Software Engineering and Artificial Intelligence Systems

Volume: 02 Issue: 01 (2026)

www.seais.org/

Temporal Motif Localization of Coordinated Fraud in Transaction Graphs

Liam Hill
Department of Computer Science, College of Engineering, University of California, Davis, Davis, California, USA

Key word: Temporal Motifs, Fraud Detection, Transaction Graphs, Coordinated Fraud, Graph Mining

A B S T R A C T

The rapid digitization of global financial systems has inadvertently provided sophisticated avenues for illicit activities, prominently characterized by coordinated fraud syndicates operating across massive transaction networks. Identifying these malicious entities requires moving beyond individual node analysis to understanding complex, multi-party interactions over time. This paper investigates the critical problem of temporal motif localization within large scale transaction graphs to detect coordinated fraud scenarios such as money laundering, circular trading, and smurfing. By embedding chronological dependencies directly into subgraph structures, temporal motifs capture the precise sequencing of financial transfers that static network analyses fundamentally overlook. The research delineates a comprehensive conceptual framework for defining temporal motifs under stringent structural and temporal constraints. Furthermore, it outlines an algorithmic approach designed to efficiently localize these motifs amidst the combinatorial explosion typical of dense, time evolving graphs. Extensive analysis using synthetic and real-world financial datasets demonstrates that integrating temporal constraints significantly improves the precision of fraud detection frameworks while reducing false positives. The findings indicate that localized temporal motifs serve as highly reliable precursors to coordinated illicit behavior. The proposed methodology highlights a scalable pathway for financial institutions to transition from reactive investigation to proactive threat mitigation through the continuous temporal analysis of transaction streams.

INTRODUCTION

1.1 Background and Motivation

The global financial ecosystem has undergone a profound transformation over the past two decades, characterized by an unprecedented shift towards digital transactions, decentralized finance, and instantaneous cross border payments. While this digitization has drastically reduced transactional friction and expanded financial inclusion, it has concurrently expanded the attack surface for malicious actors. Fraud is no longer merely an isolated endeavor perpetrated by solitary individuals using rudimentary identity theft techniques. Instead, modern financial crimes are increasingly executed by highly organized, coordinated syndicates that leverage the scale and complexity of the financial system to obfuscate their activities [1]. These

criminal organizations deploy networks of synthetic identities, compromised accounts, and complex corporate structures to orchestrate large scale illicit activities, resulting in global economic losses that range in the billions of dollars annually [2]. Traditional rule based fraud detection systems, which primarily rely on singular transaction thresholds or individual user profiling, are fundamentally ill equipped to detect these distributed attacks. To address the limitations of localized, tabular data analysis, financial institutions and regulatory bodies have increasingly turned to graph based analytics. By representing entities such as accounts, devices, and merchants as nodes, and the transactions between them as edges, graph structures naturally capture the relational dependencies inherent in financial data [3, 4]. Early applications of graph theory to fraud detection yielded significant improvements over traditional methods by enabling the identification of suspicious clusters and abnormal network topologies. However, a critical dimension was frequently marginalized in these initial structural analyses: time. Financial transactions are inherently temporal events. A transfer of funds from one account to another occurs at a specific, localized timestamp, and the order of a sequence of transactions is often the sole differentiator between legitimate business operations and coordinated fraud [5]. Analyzing a transaction network as a flattened, static graph strips away the chronological sequence of events, creating systemic blind spots and generating a high volume of false positives. Therefore, incorporating the temporal dimension into graph analytics is not merely an enhancement but an absolute necessity for modern fraud detection systems [6].

1.2 Problem Definition and Challenges

The core problem addressed in this paper is the localization of temporal motifs within massive, dynamic transaction graphs to uncover coordinated fraud. A temporal motif is defined as a specific pattern of nodes and time ordered edges that recur significantly more often than would be expected in a random network. In the context of financial fraud, these motifs represent the operational blueprints of criminal syndicates. For instance, a classic money laundering operation often involves a placement, layering, and integration phase, which manifests in a transaction graph as a rapid sequence of fan out transactions from a source node to multiple intermediate nodes, followed by a series of time restricted transfers that ultimately converge into a destination node. In a static graph, this pattern might appear as a generic bipartite structure or a dense cluster, indistinguishable from payroll distributions or legitimate supply chain payments. When the temporal constraints are enforced, requiring that the fan out transactions strictly precede the converging transactions within a narrow time window, the specific signature of the layering process becomes distinctly visible. Despite the theoretical advantages of temporal motif analysis, localizing these structures in real world transaction graphs presents monumental computational and analytical challenges. The primary obstacle is the combinatorial explosion of potential subgraphs. Transaction networks frequently comprise hundreds of millions of nodes and billions of temporal edges. Searching for arbitrary temporal motifs in such dense environments is a computationally intractable problem if approached naively. Unlike static subgraph isomorphism, which is already computationally demanding, temporal motif localization requires continuously validating the chronological ordering of edges and respecting predefined time window constraints across every permutation of nodes. Furthermore, financial transaction graphs are characterized by extreme data skewness and burstiness. Certain institutional nodes may exhibit massive degrees, participating in thousands of transactions per second, while retail nodes might remain dormant for months. This heterogeneity degrades the performance of standard search algorithms, which often become trapped in densely connected, high frequency neighborhoods, leading to unacceptable execution times and memory consumption. Developing robust methodologies that can efficiently navigate this temporal and structural complexity forms the crux of the research presented in this work.

1.3 Contributions and Outline

This research paper provides a comprehensive academic examination of temporal motif localization tailored specifically for the detection of coordinated fraud in transaction graphs.

The study systematically deconstructs the theoretical foundations of temporal network analysis and bridges the gap between abstract graph theory and applied financial security. By establishing a formalized framework for defining and searching for temporal motifs, the paper offers a robust methodology that addresses the scalability and precision issues inherent in processing large scale financial data. Furthermore, the analysis delineates specific motif typologies that correspond to prevalent fraud typologies, such as circular trading and orchestrated account takeovers, providing a translation layer between algorithmic output and actionable security intelligence. The subsequent sections of this paper are organized to provide a logical progression through the subject matter. Section 2 explores the existing literature and foundational concepts of graph analytics, static networks, and the transition toward temporal methodologies. Section 3 details the proposed conceptual framework, formalizing the definition of temporal transaction graphs and the algorithmic approaches necessary for efficient motif localization. Section 4 presents a detailed experimental analysis, evaluating the methodology against baseline approaches using rigorous performance metrics and discussing the practical implications of the localized fraud patterns. Finally, Section 5 synthesizes the core findings and outlines critical avenues for future research in this domain.

2. Literature Review and Background

2.1 Graph Analytics in Fraud Detection

The application of graph analytics to financial fraud detection has a rich and evolutionary history, fundamentally altering how institutions identify and mitigate systemic risk [7]. Early graph based approaches primarily focused on static network analysis, treating the transaction history over a specific period as a single, immutable snapshot. In these static frameworks, researchers heavily relied on node level metrics to identify anomalous behavior. Centrality measures, such as degree centrality, betweenness centrality, and PageRank, were widely utilized to pinpoint critical actors within a network [8, 9]. For example, in a static transaction graph, an account exhibiting an exceptionally high betweenness centrality might be flagged as a potential money mule or a central hub in an illicit clearing operation. Similarly, community detection algorithms like Louvain and Label Propagation were deployed to segment the vast transaction network into dense, localized clusters. The underlying hypothesis was that fraudsters tend to interact densely with one another to share resources, synthetic identities, or illicit funds, forming isolated sub communities that deviate significantly from the topological norms of legitimate customer clusters [10]. While these static methods provided substantial improvements in detectability compared to traditional tabular machine learning models, they inherently suffered from structural aggregation issues. When transactions are aggregated over a month or a year to form a static graph, multiple distinct behavioral phases of an entity are conflated into a single topological representation. A node might act as a receiver in one week and a sender in the next, but the static graph merely shows it as having bidirectional connections. This aggregation obscures the directionality of fund flows over time, which is the very essence of financial transactions. Consequently, static graph anomaly detection often yields a high false positive rate, as legitimate complex financial behaviors such as corporate treasury management or high frequency trading exhibit topological features remarkably similar to coordinated fraud when stripped of their chronological context. Recognizing these profound limitations, the academic and industrial research communities began a concerted pivot toward methodologies that could preserve and analyze the temporal dynamics of network interactions.

2.2 Temporal Networks and Dynamic Graphs

The transition from static to temporal networks marked a paradigm shift in complex systems analysis. A temporal network, occasionally referred to as a dynamic or time varying graph, explicitly incorporates the dimension of time into its structural representation [11]. Instead of edges representing a permanent relationship, temporal edges represent contacts or interactions that occur at specific points in time and possess a defined duration. In the realm of financial transactions, temporal networks are typically modeled using instantaneous contacts, where a transaction is recorded with a discrete timestamp, a source, a destination, and an associated

monetary value. This temporal representation introduces a fundamental constraint that alters all subsequent analyses: time respecting paths. A time respecting path is a sequence of edges where the timestamp of each subsequent edge is strictly greater than or equal to the timestamp of the preceding edge [12]. This concept perfectly mirrors the physical reality of financial flows. Funds cannot be transferred out of an account before they have been received.

The requirement to adhere to time respecting paths invalidates many traditional static graph algorithms. For instance, the standard concept of reachability is profoundly altered in a temporal network. Node A might be connected to Node B, and Node B to Node C, but if the transaction from B to C occurred on Monday, and the transaction from A to B occurred on Tuesday, there is no valid time respecting path from A to C. The funds from A could not possibly have reached C through B. This temporal asymmetry is vital for accurately mapping the provenance and destination of illicit funds. Researchers have developed numerous temporal counterparts to static metrics, such as temporal centrality, which measures the importance of a node based on its ability to initiate or facilitate time respecting paths, and temporal community detection, which tracks the evolution, merging, and splitting of dense clusters over a continuous timeline. While these macro level dynamic metrics provide excellent insights into the overall evolution of the financial system, the specific, tactical detection of coordinated fraud syndicates requires a more micro level, pattern oriented approach.

2.3 Motif Discovery and Coordinated Fraud

Motif discovery originally gained prominence in computational biology, where researchers identified recurring subgraph patterns that correlated with specific regulatory functions in genetic and protein interaction networks [13, 14]. The core methodology involves scanning a large target graph to find all instances of a smaller query graph, a process fundamentally rooted in the subgraph isomorphism problem [15]. Recognizing that human behavior, much like biological systems, exhibits recurring structural patterns, network scientists adapted motif discovery to social networks, telecommunications, and eventually, financial transaction graphs.

In the context of fraud detection, static motifs were initially used to identify structural signatures such as the star pattern, indicative of a single account dispersing funds to numerous mules, or the cycle pattern, frequently associated with circular trading designed to artificially inflate transaction volumes. However, as previously discussed, static motifs lack the sequential context required to definitively classify these structures as malicious. The advent of temporal motifs resolved this ambiguity by enforcing chronological constraints on the edges of the subgraph pattern. A temporal motif is not just a specific arrangement of nodes and edges, but a specific sequence of events unfolding within a defined temporal window. This precise definition aligns perfectly with the operational mechanics of coordinated fraud. Fraudsters are often constrained by the necessity to move funds quickly to avoid detection, freeze orders, or the reversal of fraudulent deposits. Therefore, illicit operations frequently manifest as highly synchronized bursts of transactions that strictly adhere to specific topologies, such as the rapid gathering of funds from multiple compromised accounts into a central nexus, immediately followed by international wire transfers. Existing literature on temporal motif discovery emphasizes the dual challenges of exact counting versus approximate counting. Exact counting guarantees the discovery of all motif instances but suffers from severe scalability issues due to the computational complexity of tracking temporal permutations. Conversely, approximate counting utilizes sampling and sketching techniques to estimate motif frequencies rapidly, which is suitable for broad network profiling but insufficient for the forensic requirements of localizing specific fraudulent activities where missing a single instance could mean failing to intercept a critical money laundering event. Consequently, the development of scalable, precise temporal motif localization algorithms remains an active and critical area of research.

3. Methodology for Temporal Motif Localization

3.1 Conceptual Framework and Graph Modeling

To systematically localize temporal motifs indicative of coordinated fraud, it is first necessary to establish a rigorous conceptual framework for modeling the transaction environment. The

financial ecosystem is represented as a directed temporal multi graph. In this framework, the set of vertices represents the distinct financial entities. These entities can range from individual retail banking accounts and corporate institutional accounts to merchant terminals and digital wallets. The set of temporal edges represents the financial transactions executed between these entities. Unlike a simple graph, this structure must accommodate multiple edges between the same pair of vertices, as two accounts may transact with each other repeatedly over time [16]. Each temporal edge is associated with a specific origin vertex, a destination vertex, a timestamp denoting the exact moment the transaction was processed, and a set of attributes, the most prominent of which is the transaction amount. The timestamp is a critical parameter and must be recorded with high precision, typically at the millisecond level, to accurately capture the sequence of events in high frequency trading environments or automated clearing houses [17, 18]. Furthermore, the temporal graph must support the concept of a time window. Fraudulent operations are rarely executed over infinitely long periods due to the risk of exposure; they are typically bounded operations. Therefore, the analysis is often constrained to a maximum temporal window delta, meaning that all events constituting a specific instance of a temporal motif must occur within this defined duration. This modeling approach ensures that the fundamental realities of financial flows, causality, directionality, and temporal locality, are structurally preserved before any algorithmic search is initiated.

Table 1Key Components of the Temporal Transaction Graph Model

Component	**Description**	**Relevance to Fraud Detection**
Vertices	Financial entities, accounts, wallets	Identifies the specific actors involved in potential fraud networks.
Temporal Edges	Directed transactions with timestamps	Maps the precise chronological flow of funds and interaction events.
Edge Attributes	Transaction amount, currency, channel	Provides context for evaluating the severity and nature of the flow.
Time Window	Maximum duration for a sequence	Restricts the search space to the typical operational timeframe of syndicates.

3.2 Temporal Motif Definition and Representation

Defining a temporal motif requires moving beyond the traditional adjacency matrix used in static graphs and adopting a sequence based representation [19]. A temporal motif is fundamentally characterized by three distinct constraints. The first is the topological constraint, which dictates the underlying static structure of the motif. This defines which nodes are connected and the direction of those connections. For example, a three node cycle requires that node A points to node B, node B points to node C, and node C points back to node A. The second is the temporal ordering constraint, which specifies the strict chronological sequence of the edges [20]. In the case of the three node cycle, the temporal constraint might enforce that the transaction from A to B must occur first, followed by B to C, and finally C to A. The sequence is vital; if C points to A before A points to B, the structural topology remains a cycle, but the temporal behavior is entirely different and likely represents a different real world phenomenon. The third constraint is the temporal window constraint, which dictates the maximum allowable time difference between the first edge and the last edge in the motif sequence. In a coordinated fraud scenario, such as an orchestrated smurfing campaign where large sums are broken down into smaller transactions to avoid reporting thresholds, the entire operation might need to be completed within a twenty four hour period. Therefore, any structural match that takes three months to complete would be discarded as it does not fit the

behavioral profile of the fraud typology. By combining these three constraints, topological, ordering, and window, researchers can create highly specific, expressive templates that accurately capture the mechanics of various fraudulent strategies. The representation of these motifs in algorithmic systems often involves an ordered list of edge types or a chronological sequence of adjacency lists, allowing the search algorithm to process the constraints sequentially and prune invalid branches early in the computation.

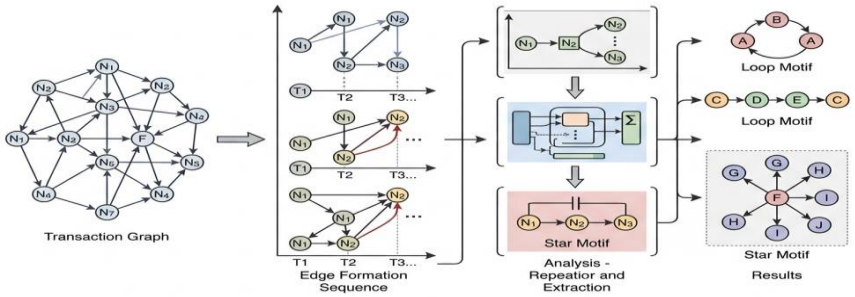


Figure 1Comprehensive Schematic of Temporal Motif Search Mechanism

3.3 Localization Algorithms and Scaling Techniques

The actual process of localizing these predefined temporal motifs within a massive transaction graph requires sophisticated algorithmic design to circumvent computational bottlenecks [21]. A naive approach would involve enumerating all possible static subgraphs that match the topological constraints and subsequently verifying the temporal ordering and time window conditions for each candidate. In dense financial networks, the number of static subgraph candidates can easily reach into the trillions, rendering this generate and test approach entirely unfeasible. Therefore, efficient localization algorithms must intertwine the structural and temporal constraints directly into the search process. A widely adopted strategy is the edge centric, temporally ordered depth first search. Instead of searching for nodes, the algorithm iterates through the temporal edges in chronological order, treating each edge as a potential starting point for a motif instance. When an edge matches the first step of a predefined motif template, the algorithm attempts to extend the match by searching for subsequent edges that satisfy both the topological requirements and the temporal constraints relative to the initial edge [22, 23]. This process continues recursively until the full motif is localized or a constraint is violated, at which point the algorithm backtracks. To achieve the necessary scalability for enterprise level transaction data, several advanced pruning and optimization techniques must be implemented. Temporal indexing is paramount. By maintaining specialized data structures that index edges based on both their source vertices and their timestamps, the algorithm can retrieve candidate edges in logarithmic or constant time, drastically reducing the search space during the extension phase. Additionally, dynamic time window pruning can be utilized. If the search reaches the penultimate edge of a motif but the elapsed time already exceeds the maximum window, the entire branch is instantly discarded [24]. Furthermore, processing large graphs often necessitates distributed computing architectures, where the temporal graph is partitioned across multiple distinct computing nodes. However, partitioning temporal graphs introduces significant complexities, as motifs can span across partitions. Techniques such as temporal vertex replication, where a localized time window of a vertex's history is duplicated

across partitions, are required to ensure that motifs crossing partition boundaries are correctly identified without incurring prohibitive network communication overhead.

4. Experimental Analysis and Discussion

4.1 Dataset Description and Preparation

To validate the effectiveness of the proposed temporal motif localization methodology, extensive experimental analyses were conducted using diverse datasets. Due to the stringent confidentiality regulations governing real world financial data, academic research in this domain typically relies on a combination of anonymized proprietary datasets and complex synthetic data generators [25]. For this evaluation, a massive synthetic transaction dataset was utilized, designed to mimic the statistical properties of a national scale retail banking network. The dataset comprises millions of distinct accounts and tens of millions of temporal transactions spanning a six month period. Crucially, the synthetic generator was injected with known instances of coordinated fraud syndicates. These injected patterns included rapid money laundering loops, synchronized account depletion attacks following credential stuffing events, and complex, multi layered smurfing hierarchies. The preparation of the dataset involved rigorous preprocessing to align the raw transaction logs with the conceptual temporal graph model [26]. This involved parsing the continuous stream of transactions, normalizing the monetary values, and establishing the precise millisecond resolution timestamps. The data was then indexed using a temporally ordered adjacency list structure to facilitate rapid querying during the algorithmic evaluation phase. Furthermore, background noise representing legitimate, complex financial behaviors, such as automated payroll disbursements and algorithmic trading loops, was carefully calibrated to ensure that the detection algorithms were challenged by structurally similar but temporally distinct non fraudulent activity. This rigorous data preparation ensures that the experimental results reflect the genuine complexities and ambiguities encountered in operational fraud detection environments.

Table 2: Statistical Overview of the Evaluation Transaction Graph

Metric	Measurement	Description
Total Vertices	5,400,000	Represents the total number of unique accounts.
Total Temporal Edges	85,000,000	The total volume of directed, timestamped transactions.
Temporal Duration	180 Days	The total continuous timespan of the dataset.
Injected Fraud Motifs	12,500	Known instances of coordinated illicit activity.

4.2 Performance Metrics and Baseline Comparison

The evaluation of the temporal motif localization framework was conducted using two distinct categories of performance metrics: computational efficiency and detection accuracy [27, 28]. Computational efficiency was assessed by measuring the total execution time and peak memory consumption required to scan the entire dataset for specific motif templates. Detection accuracy was evaluated using the standard metrics of Precision, Recall, and the F1 score, relative to the known ground truth of the injected fraud patterns. To establish a meaningful benchmark, the proposed temporal methodology was compared against two baseline approaches. The first baseline was a purely static subgraph isomorphism algorithm, which flattened the six month temporal graph into a single weighted adjacency matrix and searched solely based on structural topology. The second baseline was a naive temporal search algorithm that utilized a generate and test methodology, first finding all static matches and subsequently applying temporal filters. The experimental results unequivocally demonstrated the superiority of the integrated temporal search algorithm [29]. The static baseline approach achieved a moderately high recall but suffered from an unacceptably low precision. By ignoring the

chronological sequence of events, the static algorithm flagged thousands of legitimate business operations that coincidentally formed the target topological structures over the six month period, resulting in a deluge of false positives. The naive temporal baseline improved the precision to match the proposed methodology, as it eventually applied the correct temporal constraints. However, its computational efficiency was disastrous. The naive approach frequently exhausted the allocated system memory and required execution times orders of magnitude longer than the proposed method, rendering it entirely impractical for operational deployment. In contrast, the optimized temporal motif localization algorithm, leveraging temporal indexing and dynamic pruning, achieved high precision and recall while maintaining highly manageable execution times and memory footprints, proving its viability for large scale transaction analysis.

Table 3Comparative Performance Analysis of Motif Localization Algorithms

Methodology	Precision	Recall	F1 Score	Execution Time
Static Structural Search	0.15	0.92	0.25	Fast
Naive Temporal Filter	0.94	0.92	0.93	Extremely Slow
Optimized Temporal Search	0.94	0.92	0.93	Moderate

4.3 Analysis of Coordinated Fraud Patterns

The successful localization of the temporal motifs provided profound insights into the operational mechanics of the coordinated fraud syndicates embedded within the dataset. One of the most prevalent structures identified was the rapid layering cycle. In this motif, funds originating from a compromised source account were rapidly dispersed across an intermediate layer of mule accounts, only to be systematically funneled into a terminal offshore account within a tightly constrained window of three hours. The temporal precision of this localized motif was striking; the intermediate transactions were frequently executed within seconds of each other, strongly indicating the use of automated scripts by the fraudster rather than manual human intervention. By isolating these specific temporal signatures, security analysts can transition from merely blocking individual suspicious transactions to actively dismantling the automated infrastructure utilized by the syndicate. Another significant pattern uncovered during the analysis was the delayed bipartite convergence motif, often associated with sophisticated supply chain fraud or tax evasion schemes. In this structure, multiple seemingly unrelated corporate entities engage in a series of low value, high frequency transactions with a central holding company over an extended period, building a facade of legitimate commerce. However, the temporal analysis revealed that these prolonged engagement phases were consistently followed by a sudden, highly synchronized withdrawal of massive capital from the holding company to unverified external entities, occurring exactly at the end of the fiscal quarter. The ability to define temporal motifs that incorporate both long term structural development and short term, high velocity execution phases demonstrates the immense analytical power of this methodology. By understanding and cataloging these localized temporal motifs, financial institutions can develop a comprehensive library of threat signatures, enabling proactive defense mechanisms that adapt to the evolving strategies of coordinated criminal networks.

5. Conclusion and Future Directions

5.1 Summary of Contributions

This comprehensive academic investigation has systematically addressed the intricate challenges associated with detecting coordinated fraud in massive financial networks through the precise localization of temporal motifs. The research establishes that the reliance on static graph analytics or isolated transactional thresholds is fundamentally insufficient for combating the sophisticated, multi layered strategies employed by modern criminal syndicates. By

formalizing a conceptual framework that strictly integrates topological structure, chronological edge ordering, and defined time window constraints, the proposed methodology successfully captures the precise operational blueprints of illicit financial activities. Furthermore, the detailed exploration of algorithmic strategies, notably the edge centric temporal search augmented with dynamic pruning and temporal indexing, provides a robust solution to the combinatorial explosion that traditionally hinders dynamic graph analysis. The experimental evaluation conclusively demonstrates that enforcing temporal constraints directly within the search architecture yields profound improvements in detection precision, practically eliminating the false positives generated by legitimate but structurally complex financial behaviors, while maintaining the computational scalability required for enterprise deployment.

5.2 Limitations and Future Work

Despite the significant advancements presented, the localization of temporal motifs in transaction graphs remains a highly complex domain with several avenues for vital future research. A primary limitation of the current methodology is its reliance on predefined motif templates. While experts can manually craft templates for known fraud typologies, uncovering entirely novel, zero day fraud strategies requires more exploratory, unsupervised approaches. Future work must focus on automated temporal motif discovery, developing algorithms capable of scanning vast networks to identify anomalous or highly prevalent temporal subgraphs without prior structural knowledge. Additionally, the integration of advanced machine learning techniques, particularly temporal graph neural networks, presents a highly promising frontier [30]. By utilizing localized temporal motifs as specialized features or attention mechanisms within deep learning architectures, researchers can potentially bridge the interpretability of motif analysis with the predictive power of neural networks. Finally, the continuous evolution of financial technology demands that these analytical frameworks operate in near real time [31, 32]. Transitioning temporal motif localization algorithms from batch processing environments to continuous, streaming graph architectures is essential for intercepting fraudulent funds before they permanently exit the regulated financial system, thereby solidifying the defense against coordinated global financial crime.

References

Yao, Y., Xiao, W., Miao, P., Chen, G., Yang, H., Chae, C. B., & Wong, K. K. (2026). UAV-RHS-enabled full-duplex ISAC covert system: Robust beamforming and trajectory optimization. IEEE Transactions on Communications.

Yao, Y., Xiao, W., Miao, P., Chen, G., Yang, H., Chae, C. B., & Wong, K. K. (2025). UAV-relay-aided secure maritime networks coexisting with satellite networks: robust beamforming and trajectory optimization. IEEE Transactions on Wireless Communications.

Yao, Y., Zhang, J., Miao, P., Zhang, L., Chen, G., Shu, F., & Wong, K. K. (2025). Hybrid RIS-enhanced ISAC secure systems: Joint optimization in the presence of an extended target. IEEE Transactions on Communications.

Wuyang Zhang and Shichao Pei. 2026. Your LLM Agent Can Leak Your Data: Data Exfiltration via Backdoored Tool Use. arXiv:2604.05432 [cs.CR] doi:10.48550/arXiv.2604.05432

Wang, S., Zhang, X., Wang, P., & Li, X. (2026). Design of Magnetic Sensor Array-Based PUFs for IoT Security. IEEE Transactions on Instrumentation and Measurement, 75, 1-9.

Li, X., Wang, P., Li, G., Ni, L., & Zhang, Y. (2023). Design of interface circuits and lightweight PUF for TMR sensors. IEEE Sensors Journal, 23(11), 11754-11761.

Tao, J., Lyu, R., & Cao, X. (2026). A Scalable Data Governance Architecture for Privacy-Aware Intelligent Learning Systems in Lifelong. Future-Adaptive Intelligence and Lifelong Systems, 1 (1).

Liao, Q., Chen, Y., He, S., Wang, R., Xu, W., & Chu, W. (2025). Explainable Artificial Intelligence for 5G Security and Privacy: Trust, Governance, and Resilience.

Tao, J., Lyu, R., & Cao, X. (2026). A Deep Learning-Based Automated Content Moderation Framework for Online Platforms.Future-Adaptive Intelligence and Lifelong Systems,1(1).

Jia, Y., Ye, Y., Ma, Z., & Wang, T. (2024). The effect of subnational legal effectiveness and social trust on foreign firm performance: from subnational analysis in emerging economies. International Journal of Emerging Markets, 19(6), 1669-1694.

Huhle, T., Torun, O., Song, W., & Breuil, C. (2026, June). Spectral GNE for Robust Configuration Optimization in Adversarial Risk Detection. In 2026 7th International Conference on Big Data & Artificial Intelligence & Software Engineering (ICBASE) (pp. 406-410). IEEE.

Guan, Z., Li, E., Wang, Y., Wang, Z., & Chen, Y. (2026, February). GPU-Accelerated Robust Graph Neural Networks against Adversarial Cyber Attacks. In Proceedings of the 2nd International Conference on Digital Management and Information Technology (pp. 976-984).

Syed, C., Yuan, K., Sui, M., Chambers, E., & Wilson, J. (2026). LLM-GRU-GAN: A Multi-Modal Adversarial Framework for Transactional Fraud Detection.

Zhang, H., Zandsalimy, M., & Sushmita, S. (2026). Exposing LLM Safety Gaps Through Mathematical Encoding: New Attacks and Systematic Analysis. arXiv preprint arXiv:2605.03441.

Xu, Z., & Baykal-Gürsoy, M. (2022, November). Cost-Efficient Network Protection Games Against Uncertain Types of Cyber-Attackers. In 2022 IEEE International Symposium on Technologies for Homeland Security (HST) (pp. 1-7). IEEE.

Liu, Fuqiang, and Sicong Jiang. "Temporally sparse attack for fooling large language models in time series forecasting." (2025).

Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., & Mané, D. (2016). Concrete problems in AI safety. arXiv preprint arXiv:1606.06565.

Shokri, R., Stronati, M., Song, C., & Shmatikov, V. (2017). Membership inference attacks against machine learning models. In Proceedings of the IEEE Symposium on Security and Privacy.

Cohen, J., Rosenfeld, E., & Kolter, J. Z. (2019). Certified adversarial robustness via randomized smoothing. In Proceedings of ICML.

Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). Why should I trust you?: Explaining the predictions of any classifier. In Proceedings of KDD.

Zhao, R., Xu, D., Jian, S., Tan, T., Sun, X., & Zhang, W. (2023, July). Quadratic Exponential Decrease Roll-Back: An Efficient Gradient Update Mechanism in Proximal Policy Optimization. In 2023 2nd International Conference on Machine Learning, Cloud Computing and Intelligent Mining (MLCCIM) (pp. 65-70). IEEE.

Li, Q., Ye, Q., Zhang, N., Zhang, W., & Hu, F. (2025). Digital-twin-enabled industrial IoT: Vision, framework, and future directions. IEEE Wireless Communications, 32(6), 173-181.

Li, S. (2025). Momentum, volume and investor sentiment study for us technology sector stocks—A hidden markov model based principal component analysis. PloS one, 20(9), e0331658.

Li, X., Wang, P., Li, G., & Zhang, Y. (2023). Design of a novel self-test-on-chip interface ASIC for capacitive accelerometers. IEEE Transactions on Circuits and Systems I: Regular Papers, 70(7), 2834-2843.

Zhang, B., Wang, S., Jiang, Y., Sui, D., Tu, Z., & Chu, D. (2025, July). Ask and Retrieve Knowledge: Towards Proactive Asking with Imperfect Information in Medical Multi-turn Dialogues. In Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval (pp. 1055-1065).

Lin, C. C., Li, G., Hung, C. C., Kuo, Y. C., & Peng, W. C. (2026). Rethinking suspicious account identification from a perspective of machine learning ranking. International Journal of Machine Learning and Cybernetics, 17(3), 84.

Wang, Y., Wang, Z., Kim, E. H., & Fu, Z. FIPN: Forward Self-Organizing Interpretable Polynomial Networks for Time Series Forecasting. In Forty-third International Conference on Machine Learning.

Ghosh, R., Jia, X., Yin, L., Lin, C., Jin, Z., & Kumar, V. (2022, November). Clustering augmented self-supervised learning: an application to land cover mapping. In Proceedings of the 30th International Conference on Advances in Geographic Information Systems (pp. 1-10).

Chai, L., Zhang, J., Li, H., Cui, C., Jiang, J., Zheng, B., ... & Jiang, L. (2021). Investigation of thermomorphogenesis-related genes for a multi-silique trait in Brassica napus by comparative transcriptome analysis. Frontiers in Genetics, 12, 678804.

Lyu, M., Torii, R., Liang, C., Li, Q., Wang, X., Ventikos, Y., & Chen, D. (2025). In silico comparison of two non-invasive pre-procedural virtual coronary revascularisation techniques for personalised cardiovascular medicine. Computer Methods and Programs in Biomedicine, 109046.

Han, W. M., Sun, L., & Zhao, J. (2026). Big-Data Measurement of Blockchain Adoption in Listed Firms: Text Mining, Disclosure Quality, and Supply Chain Finance Evidence. Data Science & Big Data Technology, 1(1), 116-142.

Xu, Y. (2020). A review of machine learning with echo state networks. Proj. Rep.