Analysis of Graph-Based Detection [9, 10]

Conclusion

Graph-based anomaly detection represents a paradigm shift in financial crime prevention, providing unprecedented capabilities for identifying sophisticated schemes that traditional systems miss entirely. The representation of financial activities as interconnected networks aligns naturally with how financial crimes actually operate, enabling the detection of coordinated behaviors distributed across multiple entities. Cloud-native architectures deliver the computational power and scalability necessary to maintain comprehensive visibility across global transaction networks while enabling real-time intervention before significant damage occurs. Future advancements in federated graph learning will extend these capabilities across institutional boundaries while preserving data privacy, creating collaborative defense networks against financial crime. Regulatory frameworks increasingly recognize the value of relationship-based detection approaches, with compliance guidelines evolving to encourage adoption of advanced analytics. As graph technologies mature, financial institutions implementing these systems gain not just improved compliance outcomes but genuine competitive advantages through reduced fraud losses, operational efficiencies, and enhanced customer trust.

Funding: This research received no external funding.
Conflicts of Interest: The authors declare no conflict of interest.
Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers.

References

[1] Sanction Scanner, "2023-2024 Financial Crime & Compliance Report". [Online]. Available: https://www.sanctionscanner.com/content/report/2023-2024-financial-crime-and-compliance-report.pdf

[2] Walter Didimo et al., "Network visualization for financial crime detection," ScienceDirect, 2014. [Online]. Available: https://www.sciencedirect.com/science/article/abs/pii/S1045926X1400010X

[3] Soroor Motie and Bijan Raahemi, "Financial fraud detection using graph neural networks: A systematic review," ScienceDirect, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/abs/pii/S0957417423026581

[4] Tahereh Pourhabibi et al., "Fraud detection: A systematic literature review of graph-based anomaly detection approaches," ScienceDirect, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0167923620300580

[5] Chaoqi Chen et al., "A Survey on Graph Neural Networks and Graph Transformers in Computer Vision: A Task-Oriented Perspective," arXiv:2209.13232, 2024. [Online]. Available: https://arxiv.org/abs/2209.13232

[6] Sahini Dyapa, "Real-Time Fraud Detection: Leveraging Apache Kafka and Spark for Financial Transaction Processing," IJSAT, 2025. [Online]. Available: https://www.ijsat.org/papers/2025/1/2654.pdf

[7] Qian Huang et al., "Combining Label Propagation and Simple Models Out-performs Graph Neural Networks," arXiv:2010.13993, 2020. [Online]. Available: https://arxiv.org/abs/2010.13993

[8] Sudhansu Ranjan Lenka and Dr. Bikram Kesari Ratha, "Financial Fraud Detection using Data Mining: A Survey," IJCSNS, 2024. [Online]. Available: http://paper.ijcsns.org/07_book/202409/20240920.pdf

[9] Neil Shah et al., "EdgeCentric: Anomaly Detection in Edge-Attributed Networks," arXiv:1510.05544, 2015. [Online]. Available: https://arxiv.org/abs/1510.05544

[10] Jie Chen et al., "FastGCN: Fast Learning with Graph Convolutional Networks via Importance Sampling," arXiv:1801.10247, 2018. [Online]. Available: https://arxiv.org/abs/1801.10247