



From Rule-Based Detection to Adaptive Intelligence: A Critical Review of Artificial Intelligence Approaches to Financial Fraud Detection

Soraya Skavinski

Polícia Civil do Estado de São Paulo, São Paulo, Brazil

Email: soraya1br@gmail.com

How to cite this paper: Skavinski, S. (2026) From Rule-Based Detection to Adaptive Intelligence: A Critical Review of Artificial Intelligence Approaches to Financial Fraud Detection. *Open Access Library Journal*, **13**: e16066. <https://doi.org/10.4236/oalib.1116066>

Received: September 22, 2026

Accepted: October 20, 2026

Published: October 23, 2026

Copyright © 2026 by author(s) and Open Access Library Inc.

This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

<http://creativecommons.org/licenses/by/4.0/>



Open Access

Abstract

Financial fraud has evolved from isolated transactional anomalies into organized, adaptive networks that challenge traditional rule-based detection. Deterministic frameworks remain transparent but fail under dynamic adversarial conditions, extreme class imbalance, and rapid behavioral shifts. Recent advances in artificial intelligence (AI), machine learning (ML), deep learning (DL), graph analytics, federated learning (FL), and explainable AI (XAI) have transformed fraud prevention from retrospective analysis to continuous, adaptive risk assessment. This review synthesizes literature from 2020-2026 across five methodological paradigms: tree-based ensemble models for transactional scoring, sequence architectures for behavioral monitoring, graph neural networks for network-level fraud rings, federated learning for privacy-preserving collaboration, and explainable AI for regulatory compliance. It identifies persistent challenges including concept drift, false-positive costs, adversarial adaptation, and privacy risks. Evidence suggests that single-algorithm systems are insufficient for modern threat landscapes. Sustainable fraud prevention requires multi-layered, human-in-the-loop architectures integrating transactional, behavioral, relational, collaborative, and interpretive intelligence.

Subject Areas

Computational Finance, Machine Learning

Keywords

Financial Fraud Detection, Artificial Intelligence, Machine Learning, Deep Learning, Graph Neural Networks, Federated Learning, Explainable AI,

1. Introduction

Financial fraud presents a continuously shifting challenge for financial institutions, payment processing networks, regulatory bodies, and law enforcement agencies. Unlike traditional computational tasks where target parameters remain stable over time, financial crime operates within an active adversarial environment. Perpetrators dynamically modify their tactics in response to deployed security controls, technological shifts, and emerging payment infrastructure. Consequently, detection systems must operate under conditions where target behavior mutates precisely because detection mechanisms are active.

Historically, anti-fraud architectures relied on deterministic, rule-based logic—such as single-transaction velocity rules, geographic restrictions, hard monetary thresholds, and merchant category blocks. While these legacy mechanisms offer total transparency and predictable execution latency, they exhibit two structural flaws. First, rigid thresholds generate substantial volumes of false positives by misinterpreting benign, volatile consumer actions as high-risk events. Second, deterministic rules are strictly reactive: a novel fraud mechanism evades detection until human analysts recognize the pattern, measure its financial impact, and manually write a corresponding rule. These operational bottlenecks have driven the rapid adoption of adaptive machine learning, deep learning, graph analytics, and privacy-preserving distributed architectures [1] [2].

While AI models overcome rigid logic by identifying complex statistical relationships across streaming data, transitioning to probabilistic models introduces distinct engineering and operational challenges. Machine learning models depend on representative training data; however, fraud labels are sparse, leading to extreme class imbalance. Furthermore, constant shifts in criminal behavior induce rapid concept drift, degrading model performance over time. In addition, strict privacy regulations restrict cross-institutional data sharing, limiting an individual firm's visibility into broader fraud rings.

To address these challenges, this critical review examines the core question: *How are AI-based approaches transforming financial fraud detection, and what technical, operational, privacy, and investigative limitations continue to prevent fully adaptive fraud prevention?*

Operational Scope and Crime Taxonomy

To evaluate model performance accurately, financial fraud must be disaggregated into distinct operational domains, as each involves unique labels, underlying data structures, and operational goals:

- **Payment & Credit Card Fraud:** Focuses on card-present and card-not-present (CNP) transactions. Datasets consist of high-throughput tabular records. The

primary goal is real-time binary scoring under stringent sub-second latency constraints.

- **Account Takeover (ATO):** Focuses on unauthorized access to active accounts via credential stuffing, phishing, or session hijacking. Datasets rely on sequential event logs (authentication velocity, device fingerprints, IP shifts), where temporal changes dictate risk.
- **Money Laundering (AML):** Involves multi-stage transfers designed to integrate illicit funds into legal financial systems. Data structures consist of heterogeneous entity graphs (accounts, shell companies, intermediaries), evaluated through relational topologies rather than isolated events.
- **Coordinated Financial Crime & Fraud Rings:** Involves collusive networks sharing infrastructure (e.g., synthetic identities, shared devices, IP addresses). Operational objectives center on graph clustering and community detection.

2. Review Scope and Methodology

To ensure methodological rigor, transparency, and reproducibility, this literature review follows the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA 2020) guidelines.

2.1. Literature Search and Search String Formulation

A systematic search was conducted across five electronic literature databases: Web of Science, IEEE Xplore, ACM Digital Library, Google Scholar, and arXiv. The search targeted literature published between January 2020 and March 2026. To capture relevant technical implementations, the primary Boolean search query was structured as follows:

“financial fraud” OR “credit card fraud” OR “money laundering” OR “account takeover”) AND (“machine learning” OR “deep learning” OR “graph neural networks” OR “federated learning” OR “explainable AI”)

2.2. Screening Protocol and Eligibility Criteria

The initial database search yielded 842 records. Following the removal of 185 duplicate entries, 657 unique records underwent title and abstract screening. Record selection was governed by explicit, non-subjective eligibility criteria:

Inclusion Criteria (IC):

- **IC1 (Publication Standard):** Peer-reviewed journal articles, indexed conference proceedings, or seminal foundational preprints with verifiable citations.
- **IC2 (Domain Specificity):** Explicit technical focus on financial transaction scoring, anti-money laundering (AML), credit card fraud, or account takeover (ATO) architectures.
- **IC3 (Quantitative Evaluation):** Empirical reporting of model performance using benchmark metrics (e.g., PR-AUC, Precision@K, F1-score, or inference latency) on real-world or standard financial benchmark datasets.

Exclusion Criteria (EC):

- **EC1 (Domain Irrelevance):** Studies focused exclusively on non-financial cyber-attacks, general network intrusion detection systems (NIDS), or software vulnerability analysis.
- **EC2 (Methodological Ambiguity):** Papers lacking explicit mathematical formulations, baseline comparisons, evaluation metric definitions, or detailed dataset parameters.
- **EC3 (Language and Format):** Non-English publications or short extended abstracts (<4 pages) lacking detailed experimental validation.

During initial title/abstract screening, 512 records were excluded under criteria EC1 - EC3. The remaining 145 full-text articles were evaluated against full inclusion criteria. Of these, 135 articles were excluded primarily due to missing quantitative parameters or lack of experimental validation on financial datasets (EC2). Ultimately, 10 primary studies met all criteria and were selected for full synthesis.

2.3. Synthesis Matrix and Operational Mapping

Selected studies were mapped into five complementary operational pillars: 1) transactional tabular scoring, 2) sequential behavioral tracking, 3) relational graph analytics, 4) collaborative federated learning, and 5) post-hoc model interpretability engines.

3. The Evolution of Financial Fraud Detection

3.1. Legacy Rule-Based Architectures

Traditional fraud engines rely on hardcoded conditional logic. A transaction flags for manual review if it violates established heuristic limits—such as exceeding a specific transaction amount within a short time window. The primary operational advantage of rule-based engines is explainability. Investigators and auditors can trace the exact logic behind an alert. Furthermore, rule evaluations require low computational overhead, facilitating fast execution.

However, the primary vulnerability of rule engines is their reliance on historical insights. Rules must be manually crafted after a fraud pattern has caused measurable losses. Consequently, rule engines fail to detect zero-day fraud mechanisms and generate excessive false positives when benign customer actions mimic rigid risk thresholds.

3.2. Probabilistic Machine Learning Systems

Machine learning replaces hardcoded thresholds with statistical patterns learned from historical data. For structured transaction data, supervised learning algorithms analyze high-dimensional feature spaces simultaneously, assessing subtle interactions across numerical, categorical, and behavioral variables.

Supervised systems, however, introduce a critical dependency on labeled historical data. If historical training sets fail to capture emerging fraud tactics, the resulting models perpetuate systemic blind spots. This dynamic highlights a fundamental paradox in automated fraud prevention: machine learning models train

on historical fraud patterns, while criminal actors actively design novel techniques. Maintaining model efficacy therefore requires continuous monitoring, re-training, and dynamic feature re-engineering.

4. Major AI Approaches to Financial Fraud Detection




Figure 1. Layered multi-pillar AI architecture for financial fraud detection.

In **Figure 1**, the conceptual end-to-end processing pipeline illustrates the integration of tabular scoring, temporal behavioral analysis, relational graph analysis, collaborative federated learning, explainability engines, and human investigative decision-making (adapted from [2]-[4]).

4.1. Tree-Based Ensemble Models

Tree-based ensemble architectures—specifically Extreme Gradient Boosting

(XGBoost), Light Gradient Boosting Machine (LightGBM), and Random Forests—remain dominant for scoring structured tabular transaction data. These algorithms offer strong predictive performance on non-linear numerical/categorical distributions while maintaining low inference latency compatible with sub-second payment Authorization processing.




Figure 2. Federated learning architecture for cross-institutional fraud intelligence.

Figure 2 decentralized federated learning architecture depicting local parameter training across autonomous financial institutions and secure parameter aggregation without centralizing raw customer transaction records (adapted [5] [6]).

Despite their efficacy, tree-based models require continuous feature engineering to stay aligned with shifting behavioral patterns. Furthermore, they evaluate each transaction row independently, leaving them blind to complex multi-step temporal anomalies and relational networks.

4.2. Sequence Models and Behavioral Analytics

Fraudulent activity often unfolds across sequences of non-financial events rather than isolated transactions. Account takeover (ATO), for example, may involve credential stuffing, device changes, low-value test transfers, and eventual capital drains. Sequence models capture these temporal dependencies by conditioning risk scores on historical account behavior rather than evaluating events in isolation.

Long Short-Term Memory (LSTM) networks and Transformer-based architectures have demonstrated strong performance in modeling sequential fraud patterns. Jurgovsky *et al.* [7] applied Transformer encoders to authentication logs, showing improved detection of gradual account compromise compared to static classifiers. Similarly, Aljunaid *et al.* [6] integrated sequential features into federated learning pipelines, highlighting the importance of temporal context in collaborative fraud detection.

Despite their effectiveness, sequence models introduce operational challenges.

Training requires substantial computational resources, and real-time inference can strain production systems. Moreover, their deep representations are difficult to interpret directly, necessitating integration with explainable AI (XAI) frameworks to generate investigator-friendly reason codes. These trade-offs underscore the need for hybrid architectures where lightweight models handle immediate scoring, while sequence models provide deeper behavioral analysis in background pipelines.

4.3. Comparative Evaluation Metrics in Production

Assessing AI models in live financial systems requires metrics beyond standard offline classification accuracy or standard Area Under the ROC Curve (AUC-ROC).

- **Temporal Validation:** Standard k -fold cross-validation introduces temporal data leakage by evaluating historical events using models exposed to future data. Production architectures require time-based rolling split validation (e.g., training on week N , evaluating on week $N + 1$) to capture performance under concept drift.
- **Precision-Recall Performance & Top-K Metrics:** Given severe class imbalance, Precision-Recall AUC (PR-AUC) and Precision@K (precision within the top K highest-risk alerts generated daily) serve as primary operational benchmarks.
- **False-Positive Burden & Cost-Sensitive Thresholds:** Excessive false positives degrade user experience and strain investigative capacity. Models are evaluated using cost-sensitive decision thresholds that weigh false positives (customer friction, manual review costs) against false negatives (direct fraud losses, regulatory fines).

5. Graph Neural Networks and Financial Crime Networks

5.1. Relational Intelligence via Network Topologies

Financial crime frequently operates through networks of interconnected entities rather than isolated accounts. Graph-based frameworks model these relationships as heterogeneous graphs $G = (V, E)$, where nodes represent accounts, devices, or merchants, and edges represent interactions such as transfers or shared infrastructure. Graph Neural Networks (GNNs) generate embeddings that capture both local topology and node attributes, enabling detection of structural anomalies across shared digital ecosystems.

Recent studies illustrate the potential of GNNs in fraud detection. Di Gennaro *et al.* [8] employed Temporal Graph Attention Networks (TGAT) to identify money laundering rings, demonstrating that multi-hop relational analysis uncovers collusive structures invisible to row-based classifiers. Alhasawi *et al.* [5] further showed that federated GNNs can detect cross-institutional fraud rings while preserving data privacy. These findings validate the role of relational intelligence in exposing synthetic identity networks and coordinated fraud rings.

Operational deployment, however, remains challenging. Dynamic financial graphs expand rapidly, and computing embeddings in real time requires significant memory and processing power. Hybrid architectures are therefore essential: lightweight models provide immediate scoring, while GNNs operate asynchronously to update relational risk profiles. This layered approach balances scalability with analytical depth, ensuring fraud rings are detected without compromising transaction latency.

5.2. Exposing Collusion and Money Laundering

GNNs excel at exposing collusive networks, synthetic identity rings, and complex money laundering schemes that evade row-based classifiers. A single account may appear benign when evaluated independently; however, when its multi-hop connections reveal shared devices, IP addresses, or cyclic transaction paths with known high-risk clusters, the overall network topology indicates elevated risk.

Despite their analytical power, dynamic graph frameworks face operational challenges in high-throughput environments. Financial graphs expand rapidly, and calculating real-time node embeddings as new edges arrive requires significant memory and computational resources. Consequently, deploying GNNs often requires hybrid architectures where lightweight models handle real-time scoring, while dynamic GNNs process background streaming updates for relational analysis.

6. Federated Learning and Cross-Institutional Intelligence

Financial crime operates across institutional borders. Perpetrators frequently execute coordinated attacks against multiple banks, payment gateways, and fintech platforms. However, strict privacy regulations (e.g., GDPR, LGPD), competitive sensitivities, and bank secrecy laws prohibit financial institutions from centralizing raw customer records into shared external repositories.

Federated Learning (FL) addresses this constraint by enabling collaborative model training without centralizing sensitive underlying transaction datasets. In a federated setup, participating institutions train local models on their private data infrastructure. Local model parameter updates (gradients or weights) are transmitted to a central coordinator, aggregated (e.g., via Federated Averaging), and redistributed back to participants.

In **Figure 3**, a continuous active learning feedback loop demonstrating how explainable risk scores guide human investigators during case triage, whose final determinations generate verified labels to retrain models against emerging concept drift (adapted from [1] [9] [10]).

However, FL introduces its own security challenges. Parameter updates can be vulnerable to gradient-inversion attacks, allowing adversaries to reconstruct underlying private training data. Furthermore, compromised participating nodes can launch model poisoning attacks by sending malicious parameter updates designed to degrade global model accuracy or insert deliberate evasion backdoors.

Deploying FL in financial environments requires robust security protocols, including Differential Privacy (DP) noise injection to prevent gradient leaks, Secure Aggregation (SecAgg) cryptographic protocols, and Byzantine-tolerant aggregation algorithms to isolate malicious updates.




Figure 3. Operational human-in-the-loop intelligence and active learning feedback Loop.

7. Explainable Artificial Intelligence (XAI) and Governance

7.1. Regulatory Compliance and Investigative Transparency

As financial institutions deploy deep, complex machine learning architectures, model interpretability becomes essential. In many jurisdictions, automated financial decisions that block transactions, decline credit, or freeze accounts are subject to regulatory oversight requiring clear reason codes.

Explainable AI (XAI) frameworks—specifically SHapley Additive exPlanations (SHAP) and Local Interpretable Model-agnostic Explanations (LIME)—bridge the gap between complex model outputs and operational requirements. XAI mechanisms provide feature-level attribution scores that explain why a specific transaction or account triggered a high risk score. These outputs generate actionable reason codes for human investigators, accelerate alert triage, support regulatory audits, and ensure compliance with anti-discrimination mandates.

7.2. The Latency-Interpretability Trade-Off

Incorporating post-hoc explainability techniques introduces computational overhead. Calculating exact SHAP values for complex ensemble models or deep networks during real-time transaction scoring can introduce processing latency in-

compatible with sub-second payment approval SLAs.

To balance this trade-off, modern anti-fraud architectures often separate real-time scoring from explanation generation. Real-time engines leverage fast tree ensembles or pre-computed models to assign immediate risk scores, while background XAI pipelines generate feature attributions only when alerts are escalated to human investigative queues.

8. Persistent Operational Challenges

8.1. Extreme Class Imbalance and Data Reality

Financial fraud detection operates under severe class imbalance, where fraudulent transactions represent only a tiny fraction of overall activity. In the widely referenced European credit card dataset, fraudulent transactions account for approximately **0.172% of total volume** [9]. Broader retail banking and interbank flows report fraud prevalence between **0.01% and 0.1%**, based on operational estimates from Thivaos [2] and Yang [3]. These figures highlight the rarity of fraud events relative to legitimate transactions, though precise rates vary across institutions, geographies, and time periods.

Such imbalance creates a fundamental challenge: standard supervised algorithms tend to over-predict the majority class, leading to missed fraud cases. Techniques such as Synthetic Minority Over-sampling Technique (SMOTE), random undersampling, and focal loss functions can partially rebalance training distributions. However, these methods do not resolve the deeper issue that historical labels are incomplete proxies for future fraud patterns. Moreover, reported prevalence rates often lack confidence intervals or denominators, underscoring the uncertainty inherent in fraud measurement. Effective fraud detection therefore requires architectures that combine statistical learning with anomaly detection and human-in-the-loop validation to mitigate the risks of imbalance and incomplete labeling.

8.2. Emerging Fraud, Adversarial Adaptation, and Concept Drift

Financial fraud prevention functions as an active adversarial domain. When perpetrators notice that specific transaction amounts, velocity patterns, or locations trigger automated security blocks, they adapt their tactics. This ongoing adaptation produces severe concept drift, causing statistical relationships learned from historical data to degrade.

Addressing novel zero-day fraud mechanisms requires moving beyond supervised models on historical labels. While unsupervised anomaly detection and graph frameworks can identify deviations from normal behavior, increasing system sensitivity to novel anomalies can elevate false-positive rates, as legitimate consumer actions (such as travel or sudden purchases) often mimic abnormal profiles. Modern defense systems manage this trade-off by using anomaly alerts to dynamically adjust risk scores and prioritize cases for human review, rather than relying solely on automated blocks.

9. Human-in-the-Loop Fraud Intelligence

Automated AI models are most effective when integrated with human expertise. Rather than replacing human investigators, AI models serve as intelligence filters that optimize how investigative attention is allocated.

An operational human-in-the-loop framework operates as a continuous intelligence loop:

Data → Detection → Risk Scoring → XAI Reason Codes → Investigative Review → Feedback → Retraining

Table 1. Lifecycle mapping of the adaptive financial fraud detection pipeline.

Phase	Pipeline Role	Primary Technology/Method
Data → Detection	Automated initial screening	Data pipelines & ML/DL models
Risk Scoring → XAI	Decision support & transparency	Risk calibration & SHAP/LIME
Review → Feedback	Operational validation	Analyst interface (Human-in-the-loop)
Retraining → Data	Continuous adaptive learning	MLOps & Supervised learning update

Table 1 details the end-to-end lifecycle of an adaptive financial fraud detection system, illustrating how raw transactional data translates into continuous model improvement through Explainable AI (XAI) and human expertise (synthesized from [1] [4] [7] [9] [11]).

Human investigators bring contextual reasoning to ambiguous, edge-case alerts generated by AI systems. When an investigator confirms or dismisses a complex alert, that decision generates a high-quality ground-truth label. Feeding these human-verified labels back into retraining pipelines helps update models against emerging fraud patterns, stabilizing system performance over time.

10. Integrated Anti-Fraud Architectures

To address multi-faceted attack vectors, modern fraud prevention frameworks combine multiple complementary analytical layers into a unified defense system:

- **Layer 1—Real-Time Transaction Scoring:** High-speed tree-based ensemble models evaluate individual structured transactions under tight sub-second latency constraints.
- **Layer 2—Sequential Behavioral Analysis:** Sequence architectures track account event histories to detect account takeovers and abnormal behavioral changes.
- **Layer 3—Relational & Graph Intelligence:** Dynamic Graph Neural Networks analyze multi-hop relationships across devices, entities, and accounts to uncover collusive fraud rings.
- **Layer 4—Cross-Institutional Collaboration:** Federated Learning frameworks enable privacy-preserving model updates across participating institutions.
- **Layer 5—Explainability Engine:** XAI modules generate feature attributions and reason codes to explain automated risk scores.
- **Layer 6—Human Triage & Active Learning:** Investigators review escalated

alerts, make final determinations, and feed verified outcome labels back into training pipelines.

11. Regulatory, Governance, and Future Directions

The deployment of AI in financial systems must align with international regulatory frameworks such as the Financial Action Task Force (FATF) guidance, the EU Payment Services Directive (PSD2/PSD3), and emerging AI governance standards. These frameworks emphasize model transparency, auditability, and protection against algorithmic bias, underscoring the need for robust governance structures.

Looking forward, four research directions appear most critical. First, **continuous active learning** pipelines must be developed to integrate investigator feedback while filtering out noisy or adversarial labels. Second, **scalable dynamic graph processing** is required to compute embeddings efficiently across massive, high-throughput financial networks. Third, **low-latency explainability** remains a pressing challenge, as institutions must balance regulatory demands for interpretability with operational constraints on transaction speed. Finally, **Byzantine-robust federated networks** will be essential to secure collaborative intelligence sharing against poisoning and inversion attacks.

Together, these directions highlight the dual imperative of technical innovation and regulatory compliance. Future fraud detection systems must not only advance algorithmic sophistication but also embed resilience, transparency, and trustworthiness at their core.

12. Conclusions

Artificial intelligence has transformed financial fraud detection, enabling a shift from rigid rule engines to adaptive systems capable of analyzing complex transactional, temporal, relational, and collaborative data streams. Tree-based ensembles excel at low-latency tabular scoring, sequence models track temporal behavior, graph neural networks uncover collusive networks, federated learning enables private cross-institutional collaboration, and explainable AI provides necessary transparency for governance and review.

However, operational challenges remain. Extreme class imbalance, concept drift, adversarial adaptation, dynamic graph scaling, and parameter-poisoning risks continue to complicate automated defense efforts. No single algorithmic approach addresses all operational requirements. Sustainable financial crime prevention depends on integrated architectures that combine transactional, behavioral, relational, collaborative, and explainable AI paradigms with human investigative expertise. By structuring AI as an intelligence-amplification layer within a human-in-the-loop framework, financial institutions can build resilient systems capable of adapting to evolving financial crime patterns.

Conflicts of Interest

The author declares no conflicts of interest.

References

- [1] Canhoto, A.I. (2021) Leveraging Machine Learning in the Global Fight against Money Laundering and Terrorism Financing: An Affordances Perspective. *Journal of Business Research*, **131**, 441-452. <https://doi.org/10.1016/j.jbusres.2020.10.012>
- [2] Thivaivos, S., Kostopoulos, G., Stefani, A. and Kotsiantis, S. (2026) A Survey of Machine Learning and Deep Learning for Financial Fraud Detection: Architectures, Data Modalities, and Real-World Deployment Challenges. *Algorithms*, **19**, Article 354. <https://doi.org/10.3390/a19050354>
- [3] Yang, H., Shukur, Z. and Sahran, S. (2026) A Review of Artificial Intelligence for Financial Fraud Detection. *Applied Sciences*, **16**, Article 1931. <https://doi.org/10.3390/app16041931>
- [4] Zhou, Y., Li, H., Xiao, Z. and Qiu, J. (2023) A User-Centered Explainable Artificial Intelligence Approach for Financial Fraud Detection. *Finance Research Letters*, **58**, 104309. <https://doi.org/10.1016/j.frl.2023.104309>
- [5] Alhasawi, Y., Almtrf, A.A. and Asad, M. (2025) A Federated Approach to Scalable and Trustworthy Financial Fraud Detection. *Security and Privacy*, **8**, e70099. <https://doi.org/10.1002/spy2.70099>
- [6] Aljunaid, S.K., Almheiri, S.J., Dawood, H. and Khan, M.A. (2025) Secure and Transparent Banking: Explainable AI-Driven Federated Learning Model for Financial Fraud Detection. *Journal of Risk and Financial Management*, **18**, Article 179. <https://doi.org/10.3390/jrfm18040179>
- [7] Jurgovsky, J., Granitzer, M., Ziegler, K., Calabretto, S., Portier, P., He-Guelton, L., *et al.* (2018) Sequence Classification for Credit-Card Fraud Detection. *Expert Systems with Applications*, **100**, 234-245. <https://doi.org/10.1016/j.eswa.2018.01.037>
- [8] Di Gennaro, M., Panebianco, F., Pianta, M., Zanero, S. and Carminati, M. (2024) Amatriciana: Exploiting Temporal Gnn's for Robust and Efficient Money Laundering Detection. 2024 *IEEE International Conference on Data Mining Workshops (ICDMW)*, Abu Dhabi, 9 December 2024, 254-261. <https://doi.org/10.1109/icdmw65004.2024.00039>
- [9] Pozzolo, A.D., Caelen, O., Johnson, R.A. and Bontempi, G. (2015) Calibrating Probability with Undersampling for Unbalanced Classification. 2015 *IEEE Symposium Series on Computational Intelligence*, Cape Town, 7-10 December 2015, 159-166. <https://doi.org/10.1109/ssci.2015.33>
- [10] Psychoula, I., Gutmann, A., Mainali, P., Lee, S.H., Dunphy, P. and Petitcolas, F. (2021) Explainable Machine Learning for Fraud Detection. *Computer*, **54**, 49-59. <https://doi.org/10.1109/mc.2021.3081249>
- [11] Purwar, A. and Manju, M. (2023) Credit Card Fraud Detection Using XGBoost for Imbalanced Data Set. *Proceedings of the 2023 Fifteenth International Conference on Contemporary Computing*, Noida, 3-5 August 2023, 216-219. <https://doi.org/10.1145/3607947.3607986>