Signal-Aware Fraud Detection: A Taxonomy for Sustainable and Lifecycle-Oriented Detection Architectures

Sara Santedicola Ribeiro[1], **Fabio Lucio Lopes de Mendonça**[1], **Laerte Peotta**[1]

[1]Programa de Pós-Graduação em Engenharia Elétrica (PPEE),
Universidade de Brasília (UnB), Brasília-DF, Brazil

sarasantedicola@gmail.com, fabio.mendonca@redes.unb.br, peotta@gmail.com

Abstract. *Online financial fraud is growing in scale, coordination and operational complexity, challenging detection systems to balance accuracy, interpretability, scalability and computational efficiency. This work proposes a signal-aware perspective for fraud detection centered on the observable evidence generated throughout fraudulent activity. Fraud-relevant signals are organized into temporal, distributional, structural and contextual families, which may be operationalized explicitly, implicitly or through hybrid representations. To contextualize signal emergence, the paper introduces a conceptual fraud lifecycle model that distinguishes strategic and operational reconnaissance activities and highlights the asymmetric observability of fraud across lifecycle stages. Building on this perspective, the paper analyzes how different detection paradigms operationalize signal families and examines the associated trade-offs involving computational complexity, latency, scalability and interpretability. By reframing fraud detection as a process-aware and systems-oriented problem grounded in signal emergence and operationalization, this work contributes towards the design of more explainable, efficient and operationally robust fraud detection systems.*

1. Introduction

Financial fraud detection has become a central challenge in modern digital ecosystems due to the accelerated expansion of online financial services, instant payment infrastructures and large scale transactional platforms. The economic impact of digital fraud has reached unprecedented levels worldwide, with recent cybercrime reports estimating losses exceeding US$ 20.88 billion annually and indicating sustained growth in financially-driven digital attacks, with an increase of 26% over the previous year [FBI 2026]. In parallel, the rapid digitalization of financial ecosystems has dramatically increased transaction volume and operational exposure. In Brazil, the widespread adoption of the Pix instant payment system has intensified this scenario considerably. According to the 2025 Febraban Banking Technology Survey, the Brazilian banking sector processed more than 208 billion transactions in a single year, while Pix operations in mobile banking alone approached 25 billion transactions in 2024, representing annual growth above 40%. The same report also highlights a 38% increase in heavy Pix users and identifies fraud prevention, risk management and cybersecurity as strategic priorities across the banking sector [FEBRABAN 2025]. This combination of transaction scale, operational immediacy and increasingly sophisticated social engineering strategies has transformed fraud detection into not only a security problem, but also a critical economic and infrastructural challenge.

At the same time, adversarial behavior has evolved from isolated transactional anomalies into coordinated and adaptive processes involving statistical deviations or learned patterns. In this context, detection systems are required not only to identify suspicious events accurately, but also to operate under tight constraints involving latency, scalability, interpretability and operational cost. Traditional fraud detection approaches have predominantly framed the problem as either classification or anomaly detection, focusing on identifying deviations from expected transactional behavior through statistical inference, rule-based mechanisms or supervised learning techniques [Hilal et al. 2022, FBI 2026]. Although such approaches remain operationally relevant, they often emphasize the final manifestation of fraudulent activity rather than the broader sequence of interactions and behavioral signals that encompass fraud execution. As fraudulent operations increasingly exhibit coordinated and multi-stage characteristics, pure transaction-centric perspectives may become insufficient for capturing the structural and temporal dimensions of adversarial behavior [Chen et al. 2025].

Recent advances in machine learning have introduced more expressive models capable of capturing temporal and relational dependencies. Temporal signals derived from behavioral evolution over time, distributional signals apprehended from frequency space, structural signals extracted from interaction networks and contextual signals associated with heterogeneous transactional metadata are now frequently combined within integrated detection architectures [Lin et al. 2024, Motie and Raahemi 2024]. Furthermore, recent studies demonstrate the growing adoption of dynamic graph representations capable of jointly capturing temporal and structural signals in financial transaction networks [Saldaña-Ulloa et al. 2024], while hybrid architectures increasingly integrate contextual or semantic information to improve robustness against adaptive and coordinated fraud behavior [Wu et al. 2024]. These developments suggest a gradual transition from isolated signal approaches towards signal integration across contemporary fraud detection architectures.

However, this increase in model sophistication introduces new challenges. As models become more expressive, they also become more computationally costly, raising concerns about scalability, latency, energy consumption and reduced interpretability. Prior work has highlighted the growing environmental and economic implications of large-scale advanced machine learning systems, emphasizing the need to balance performance with efficiency and sustainability [Strubell et al. 2019, Schwartz et al. 2020, Castaño et al. 2026]. In the fraud detection field, where systems often operate continuously and under real-time constraints, these trade-offs become particularly critical. The challenge is therefore not limited to maximizing predictive performance, but also involves determining which manners of signal extraction, representation and integration provide the highest operational value relative to their computational, infrastructural and energetic cost.

This paper argues that fraud detection architectures can be more systematically understood under the lens of signals and their operationalization across different detection paradigms. Rather than analyzing methods solely according to algorithmic families, we propose a signal-aware perspective that emphasizes the types of information exploited by detection systems and the degree to which such signals are modeled explicitly or implicitly. Under this perspective, temporal occurrences, distributional

deviations, relational structures and contextual attributes can be interpretated as distinct but complementary signal families that shape the detection process and influence how fraud manifests across operational environments.

In addition to the proposed taxonomy, this work introduces a conceptual fraud lifecycle model that enables connection between observable signals and different stages of fraudulent activity. By connecting detection mechanisms to the progression of fraud operations, the proposed model provides a process-oriented interpretation of how distinct architectural paradigms capture adversarial behavior over time. This notion also enables a broader discussion regarding sustainability in fraud detection, particularly concerning energy expenditure, computational efficiency and operational robustness in continuously deployed systems. Rather than proposing a novel predictive architecture, this work focuses on providing an analytical and organizational framework for understanding and building fraud detection systems under operational and sustainability constraints

To guide our analysis and frame the contributions of this paper, we define three central research questions. Each question corresponds to a core dimension of our study: organizing fraud detection methods, evaluating their trade-offs, and considering sustainability in system design. These questions are presented as follows:

RQ1 – Signal Organization: How can existing fraud detection methods be systematically organized based on the signal families they exploit and the manner in which these signals are leveraged?

RQ2 – Trade-offs and Process Awareness: What are the key trade-offs between detection capability, process awareness and computational cost across different signal-aware detection strategies?

RQ3 – Sustainability and Practical Design: How can fraud detection systems balance predictive performance with efficiency, scalability, interpretability and real-time deployability?

2. Fraud Lifecycle and Signal Emergence

Financial fraud rarely manifests as an isolated event. Instead, fraudulent activity tipically unfolds as a progressive and adaptive process involving preparation, probing, execution, expansion and concealment strategies across transactional ecosystems. This perspective aligns with prior lifecycle-oriented approaches such as the Fraud Kill Chain, which adapts adversarial modeling concepts from cibersecurity to financially motivated attacks [Moloney and Driehuis 2023, Hutchins et al. , Strom et al. 2018]. Rather than treating fraud detection solely as a classification problem, lifecycle-oriented perspectives emphasize how observable evidence emerges throughout different operational stages of adversarial activity.

In practice, fraud operations often involve both long-term strategic preparation and short-term adaptive interaction with transactional systems. Motivated by this observation, the lifecycle adopted in this work distinguishes between two complementary reconnaissance dimensions. While strategic reconnaissance involves market intelligence gathering, infrastructure preparation, target selection and long term operational planning, operational reconnaissance, in contrast, refers to real-time probing activities such as transaction testing, behavioral exploration, credential validation or interaction

experimentation conducted during active fraud execution. This distinction reflects the increasingly adaptive nature of modern fraud ecosystems, where adversaries continuously adjust strategies based on system responses and detection pressure.

Unlike traditional intrusion-oriented threat models, infiltration in financial fraud is not always a discrete operational stage. Fraudsters frequently operate through legitimate interfaces, compromised credentials, social engineering or otherwise authorized transactional channels without necessarily bypassing technical access controls [Strom et al. 2018, Hilal et al. 2022]. Recent fraud analyses further emphasize that contemporary financial attacks increasingly exploit trusted operational workflows and human-mediated interactions rather than purely technical intrusion vectors [ACFE 2026, FBI 2026]. Consequently, infiltration is treated here as an optional or implicit operational condition associated with the establishment or abuse of interaction capabilities rather than as a mandatory stage preceding execution.

Figure 1 summarizes the proposed lifecycle model and illustrates the interaction between adversarial stages, reconnaissance activities, and iterative adaptation mechanisms throughout fraud operations. The framework provides the operational foundation upon which different observable signal manifestations may emerge across the fraud lifecycle.

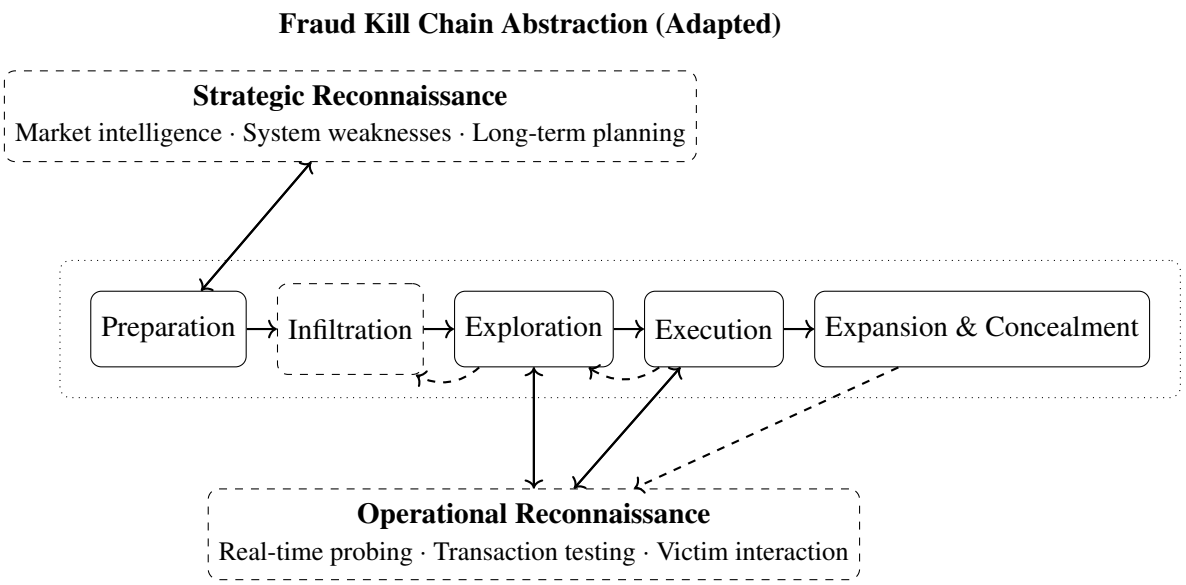

Figure 1. Revised fraud lifecycle model grounded in the Fraud Kill Chain. The framework extends prior models by distinguishing strategic and operational reconnaissance and explicitly modeling iterative feedback loops across stages.

A central implication of the lifecycle perspective is that observable fraud evidence emerges asymmetrically throughout different operational stages. Early stages such as reconnaissance and probing tend to produce weak, sparse or ambiguous signals, often expressed through contextual inconsistencies, low-intensity temporal patterns or subtle structural anomalies. Conversely, later stages such as execution, expansion and concealment typically generate stronger distributional

and structural patterns, including anomalous transaction distributions, coordinated interaction networks, repeated transfer structures or multi-account propagation behavior [Akoglu et al. 2015, Motie and Raahemi 2024]. Recent graph-based fraud studies similarly show that coordinated fraud schemes become increasingly observable through relational propagation patterns and dynamic interaction topology as attacks mature operationally [Lin et al. 2024, Saldaña-Ulloa et al. 2024]. Such asymmetry creates a fundamental challenge in fraud detection: the earlier a fraud operation is observed, the weaker and noisier its observable evidence tends to be.

In this work, signals are understood as fraud-relevant patterns emerging from transactional activities, relational interactions and operational context. Hence signals do not correspond to features, models or architectures themselves, but rather to forms of observable evidence that may later be operationalized computationally. Four dominant signal families are considered: temporal, distributional, structural and contextual. These categories are analytical rather than mutually exclusive and represent dominant forms through which fraud-relevant evidence becomes observable across financial systems [Lin et al. 2024, Wen 2024, Wu et al. 2024].

Under this perspective, temporal signals emerge through sequential or time-dependent activity patterns, such as transaction bursts, repeated retries, evolving interaction dynamics or adaptive probing behavior. Sequential anomaly detection methods and temporal representation learning approaches have demonstrated strong capability in capturing such evolving transactional dependencies [Malhotra et al. 2015, Zioviris et al. 2024, Ma et al. 2025]. Distributional signals represent deviations from expected statistical behavior, including anomalous transaction values, irregular frequency distributions or density-based outliers, remaining central to classical anomaly detection and probabilistic fraud modeling approaches [Bhattacharyya et al. 2011, Hilal et al. 2022]. Structural signals emerge from interaction topology and relational organization, apprehending coordinated activity, transaction propagation, shared infrastructure usage or collusive network structures. These signals are particularly relevant in graph-based fraud detection, where graph neural networks and graph transformers operationalize relational dependencies across accounts, devices, merchants and transactional flows [Akoglu et al. 2015, Motie and Raahemi 2024, Lin et al. 2024]. Finally, contextual signals incorporate auxiliary operational information associated with transactions, including device metadata, geolocation inconsistencies, access environments, communication channels or cross-platform interaction context. Recent multimodal and heterogeneous fraud detection architectures increasingly integrate such contextual evidence to improve robustness against adaptive and socially engineered fraud scenarios [Hilal et al. 2022, Wu et al. 2024].

Importantly, these signal families describe observable manifestations rather than specific modeling paradigms. A single fraud event may generate simultaneous temporal, distributional, structural and contextual evidence, while different detection architectures may operationalize the same signal family through distinct computational mechanisms. This distinction motivates a shift away from architecture-centric categorization towards a signal-aware perspective centered on how fraud-relevant patterns emerge, evolve and become observable across the fraud lifecycle. The relationship between fraud lifecycle stages and signal emergence also reveals important operational chokepoints. Certain stage

transitions between lifecycle stages naturally concentrate observable evidence, making them more vulnerable to detection or disruption. For example, operational probing may expose weak temporal irregularities prior to execution, while expansion stages frequently expose stronger structural evidence associated with coordination or laundering behavior. These chokepoints are particularly relevant because they connect adversarial process awareness with practical detection opportunities.

3. Signal-Aware Detection and Operational Trade-offs

Traditional fraud detection literature commonly categorizes methods according to predictive paradigms such as rule-based systems, statistical learning, deep learning, graph-based architectures or federated detection systems [Hilal et al. 2022, Motie and Raahemi 2024]. While useful from an implementation perspective, such classifications often emphasize model architecture more strongly than the forms of observable evidence being operationalized throughout the detection process. Recent graph-enhanced and multimodal fraud detection approaches further blur the boundaries between architectural categories by jointly integrating temporal, distributional, structural and contextual informationa within unified learning systems [Lin et al. 2024, Wen 2024, Wu et al. 2024]. Under the signal-aware perspective proposed in this work, detection methods are therefore analyzed according to the dominant signal families they operationalize and the computational trade-offs associated with their deployment.

Recent discussions in machine learning sustainability emphasize that predictive performance alone is insufficient for evaluating operational AI systems [Strubell et al. 2019, Schwartz et al. 2020, Ares-Robledo 2026]. Large-scale learning architectures frequently impose substantial computational costs during both training and inference, particularly when operating under real-time constraints or high-throughput environments. In fraud detection, this challenge becomes especially critical because systems must process massive transaction volumes continuously while maintaining strict latency requirements and high operational availability. Consequently, computational sustainability emerges not as a secondary optimization criterion, but as a central systems-level design constraint.

Under the signal-aware perspective proposed in this work, sustainability trade-offs can be interpreted as partially dependent on the forms of observable evidence being operationalized. Different signal families naturally require different forms of computational processing, storage organization, representation complexity, and inference orchestration. Distributional signals, for example, are frequently operationalized through lightweight statistical modeling or engineered feature evaluation, often allowing efficient inference with comparatively low computational overhead. In many cases, these approaches scale approximately linearly with transaction volume, particularly in streaming or threshold-based systems.

Signal operationalization may occur explicitly, implicitly or through hybrid representations. Explicit operationalization directly models a signal family through predefined metrics, statistical assumptions or interpretable relational structures. Examples include transaction velocity thresholds, probabilistic anomaly scoring, temporal sequence windows or graph connectivity metrics commonly employed in anomaly detection and graph-based fraud analysis [Akoglu et al. 2015, Bhattacharyya et al. 2011,

Hilal et al. 2022]. Implicit operationalization, in contrast, captures signals through learned latent representations without being individually specified during modeling. Deep neural architectures frequently operate in this manner, capturing temporal, contextual or relational regularities through representation learning rather than through manually defined signal construction [Hochreiter and Schmidhuber 1997, Zioviris et al. 2024]. Hybrid operationalization combines engineered signals with learned representations or jointly modeling multiple signal families across heterogeneous architectures, as increasingly observed in graph transformer and multimodal fraud detection systems [Lin et al. 2024, Wen 2024, Wu et al. 2024].

Temporal operationalization focuses on signals emerging from sequential activity evolution and is particularly relevant in scenarios involving probing behavior, transaction escalation, repeated retries or adaptive fraud progression. Classical sequence modeling approaches such as Long Short-Term Memory (LSTM) networks introduced mechanisms for capturing long-range dependencies in evolving transactional data [Hochreiter and Schmidhuber 1997]. More recent work extends this paradigm through transformer-based sequence modeling and temporal graph representations capable of jointly capturing temporal and relational dependencies across transaction streams [Saldaña-Ulloa et al. 2024, Zioviris et al. 2024, Ma et al. 2025]. While temporal operationalization substantially improves process awareness, it also introduces higher inference latency and increased streaming complexity in real-time environments.

Distributional operationalization focuses on deviations from expected transactional behavior and remain foundational across both classical and modern fraud detection systems. These signals are commonly operationalized explicitly through statistical modeling, probabilistic scoring, density estimation or anomaly detection mechanisms [Bhattacharyya et al. 2011, Hilal et al. 2022]. Traditional supervised learning methods, including logistic regression, decision trees, and ensemble approaches, frequently rely on distributional representations derived from engineered transaction statistics [Dal Pozzolo et al. 2015]. Despite their comparatively lower computational cost and strong interpretability, distributional approaches often exhibit limited sensitivity to coordinated or adaptive fraud patterns that extend beyond isolated statistical irregularities.

Structural operationalization models fraud through relational organization and interaction topology across transactional ecosystems. Rather than analyzing transactions independently, these approaches represent accounts, devices, merchants or users as interconnected entities whose relational dynamics may reveal coordinated fraud behavior. Graph-based anomaly detection and graph neural networks have become particularly important in this context due to their ability to capture collusive structures, laundering chains, shared infrastructure usage and multi-account propagation patterns [Akoglu et al. 2015, Motie and Raahemi 2024, Ares-Robledo 2026]. Recent graph transformer architectures further extend this paradigm by integrating relational propagation with attention mechanisms and heterogeneous feature aggregation [Lin et al. 2024]. However, these approaches typically introduce substantially higher computational and infrastructural demands associated with graph construction, iterative message propagation, dynamic topology maintenance and large-scale relational aggregation.

Contextual operationalization incorporates auxiliary evidence surrounding

transactions and interactions. Unlike distributional or structural patterns directly observable from transaction flows, contextual evidence frequently depends on heterogeneous metadata sources such as device characteristics, geolocation consistency, communication channels, access environments, authentication flows or cross-platform interaction patterns. Recent multimodal fraud detection systems increasingly integrate contextual representations with temporal and structural learning mechanisms to improve robustness against adaptive fraud and social engineering attacks [Wu et al. 2024, Wen 2024]. Nevertheless, contextual operationalization often introduces additional challenges involving data heterogeneity, integration complexity, explainability and privacy-sensitive infrastructure requirements.

Importantly, modern fraud detection architectures rarely operate on a single signal family in isolation; instead, recent research increasingly reflects a convergence towards hybrid operationalization strategies capable of integrating temporal, structural, distributional, and contextual evidence simultaneously [Lin et al. 2024, Motie and Raahemi 2024, Wen 2024]. This trend is particularly visible in graph-enhanced and multimodal architectures, where heterogeneous signal integration enables broader process awareness across different fraud lifecycle stages. However, the increased expressiveness provided by such systems is frequently accompanied by substantial computational trade-offs involving scalability, latency, energy expenditure and operational complexity.

Table 1 summarizes the dominant operational trade-offs associated with different signal families. The complexity expressions are intentionally approximate and represent dominant operational scaling behaviors commonly associated with signal extraction and inference rather than exact implementation-specific computational cost.

Table 1. Operational trade-offs across signal families. Complexity terms represent dominant operational scaling behaviors commonly associated with signal extraction and inference.

Signal Family	Dominant Operations	Typical Complexity	Primary Constraints	Typical Deployment Characteristics
Distributional	Threshold evaluation, statistical scoring, density estimation	$\mathcal{O}(n)$	Concept drift, limited relational awareness	Real-time screening, high-throughput systems
Temporal	Sequential state propagation, temporal attention, sequence modeling	$\mathcal{O}(nL)$ to $\mathcal{O}(L^2)$	Latency, sequence maintenance, streaming complexity	Sequential monitoring, adaptive fraud detection
Structural	Graph construction, neighborhood aggregation, message passing	$\mathcal{O}(\|V\| + \|E\|)$ per propagation layer	Graph propagation, topology maintenance, infrastructure cost	Coordinated fraud analysis, laundering detection
Contextual	Multimodal integration, metadata fusion, heterogeneous inference	Variable / architecture-dependent	Data integration, heterogeneity, governance overhead	Multimodal operational monitoring

An important implication of these trade-offs is that increased fraud process awareness frequently correlates with increased operational complexity. Architectures capable of capturing coordinated, adaptive, and multi-stage fraud patterns generally require richer representations and more computationally intensive inference pipelines. As a result, systems optimized exclusively for predictive performance may become economically or operationally impractical in real-world deployment scenarios, which may be particularly relevant in financial infrastructures operating under high transaction throughput and strict response deadlines, where inference delays may directly impact user experience, operational risk, and infrastructure expenditure.

Recent fraud detection research increasingly reflects attempts to balance expressiveness and efficiency through hybrid and layered architectures [Wen 2024, Zioviris et al. 2024]. In such systems, lightweight models are frequently employed for large-scale preliminary screening, while more computationally expensive graph-based or multimodal methods are selectively activated for high-risk cases requiring deeper investigation. This strategy enables more efficient resource allocation while preserving the ability to detect sophisticated coordinated fraud patterns. Similar directions also appear in recent work on efficient graph learning, sparse attention mechanisms, and scalable heterogeneous graph processing [Hamilton et al. 2017, Lin et al. 2024].

The lifecycle perspective introduced in this work suggests that fraud detection should not be interpreted as a static classification problem, but rather as a continuous process of observing evolving adversarial patterns across transactional environments. Different stages of fraudulent activity generate different forms of observable evidence, leading to important trade-offs between early detection capability, operational observability, computational cost, and deployment feasibility. Early-stage fraud activity, particularly during reconnaissance, probing, or low-intensity exploration, often produces weak and ambiguous signals. These patterns often manifest through subtle contextual inconsistencies, sparse temporal irregularities, or low-density interaction anomalies that are difficult to distinguish from legitimate user behavior. Similar observations have been reported in adversarial fraud modeling and behavioral anomaly detection literature, where early fraudulent operations exhibit low signal intensity and high uncertainty [Hilal et al. 2022, Akoglu et al. 2015, Strom et al. 2018]. Conversely, later stages such as execution, expansion, laundering, or coordinated propagation tend to generate stronger distributional and structural evidence, including anomalous transaction distributions, repeated interaction structures, coordinated account activity, and network-level propagation behavior [Motie and Raahemi 2024, Lin et al. 2024].

This asymmetry reveals a fundamental operational trade-off in fraud detection systems. Architectures optimized for early process awareness typically require richer contextualization, sequential monitoring, or relational modeling capabilities, which increase computational complexity and infrastructure requirements. In contrast, lightweight distributional approaches based on statistical deviations or transaction-level inference provide high scalability and low latency, but often remain limited to detecting fraud after stronger manifestations become observable. Recent graph-enhanced and temporal architectures partially mitigate this limitation by integrating heterogeneous signals across sequential and relational dimensions [Lin et al. 2024, Saldaña-Ulloa et al. 2024, Wu et al. 2024]. However, such systems also introduce

substantially higher operational costs associated with graph maintenance, message propagation, temporal synchronization, and high-dimensional representation learning.

These observations suggest that no single detection architecture is universally optimal across all fraud stages or operational environments. Instead, practical fraud detection increasingly relies on layered and adaptive system designs. In high-throughput financial infrastructures, lightweight distributional and contextual screening mechanisms may operate continuously as low-cost filtering layers, while more computationally intensive temporal or structural analysis is selectively applied to high-risk transactions, suspicious interaction subgraphs, or dynamically escalated cases. Similar multi-stage operational strategies have been discussed in sustainable AI and scalable fraud detection literature as mechanisms for balancing detection capability with real-time deployability and infrastructure efficiency [Schwartz et al. 2020, Zioviris et al. 2024, Castaño et al. 2026].

4. Discussion and Conclusion

This work proposed a signal-aware perspective for analyzing fraud detection systems under the combined lenses of process awareness and computational sustainability. Rather than organizing fraud detection methods exclusively according to predictive architectures, the proposed framework emphasizes how different forms of observable evidence emerge across fraud operations and how such signals are operationalized within detection systems.

The analysis suggests that contemporary fraud detection increasingly depends on the integration of heterogeneous signal families, particularly as fraud operations evolve toward adaptive, coordinated, and multi-stage behaviors. Temporal evolution, distributional deviations, structural relationships, and contextual inconsistencies frequently coexist within real-world transactional environments, motivating the growing adoption of hybrid and multimodal architectures. At the same time, increasing model expressiveness introduces important sustainability challenges involving computational complexity, latency, scalability, infrastructure requirements, and operational maintainability. These trade-offs indicate that fraud detection should be understood not only as a predictive task, but also as a systems engineering problem operating under real-world deployment constraints.

The proposed fraud lifecycle model further contributes by associating signal emergence with different operational stages of fraudulent activity. By distinguishing strategic and operational reconnaissance, incorporating iterative feedback behavior, and emphasizing lifecycle chokepoints, the framework provides a process-oriented interpretation of how observable evidence evolves throughout fraud operations. This perspective may support future efforts toward stage-aware detection, adaptive intervention strategies, and sustainability-oriented system design.

Nevertheless, several limitations should be acknowledged. The proposed taxonomy and lifecycle model remain primarily conceptual and analytical rather than empirically validated through large-scale benchmarking. The mapping between signal families, lifecycle stages and detection architectures is interpretative and may vary across application domains, institutional infrastructures, and fraud typologies. Furthermore, while computational sustainability is discussed through operational complexity and

deployment considerations, precise quantification of energy consumption, carbon footprint, or infrastructure expenditure remains outside the scope of this study.

Future research may extend this work through empirical validation of lifecycle-aware detection strategies, quantitative evaluation of signal operationalization costs, and development of benchmark environments capable of jointly evaluating predictive performance, scalability, latency, interpretability, and sustainability metrics. Additional opportunities include adaptive stage-aware architectures, multimodal graph-temporal systems, and operational frameworks capable of dynamically balancing lightweight screening with selective high-complexity analysis.

Ultimately, sustainable fraud detection requires moving beyond purely accuracy-centric evaluation towards integrated approaches that jointly consider adversarial process awareness, operational feasibility, and computational efficiency. By framing fraud detection through the emergence and operationalization of observable signals, this work contributes towards a more systematic and deployment-oriented understanding of modern fraud detection systems.

References

ACFE (2026). Antifraud technology benchmarking report. Technical report, ACFE.

Akoglu, L., Tong, H., and Koutra, D. (2015). Graph-based anomaly detection and description: a survey. *Data Mining and Knowledge Discovery*.

Ares-Robledo, F.; Rifà-Pous, H. C. R. (2026). Graph neural networks for anomaly detection: a systematic review of dynamic temporal approaches. *Artificial Intelligence Review*, 59(129).

Bhattacharyya, S., Jha, S., Tharakunnel, K., and Westland, J. C. (2011). Data mining for credit card fraud: A comparative study. *Decision Support Systems*, 50(3):602–613. On quantitative methods for detection of financial fraud.

Castaño, J., Bustillo, J., Franch, X., and Martínez-Fernández, S. (2026). Data-driven multi-objective optimization of ml inference hardware configurations for energy, performance and cost. *Sustainable Computing: Informatics and Systems*, 50:101347.

Chen, Y., Zhao, C., Xu, Y., Nie, C., and Zhang, Y. (2025). Deep learning in financial fraud detection: Innovations, challenges, and applications. *Data Science and Management*.

Dal Pozzolo, A., Caelen, O., Johnson, R., and Bontempi, G. (2015). Calibrating probability with undersampling for unbalanced classification.

FBI (2026). Internet crime report 2025. Technical report, Internet Crime Complaint Center - IC3.

FEBRABAN (2025). Pesquisa febraban de tecnologia bancária 2025. Technical report, FEBRABAN.

Hamilton, W. L., Ying, R., and Leskovec, J. (2017). Inductive representation learning on large graphs. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17, page 1025–1035, Red Hook, NY, USA. Curran Associates Inc.

Hilal, W. et al. (2022). Financial fraud: A review of anomaly detection techniques and recent advances. *IEEE Access*.

Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9:1735–1780.

Hutchins, E. M., Cloppert, M. J., and Amin, R. M. Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains. *Leading Issues in Information Warfare & Security Research*.

Lin, J., Guo, X., Zhu, Y., Mitchell, S., Altman, E., and Shun, J. (2024). Fraudgt: A simple, effective, and efficient graph transformer for financial fraud detection. In *Proceedings of the 5th ACM International Conference on AI in Finance*, ICAIF '24, page 292–300, New York, NY, USA. Association for Computing Machinery.

Ma, R., Ma, Y., and Liu, X. (2025). Time series anomaly detection via temporal relationship graphs and adaptive smoothing. *Applied Soft Computing*, 179:113298.

Malhotra, P., Vig, L., Shroff, G., and Agarwal, P. (2015). Long short term memory networks for anomaly detection in time series.

Moloney, S. and Driehuis, E. (2023). The fraud kill chain. `https://www.featurespace.com/newsroom/fraud-kill-chain`. Accessed: 2026-03-20.

Motie, S. and Raahemi, B. (2024). Financial fraud detection using graph neural networks: A systematic review. *Expert Systems with Applications*, 240:122156.

Saldaña-Ulloa, D., De Ita Luna, G., and Marcial-Romero, J. R. (2024). A temporal graph network algorithm for detecting fraudulent transactions on online payment platforms. *Algorithms*, 17(12).

Schwartz, R., Dodge, J., Smith, N. A., and Etzioni, O. (2020). Green ai. *Commun. ACM*, 63(12):54–63.

Strom, B. E. et al. (2018). Mitre att&ck: Design and philosophy. Technical report, MITRE Corporation.

Strubell, E., Ganesh, A., and McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. In Korhonen, A., Traum, D., and Màrquez, L., editors, *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3645–3650, Florence, Italy. Association for Computational Linguistics.

Wen, J., T. X.-L. J. (2024). An imbalanced learning method based on graph tran-smote for fraud detection. *Sci Rep*, 14.

Wu, J., Hu, R., Li, D., Ren, L., Hu, W., and Zang, Y. (2024). A gnn-based fraud detector with dual resistance to graph disassortativity and imbalance. *Information Sciences*, 669:120580.

Zioviris, G., Kolomvatsos, K., and Stamoulis, G. (2024). An intelligent sequential fraud detection model based on deep learning. *Journal of Supercomputing*, 80(10):14824 – 14847.