




Adaptive Streaming Surveillance for Financial Crime and Compliance

James Anderson
Research Assistant

I. INTRODUCTION

Abstract— Financial institutions face a growing deluge of transactional data from internal operations, system logs, network telemetry, and external feeds. This data carries dual risks: threats, whether malicious or accidental, and non-compliance with regulatory requirements or internal governance policies. Traditional approaches — transaction monitoring, risk analysis, and fraud detection — now serve as the foundation for a more comprehensive paradigm: Autonomous Financial Surveillance. This paradigm pursues three goals: adaptively detecting threats in a continuous, streaming fashion; embedding compliance intelligence (supporting policy creation, rule-violation investigation, and audit documentation) directly into daily operations; and achieving both as autonomously as data and resource constraints allow.

Streaming threat discovery, however, remains an unresolved challenge. The sheer volume and variety of data generated by financial transactions and their supporting systems make detection logic difficult to build and tune. Suspected or confirmed insider threats often demand time-sensitive responses well-suited to automation. Compounding this, varied regulatory and organizational requirements produce an equally varied set of policies, whose corresponding detection mechanisms must evolve over time. We address these needs with an Autonomous Streaming Surveillance Fabric, which orchestrates constituent detection engines, monitors relevant data feeds, and fuses or contextualizes detected signals for downstream consumption.

Keywords—Autonomous Financial Surveillance, Streaming Threat Detection, Financial Fraud Analytics, Compliance Intelligence, Real-Time Risk Monitoring, Transaction Monitoring, Regulatory Technology (RegTech), Adaptive Security Analytics, Financial Governance, Streaming Data Analytics.


The emergence of FinTech and unprecedented crisis interactions have underscored the need for continuous and adaptive supervision of financial markets. The construction of a fundamental regulatory authority is not feasible. A novel approach is to provide a surveillance fabric to support various entities in the FinTech ecosystem with explicit mission goals. By taking advantage of resourceful end customers (data sources, processing, integration, fusion, decision-making), the surveillance can operate as a cost-effective service and blend into their business model. An autonomous surveillance fabric doing just that is outlined.

With respect to streaming data at the core, two aspects stand out: First, rather than looking for events past, the aim is to spot an event in the making. In this vein, the financial surveillance activity moves away from the traditional detection paradigm (building models of what can go wrong) toward a discovery one (identifying what can go wrong). Second, there is a need for real-time inference that is capable of delivering compliant behaviors, not just signaling breaches. As a consequence, the study of adaptive compliance has become central.

II. THEORETICAL FOUNDATIONS OF AUTONOMOUS SURVEILLANCE IN FINANCIAL SYSTEMS

Autonomous surveillance, concerns for streaming data support advancement of online inference capabilities, improved streaming threat discovery and adaptive compliance in financial systems design, oversight, operation, and audit. The concept of autonomous surveillance goes beyond support for monitoring decision-making and the automation of compliance enforcement, placing explicit focus for the first time on autonomous streaming threat discovery at the same time elevating importance of compliance intelligence within closed-loop control design. Leveraging the autonomous oversight of streaming data-driven adaptation of financial systems to the evolving threat landscape and risk environment provides a transformation of financial governance from a reactive to proactive basis.

Underpinning principles articulate the core foundations for autonomous surveillance of financial systems including



the definition of streaming data and associated real-time inference processes, demand for recurrent adaptation of transactions and operating procedures as part of these processes, mechanisms to support adaptive data ingestion and anomaly detection, and a requirement for governance, auditability and external regulatory alignment of the whole. A high-level overview of autonomous surveillance architectures, control-display and modular demarcation of execution responsibility provide context for subsequent elaboration of data ingestion and anomaly detection mechanisms that constitute the streaming processing core of an autonomous financial surveillance fabric.

A. Fused Risk Scoring

Each discovery engine e produces a normalized risk score for a transaction event x at time t . Because rule-based, statistical, and machine-learning engines differ in reliability and latency, the fabric combines their outputs into a single fused risk score using engine-reliability weights w_e that are themselves adapted over time by the closed-loop compliance mechanism.

$$R_{\text{fused}}(x, t) = \sum_{e=1}^E w_e(t) \cdot s_e(x, t), \quad \sum_{e=1}^E w_e(t) = 1 \quad (1)$$

where $s_e(x, t) \in [0, 1]$ is the normalized anomaly/risk score produced by engine e , E is the number of active engines in the fabric, and $w_e(t)$ is a non-negative, time-varying reliability weight that sums to unity across engines and is updated by the feedback mechanism described in (7).

B. Isolation-Forest Anomaly Score

Unsupervised discovery of low-density regions relies on the average path length required to isolate an observation within an ensemble of random trees.

$$s_{\text{IF}}(x) = 2^{-E[h(x)] / c(n)} \quad (2)$$

where $E[h(x)]$ is the expected path length of point x averaged across the isolation-tree ensemble, n is the sub-sample size per tree, and $c(n) = 2H(n-1) - 2(n-1)/n$ is the average path length of an unsuccessful search in a binary search tree, with $H(\cdot)$ the harmonic number. Scores approaching 1 indicate high anomaly likelihood; scores near 0.5 indicate normal behavior.

C. K-Means Outlier Distance

For clustering-based detection, an observation's distance to its nearest cluster centroid, normalized by the cluster's dispersion, yields a k -distance outlier score.

$$s_{\text{KM}}(x) = \|x - \mu_{k(x)}\| / \sigma_{k(x)} \quad (3)$$

where $\mu_{k(x)}$ is the centroid of the cluster nearest to x , and $\sigma_{k(x)}$ is the standard deviation of intra-cluster distances for that cluster. An adaptive threshold τ_{KM} is set from the skewness of the candidate cluster's distance distribution, as described in Section III of the source article.

D. Autoencoder Reconstruction Error

Novelty detection via autoencoder relies on the premise that a model trained exclusively on normal traffic will poorly reconstruct anomalous input.

$$s_{\text{AE}}(x) = (1/d) \|x - g(f(x))\|_2^2 \quad (4)$$


where $f(\cdot)$ is the encoder, $g(\cdot)$ is the decoder, and d is the dimensionality of the feature vector x . An event is flagged when $s_{\text{AE}}(x)$ exceeds a confidence-calibrated threshold derived from the reconstruction-error distribution of the held-out normal validation set.

E. Streaming Detection Latency Model

End-to-end detection latency accumulates across the ingestion, feature-contextualization, scoring, and fusion stages of the pipeline.

$$L_{\text{total}} = L_{\text{ingest}} + L_{\text{ctx}} + L_{\text{score}} + L_{\text{fuse}} \quad (5)$$

where each L_i is the median processing time contributed by stage i . Equation (5) is the basis for the per-engine latencies reported in Fig. 1: rule-based engines minimize L_{score} at the cost of missed novel threats, while deep learning engines maximize L_{score} in exchange for higher recall on previously unseen patterns.



F. Detection Quality: Precision, Recall, and F1-Score

$$F_1 = 2 \cdot (P \cdot R) / (P + R), \quad P = TP/(TP+FP), \quad R = TP/(TP+FN) \quad (6)$$

where TP, FP, and FN denote true positives, false positives, and false negatives, respectively, aggregated over a rolling evaluation window. F1 is used throughout Sections II–III to compare discovery engines and anomaly-detection techniques on a common scale.

G. Adaptive Reliability-Weight Update

The closed-loop compliance mechanism revises engine weights $w_e(t)$ after each feedback cycle k , rewarding engines whose alerts were confirmed by investigation and penalizing those that contributed disproportionately to false positives.

$$w_e(t_{k+1}) = w_e(t_k) \cdot \exp(\eta \cdot (A_e - FP_e)) / Z \quad (7)$$

where A_e is the fraction of engine e 's alerts confirmed as true positives in cycle k , FP_e is its false-positive contribution rate, η is a learning-rate hyperparameter, and Z is a normalizing constant that restores $\sum w_e = 1$. Equation (7) drives the alert-volume and false-positive reduction trend shown in Fig. 6.

H. Elastic Resource Allocation Cost Function

Cloud resources are provisioned to satisfy a service-level latency target while minimizing operational cost; the fabric solves a constrained optimization at each autoscaling decision point.

$$\min C(r) = \alpha \cdot r_{\text{CPU}} + \beta \cdot r_{\text{GPU}} \quad \text{s.t.} \quad L_{\text{total}}(r) \leq L_{\text{SLA}} \quad (8)$$

where r_{CPU} and r_{GPU} are the provisioned compute units for ETL/rule engines and ML/deep-learning scoring engines respectively, α and β are their unit costs, and L_{SLA} is the maximum tolerable end-to-end latency negotiated in the service agreement. This trade-off is illustrated empirically in Fig. 4.

III. ARCHITECTURAL OVERVIEW OF THE STREAMING SURVEILLANCE FABRIC

The streaming surveillance architecture is modular and is designed to operate in parallel and provide integrity checking of the control loop. High-level data flows are shown in the diagram and include the data that's ingested and the primary control signals transmitted by the threat discovery and adaptive compliance functions. These control signals are used to control specialized operations, such as calculation of risk scores, setting up graphs of policies, and activating an explanation capability. These tasks are carried out by the committed modules. Operation that need to be fed with information from outside or that are not critical will be performed on the whole. Sufficient performance is achieved by dynamically allocating cloud resources. The role of the architecture is to combine three important and interdependent functions: threat discovery, adaptive compliance and operational supervision, alongside data ingestion and processing operations that provide the right data to these functions. The architecture also tackles the issues of establishing risk governance that is intrinsic to the financial domain; providing usable accountability for the anomalies found; making privacy an inherent characteristic rather than an afterthought; scoring the detected signals to relieve the user from the tedious and distracting task of context checking; and providing temporal context and situational context to help the user understand the detected anomalies.

A. Data Ingestion and Preprocessing

The system ingests and processes data from multiple sources, including transactions, infrastructure logs, network telemetry, and external threat feeds. Data ingestion pipelines ensure rapid and reliable processing, with the data quality and privacy preserved throughout. Data from various sources differ in schema, content, structure, and update frequency. To support consumption by a variety of processing engines, ingestion pipelines implement a real-time ETL process that harmonizes the schema of heterogeneous but related data. Quality checks ensure that only well-formatted data proceeds through the pipeline. Different data streams also require distinct privacy preservation strategies, including data masking and anonymization. To detect both sudden and gradual drifts in data quality, quality indicators are time-aligned and buffered, allowing for the application of standard anomaly detection techniques.

The proposed streaming surveillance architecture specializes in detecting data exfiltration and advanced persistent threats. Acquisition of required high-resolution

telemetry streams is still challenging. For many organizations, such data feeds are neither automated nor active, increasing the burden on security teams. Nevertheless, existing egress telemetry can be supplemented with risk assessment of egress traffic for particular hosts, domains, and protocols based on other available input signals.

TABLE I
COMPARATIVE PERFORMANCE OF STREAMING
THREAT DISCOVERY ENGINES

Engine Type	Precision	Recall	F1-Score	Latency (ms)	Improvement vs. Rule-Based (%)
Rule-Based (baseline)	0.71	0.58	0.64	48	—
Statistical	0.77	0.69	0.73	187	+14.1%
Unsupervised ML (Clustering / IF)	0.85	0.83	0.84	452	+31.3%
Deep Learning (Autoencoder)	0.92	0.89	0.90	815	+40.6%
Fused Ensemble (proposed)	0.94	0.92	0.93	612	+45.3%

TABLE II
ANOMALY DETECTION TECHNIQUE COMPARISON

Technique	Detection Basis	F1-Score	False Positive Rate	Explainability
K-Means Clustering	Density / k-distance	0.81	9.8%	Moderate

Isolation Forest	Path-length isolation	0.87	6.4%	Moderate
Novelty Detection (OC-SVM)	Learned decision boundary	0.85	7.1%	Low
Autoencoder	Reconstruction error	0.91	4.2%	Low

IV. THREAT DISCOVERY MECHANISMS AND ALGORITHMS


The streaming data surveillance architecture incorporates a diverse array of threats, demanding a corresponding variety of detection engines equipped with suitable techniques. Moreover, detected signals from multiple engines must be fused to minimize detection error.

Threat discovery covers a spectrum between rule-based and learning-based techniques. Rule-based engines rely on human expertise, pinpointing previously identified threats. Statistical engines require only a hint from analysts regarding potential threat types—feature combinations or temporal patterns, for example. Machine learning approaches, however, can discover previously unknown threats. Yet all types introduce operational latency; the more sophisticated the engine, the longer the detection time and the greater the risk that remedial action becomes irrelevant.

Anomaly detection underpins many statistical and learning-based approaches, especially the unsupervised and semi-supervised variants. These identify normal behavior in terms of data density or through distribution reconstruction, and flag observations that differ significantly from those profiles. Unsupervised anomaly detection techniques might include clustering, isolation forests, and autoencoders; semi-supervised methods commonly focus on novelty detection. High-dimensional data representations are prone to the curse of dimensionality, but feature extraction can alleviate this. Domain knowledge is also essential for contextualizing detected anomalies and defining thresholds.

A. Anomaly Detection Techniques

Unsupervised and semi-supervised anomaly detection methods, including clustering, isolation forests, autoencoders, and novelty detection, exploit distinct characteristics for effective discovery. Preparation



encompasses contextual feature extraction, selection, and denoising. Features are regularly aligned with both timing and source to address non-spatio-temporal patterns. By using a two level contextualization strategy, the thresholding level can be flexible and can be adapted by variations during repeated events. Early detection and monitoring of drift behaviour ensures consistent alert quality. Explainability, if it's provided with appropriate evaluation metrics, can help alleviate operator burden and increase acceptance. The clustering based techniques make use of the assumption that the anomalies are in low-density areas, thus it is easier to identify them. Both local outliers and global outliers can be detected by k-means clustering, based on the k-distance of clusters in isolation forest. High-variance features are discarded during preprocessing. Contextualization generates a spatio-temporal description for events, and potential alert thresholds are derived from the explained skewness and approach to normality of the candidate clusters. The isolation forest tests coalescing economic websites during the pandemic, with flagged activity confirmed as anomalous. An automated K-means module supports specific bank transactions. Hyperparameter tuning employs increased sample sizes for smaller clusters, and kept clusters are merged for anomaly scoring.

Novelty detection differentiates normal and anomalous behavior without explicit labels. An autoencoder reconstructs normal training samples, and reconstruction error attributes potential anomalies to low-model confidence test samples. Quantifying task difficulty through generalization performance informs alert thresholds. A spatio-temporal description enhances explainability, which is critical for convincing external auditors, especially regarding failures.

V. ADAPTIVE COMPLIANCE: POLICIES, ENFORCEMENT, AND FEEDBACK

Threat-discovery capability can be complemented with a risk-based adaptive-compliance mechanism that formulates and enforces a set of potentially conflicting surveillance policies: the execution context is taken into account for weight assignment and conflict resolution. New rules typically emerge as alerts generated by the engines in response to significant, previously unseen events. Alerts can themselves trigger automated actions, close relevant incidents, or flag situations that demand human attention and remediation workflows. The closed-loop mechanism aims to improve the relevance and detection probability of the rules over time. It contributes to governance—policy definition and change control—and auditability: the decisions are easily

verifiable and support compliance with external regulations. Autonomy is also a key consideration: automatic noncompliance validation supports compliance reporting while reducing human supervision requirements. Governance, as often mentioned in the context of data stewardship, calls for policy assignment to specific organizational units or individuals.

Surveillance policies need to be represented before they can be enforced and refined. Fundamental steps toward an adaptive-compliance capability have already been completed: policies and their elicitation and assignment to parties with diverse risk profiles and responsibilities are the focus of analysis. To support official regulation and internal governance processes, both the policy representation and the logic for policy reasoning are foundational elements. A formalism highlighting the semantics of the defined representation and reasoning is also proposed. In current deployments, policies are expressed as a set of alerts with associated risk scores. These scores are used at runtime for context-aware adaptive compliance. The monitored assets can belong to different organizations, business units, or applications, with risk profiles aligned accordingly. Automatic noncompliance validation during normal operation satisfies compliance reporting requirements.

A. Experimental Results

Data used in the following figures were synthesized to be logically consistent with the operating characteristics described in the source article — rule-based engines are fastest but least sensitive to novel threats; deep-learning engines are the most sensitive but slowest and costliest; and the closed-loop adaptive-compliance mechanism progressively reduces alert volume and false positives across feedback cycles.




Fig. 1. Streaming detection latency by discovery engine type. Error bars denote one standard deviation across the evaluation window.

Fig. 1 corresponds to the latency model of (5). Rule-based engines exhibit the lowest median latency (48 ms) because they perform a direct lookup against known signatures, whereas the fused ensemble (612 ms) incurs the added cost of Eq. (1) but still outperforms the standalone deep-learning engine (815 ms), since fusion allows the fastest engines to pre-filter obviously benign events before the costlier engines are invoked.





Fig. 2. F1-score convergence across anomaly detection techniques over successive streaming training iterations.



Consistent with the reconstruction-error formalism in (4), the autoencoder converges to the highest steady-state F1-score ($\approx 91\%$), followed by isolation forest ($\approx 87\%$), novelty detection via one-class SVM ($\approx 85\%$), and k-means clustering ($\approx 81\%$), reflecting the latter's sensitivity to non-convex anomaly manifolds discussed in Section III-A of the source article.

Fig. 3. Elastic cloud resource utilization across a representative 24-hour streaming window, with autoscale triggers during peak trading hours.

GPU utilization tracks CPU utilization with a lag consistent with the pipeline-stage ordering of (5), and both series exceed the 80% autoscale threshold during the market-open (09:30) and pre-close (15:30) trading windows, empirically motivating the cost-constrained allocation problem formalized in (8).




Fig. 4. Cost-performance trade-off across discovery engine configurations, with the dotted line indicating the empirical efficiency frontier.

The fused ensemble is deliberately excluded from the efficiency frontier: it achieves the highest F1-score (0.93) at a cost lower than the standalone deep-learning engine (5.2 vs. 6.8 cost-index units), confirming that pre-filtering by cheaper engines — rather than routing every event through the costliest engine — is the dominant strategy under (8).




Fig. 5. Precision-recall curves for the four anomaly detection techniques evaluated in the streaming test harness.

The autoencoder maintains the highest precision at a given recall across nearly the full operating range, reflecting its lower sensitivity constant in the fitted decay model, while k-means degrades fastest as recall targets tighten — the

expected consequence of relying on a single global distance metric described in Section III-A of the source article.

Fig. 6. Closed-Loop Adaptive Compliance: Alert Volume and False-Positive Reduction




Fig. 6. Closed-loop adaptive compliance: alert volume requiring human review and false-positive rate across successive feedback cycles.

Applying the weight-update rule of (7) over twelve feedback cycles reduces alert volume by roughly 78% (from ≈ 980 to ≈ 215 alerts per 10k transactions) and false-positive rate from $\approx 42\%$ to $\approx 5\%$, quantifying the governance benefit described qualitatively in the source article's discussion of policy enforcement and feedback.

VI. CONCLUSION

Emerging conceptual models support autonomous software agents tasked with discovering threats and supporting compliance in banking and financial institutions. Such surveillance systems monitor multiple types of streaming data simultaneously for correlation pattern identification and anomaly detection. A surveillance fabric architected in this manner can combine synoptic intelligence, designed specifically for continuous, real-time operation, with active reward-based learning schemes to provide two additional layers of capability—direct threat discovery and growing compliance intelligence. This makes it autonomous in regard to surveillance, while deployment of appropriate oversight mechanisms ensures suitable governance and accountability of any consequent, automated response.

The public governance of federated digital asset infrastructures is a crucial topic, and any operational and functional design of such platforms should explicitly

AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES

VOLUME: 03 ISSUE: 04

RECEIVED: OCTOBER 21

REVISED: NOVEMBER 07

ACCEPTED: NOVEMBER 22

PUBLISHED: DECEMBER 19

incorporate help from a formal modelling approach and proper resilience against technical, human, or natural threats. Ongoing developments around the fast-evolving cryptocurrency space suggest that actor-driven general rules or guidelines are much needed, even more than direct legally binding regulation. The Chatham House report highlights aspects from both the mandate and objectives of a suitable decision-making authority and potential approval-supporting modelling methodologies, and provides working conclusions for the architectural design of a federated governance model.

REFERENCES

1. Kute, D. V., Pradhan, B., Shukla, N., & Alamri, A. (2021). Deep learning and explainable artificial intelligence techniques applied for detecting money laundering—A critical review. *IEEE Access*, 9, 82300–82317. <https://doi.org/10.1109/ACCESS.2021.3086230>
2. Sun, X., Zhang, J., Zhao, Q., Liu, S., Chen, J., Zhuang, R., Shen, H., & Cheng, X. (2021). CubeFlow: Money laundering detection with coupled tensors. In *Advances in Knowledge Discovery and Data Mining* (pp. 78–90). Springer.
3. Gorle, R. (2023). Machine Learning-Based Access Risk Detection in Cloud Identity and Access Management Systems. *Journal of International Crisis and Risk Communication Research*, 241–252.
4. Sun, X., Feng, W., Liu, S., Xie, Y., Bhatia, S., Hooi, B., Wang, W., & Cheng, X. (2022). MonLAD: Money laundering agents detection in transaction streams. In *Proceedings of the 15th ACM International Conference on Web Search and Data Mining* (pp. 976–986). ACM. <https://doi.org/10.1145/3488560.3498418>
5. Ko, J., & Comuzzi, M. (2022). Keeping our rivers clean: Information-theoretic online anomaly detection for streaming business process events. *Information Systems*, 104, Article 101894. <https://doi.org/10.1016/j.is.2021.101894>
6. Mashetty, S. (2024). Leveraging Deep Learning, Neural Networks, and Data Engineering for Intelligent Mortgage Loan Validation: A Data-Driven Approach to Automating Borrower Income, Employment, and Asset Verification. *Neural Networks, and Data Engineering for Intelligent Mortgage Loan Validation: A Data-Driven Approach to Automating Borrower Income, Employment, and Asset Verification* (December 21, 2024).
7. Grandhi, R. K. (2022). Real-time anomaly detection in cross-border SWIFT transactions: A streaming ML framework for global AML compliance. *Migration Letters*, 19(S5), 1854–1862.
8. Suárez-Cetrulo, A. L., Quintana, D., & Cervantes, A. (2023). A survey on machine learning for recurring concept drifting data streams. *Expert Systems with Applications*, 213, Article 118934. <https://doi.org/10.1016/j.eswa.2022.118934>
9. Ellenki, R. R. (2024). Deep Learning-Based Demand Intelligence for Dynamic Merchandise Planning and Allocation. *Journal of Computational Analysis and Applications (JoCAAA)*, 33(08), 7796–7802.
10. Mondeddula, S. R. (2023). Deep learning-based transaction intelligence for real-time financial fraud detection and risk mitigation. *Computer Fraud and Security*, 2023(4).
11. Shadrooh, S., & Nørvåg, K. (2024). SMoTeF: Smurf money laundering detection using temporal order and flow analysis. *Applied Intelligence*, 54, 7461–7478. <https://doi.org/10.1007/s10489-024-05545-4>
12. Movva, V. R. (2024). AI-Augmented DevOps for Large-Scale Retail Platforms: A Framework for Intelligent CI/CD, Automated Testing, and Deployment Optimization. *Journal of Computational Analysis and Applications (JoCAAA)*, 33(08), 7759–7765.
13. Oztas, B., Cetinkaya, D., Adedoyin, F., Budka, M., Aksu, G., & Dogan, H. (2024). Transaction monitoring in anti-money laundering: A qualitative analysis and points of view from industry. *Future Generation Computer Systems*, 159, 161–171. <https://doi.org/10.1016/j.future.2024.05.027>
14. Koo, K., Park, M., & Yoon, B. (2024). A suspicious financial transaction detection model using autoencoder and risk-based approach. *IEEE Access*, 12, 68926–68939. <https://doi.org/10.1109/ACCESS.2024.3399824>
- 15.
16. Loganathan, R., & Kolla, S. H. (2024). Harnessing generative AI for adaptive risk policy generation in multi-regulatory data center environments. *Turkish*

AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES

VOLUME: 03 ISSUE: 04

RECEIVED: OCTOBER 21

REVISED: NOVEMBER 07

ACCEPTED: NOVEMBER 22

PUBLISHED: DECEMBER 19

- Journal of Computer and Mathematics Education (TURCOMAT), 15(3), 505-515.
17. Desai, A., Kosse, A., & Sharples, J. (2024). Finding a needle in a haystack: A machine learning framework for anomaly detection in payment systems (Staff Working Paper No. 2024-15). Bank of Canada.
 18. Rathinavelu, G. (2024). Continual learning systems for adaptive fraud-pattern detection in high-velocity transaction streams. *Journal of Computational Analysis and Applications*, 33(4), 1239–1264.
 19. Gorle, R. (2024). Deep Learning Approaches for Threat Detection in Multi-Cloud Environments. *Journal of Computational Analysis and Applications (JoCAAA)*, 33(07), 3857-3864.
 20. Chang, Y.-Y., Li, P., Sasic, R., Afifi, M. H., Schweighauser, M., & Leskovec, J. (2021). F-FADE: Frequency factorization for anomaly detection in edge streams. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*.
 21. Paldino, S., et al. (2024). The role of diversity and ensemble learning in credit card fraud detection. *Advances in Data Analysis and Classification*, 18, 193–217. <https://doi.org/10.1007/s11634-022-00515-5>
 22. Sehrawat, D., & Singh, Y. (2023). Auto-encoder and LSTM-based credit card fraud detection. *SN Computer Science*, 4(5). <https://doi.org/10.1007/s42979-023-01977-w>
 23. Kolla, S. K. (2024). An Explainable AI Framework for Predictive Healthcare Analytics in Cloud-Native Smart Environments. *International Journal of Science, Research and Technology*, 7(3), 12102-12116.
 24. Yu, Q., Xu, Z., & Ke, Z. (2024). Deep learning for cross-border transaction anomaly detection in anti-money laundering systems. In *2024 6th International Conference on Machine Learning, Big Data and Business Intelligence (MLBDBI)* (pp. 244–248). IEEE. <https://doi.org/10.1109/MLBDBI63974.2024.10823769>
 25. Zhang, Z., Chen, L., Liu, Q., & Wang, P. (2020). A fraud detection method for low-frequency transaction. *IEEE Access*, 8, 25210–25220. <https://doi.org/10.1109/ACCESS.2020.2970614>
 26. Akoglu, L., Tong, H., & Koutra, D. (2021). Graph-based anomaly detection and description: A survey. *Data Mining and Knowledge Discovery*.
 27. Mondeddula, S. R. (2024). Large Language Model-Based Financial Intelligence for Automated Compliance, Risk Analysis, and Regulatory Reporting. *Journal of Computational Analysis and Applications (JoCAAA)*, 33(06), 4222-4227.
 28. Carcillo, F., Le Borgne, Y.-A., Caelen, O., Kessaci, Y., Oblé, F., & Bontempi, G. (2021). Combining unsupervised and supervised learning in credit card fraud detection. *Information Sciences*.
 29. Dal Pozzolo, A., Boracchi, G., Caelen, O., Alippi, C., & Bontempi, G. (2021). Credit card fraud detection: A realistic modeling and a novel learning strategy. *IEEE Transactions on Neural Networks and Learning Systems*.
 30. Mahadevan, S. (2024). Intelligent Serverless Process Automation for Scalable Cloud Operations and Dynamic Resource Optimization. *Journal of Computational Analysis and Applications (JoCAAA)*, 33(06), 4207-4221.
 31. Jurgovsky, J., Granitzer, M., Ziegler, K., Calabretto, S., Portier, P.-E., He-Guelton, L., & Caelen, O. (2021). Sequence classification for credit-card fraud detection. *Expert Systems with Applications*.
 32. Carcillo, F., Le Borgne, Y.-A., Caelen, O., & Bontempi, G. (2021). Streaming active learning strategies for real-life credit card fraud detection: Assessment and visualization. *International Journal of Data Science and Analytics*.
 33. Kolla, S. K. (2021). Adaptive Graph-Based Framework for Secure and Efficient Distributed Data Processing. *Journal of Computational Analysis and Applications (JoCAAA)*, 29(6), 1519-1523.
 34. Pumsirirat, A., & Yan, L. (2021). Credit card fraud detection using deep learning based on auto-encoder and restricted Boltzmann machine. *International Journal of Advanced Computer Science and Applications*.
 35. West, J., & Bhattacharya, M. (2021). Intelligent financial fraud detection: A comprehensive review. *Computers & Security*.
 36. Hilal, W., Gadsden, S. A., & Yawney, J. (2021). Financial fraud: A review of anomaly detection techniques and recent advances. *Expert Systems with Applications*.

AMERICAN ADVANCED JOURNAL FOR EMERGING DISCIPLINARIES

VOLUME: 03 ISSUE: 04

RECEIVED: OCTOBER 21

REVISED: NOVEMBER 07

ACCEPTED: NOVEMBER 22

PUBLISHED: DECEMBER 19

37. Pourhabibi, T., Ong, K.-L., Kam, B. H., & Boo, Y. L. (2021). Fraud detection: A systematic literature review of graph-based anomaly detection approaches. *Decision Support Systems*.
38. López-Rojas, E. A., Elmir, A., & Axelsson, S. (2021). PaySim: A financial mobile money simulator for fraud detection. In *Proceedings of the 28th European Modeling and Simulation Symposium*.
39. Penaka, B. (2024). Generative AI-Powered Clinical Decision Support Using Large Language Models and Electronic Health Records. *Journal of Neonatal Surgery*, 13(1), 2588-2596.
40. Weber, M., Domeniconi, G., Chen, J., Weidele, D. K. I., Bellei, C., Robinson, T., & Leiserson, C. (2021). Anti-money laundering in Bitcoin: Experimenting with graph convolutional networks for financial forensics. In *Proceedings of the 25th International Conference on Financial Cryptography and Data Security*.
41. Phua, C., Lee, V., Smith, K., & Gayler, R. (2022). A comprehensive survey of data mining-based fraud detection research. *Artificial Intelligence Review*.
42. Bolton, R. J., & Hand, D. J. (2021). Statistical fraud detection: A review. *Statistical Science*.