NOLEGEIN Journal of
Leadership and Strategic Management

ISSN: 2582-2888
Volume 9, Issue 2, 2026
July—December
DOI (Journal): 10.37591/NJLSM

https://mbajournals.in/index.php/JoKSM

Research NJLSM

The Intelligent Frontier: Orchestrating AI and Machine Learning for Advanced Financial Fraud Detection

Shipra Aggarwal*

Abstract

Traditional rule-based systems cannot match that scale. They are too slow, too rigid, and too reactive. This paper examines how AI and ML are reshaping fraud detection. It evaluates LSTMs, GNNs, and Transformers as detection tools. Mastercard reported a 20% uplift in fraud detection rates. The paper also maps two emerging frontiers: Agentic AI and Quantum Graph Neural Networks. Strategic guidance is offered throughout CFOs and finance leaders. Global financial crime reached $4.4 trillion in 2025. Fraud networks now deploy AI against the very systems designed to stop them. Static detection methods are no longer adequate. This paper provides a structured evaluation of AI and machine learning architectures for financial fraud detection. Supervised models, including Random Forest and Logistic Regression, form the operational baseline. Deep learning architecture extends this capability significantly. Long Short-Term Memory networks capture sequential transaction patterns. Graph Neural Networks expose hidden criminal networks across accounts. Transformers apply self-attention to behavioural context across full transaction histories. Generative Adversarial Networks address the persistent class imbalance problem in training data. The paper introduces the DLSG framework as a governance blueprint. It integrates deep learning performance with sector-specific compliance requirements. GDPR and CCPA obligations are embedded at the design level. Case studies from Mastercard and Visa demonstrate real-world impact. Mastercard achieved a 20% uplift in detection rates. Visa processes risk scores in under one second per transaction. Ethical risks, including algorithmic bias and explainability gaps, are examined. SHAP-based tools and federated learning are presented as mitigations. Two frontier technologies conclude the analysis: Agentic AI and Quantum Graph Neural Networks. Both will redefine fraud defence within the decade. Strategic recommendations guide CFOs and management accountants toward governance-first AI adoption.

Keywords: Fraud detection, machine learning, deep learning, graph neural networks, DLSG framework, explainable AI, financial crime, agentic AI, quantum computing, GDPR compliance

INTRODUCTION

Financial crime has never posed a greater threat. Illicit activity reached $4.4 trillion in 2025. That figure outpaced global GDP growth last year [1]. Fraud networks are not merely growing – they are adapting. They now deploy the very AI tools designed to detect them.

***Author for Correspondence**
Shipra Aggarwal
E-mail: saggarwal5@alum.babson.edu

HOD, Department of Finance, Pelham Community Pharmacy
Waltham, MA 02451, USA

Received Date: May 18, 2026
Acceptance Date: May 28, 2026
Publication Date: June 12, 2026

Citation: Shipra Aggarwal. The Intelligent Frontier: Orchestrating AI and Machine Learning for Advanced Financial Fraud Detection. NOLEGEIN Journal of Leadership and Strategic Management. 2026; 9(2): 30–37p.

For management accountants, the stakes have fundamentally changed. Retrospective analysis is no longer sufficient [2]. Static rule-based systems are outpaced daily. Machine learning has become a strategic survival tool [3]. It is no longer optional for institutional resilience.

The cost of inaction is clear and concrete. A typical organisation loses roughly 5% of revenue to fraud [4]. The median scheme runs undetected for about twelve months. Losses are intensifying as criminals adopt better technology. This paper

addresses all this directly. It covers foundational models, advanced architecture, and the DLSG framework [3]. It examines real results at Mastercard and Visa [5, 6]. It also looks ahead to Agentic AI and quantum-enhanced security [7].

THE GLOBAL PARADIGM SHIFT IN FINANCIAL CRIME

Financial fraud today presents a stark paradox. Digital systems are more transparent than ever. Yet criminal methods have grown more complex, not less. In 2025, global fraud losses reached roughly $579.4 billion [1]. Sustained compound growth signals a systemic crisis, not a cyclical blip.

The threat is further complicated by poly criminality. Financial fraud now intersects with human trafficking, cybercrime, and terrorism financing. No single institution can address those overlapping vectors alone [7]. Coordinated cross-sector responses are essential.

Quantitative Impact and Sectoral Vulnerabilities

The ACFE provides the clearest picture of fraud's true cost. Asset misappropriation is the most common occupational fraud type. It occurs in 86% of cases yet carries a median loss of just $120,000 [4]. Financial statement fraud appears in only 5% of cases. Its median loss reaches $766,000 per incident. Internal controls must be tiered by consequence, not just frequency.

Table 1 shows fraud losses by industry sector [4, 8]. Mining and manufacturing sustain the highest median losses. Their supply chains are complex and hard to monitor centrally. Online communities and gaming platforms show a different pattern. Fraud attempt rates reach 11.7% and 9.8%, respectively. Criminals are targeting account creation to build long-term false identities [8].

Table 1. Median fraud loss and suspected attempt rates by industry sector [4, 8].

Industry sector	Median fraud loss (2024–2025)	Suspected attempt rate (%)
Mining	$550,000	N/A
Wholesale Trade	$361,000	N/A
Manufacturing	$267,000	N/A
Banking & Financial Services	$200,000	3.2%
Healthcare	$100,000	N/A
Communities (Dating/Forums)	N/A	11.7%
Gaming (Sports Betting)	N/A	9.8%
Retail	$48,000	3.8%

The Proactive Shift in Monitoring

Corporate fraud detection was once entirely reactive. Retrospective reviews and whistleblowers drove most discoveries. Tips still account for 43% of cases detected today. That reliance creates dangerous lag. Fraud can run unchecked for months before it surfaces [4].

AI changes this equation fundamentally [5, 6]. Systems now score transactions in real time. Risk assessments are delivered in under 120 milliseconds. Compliance responses can adjust almost instantaneously [5]. Manual methods simply cannot operate at that speed or scale [9].

FOUNDATIONAL MACHINE LEARNING MODELS IN FINANCE

Model selection is not a minor technical decision. The wrong choice degrades detection quality significantly. Research from 2012 to 2023 reveals a dominant trend. Supervised learning is used in roughly 56.7% of applications [10].

Supervised Learning: The Workhorse of Fraud Detection

Supervised models learn from pre-labelled transaction data. Each transaction is marked fraudulent or legitimate in advance [10]. The model then learns which features distinguish between the two groups.

Table 2 ranks the most cited models in current literature. Random Forest dominates, appearing in 34 studies. It handles large, multi-variable financial datasets well. Hundreds of decision trees vote together, which reduces noise. Support Vector Machines suit high-stakes settings like mortgage lending [10]. They draw precise boundaries between safe and risky transactions.

Table 2. Supervised learning model usage in fraud detection literature [10].

Model	Citations in literature	Primary application	Key characteristic
Random Forest (RF)	34	Credit card / financial statements	High accuracy; resistant to overfitting.
Logistic Regression (LR)	32	Binary classification	Simplicity; high interpretability.
Support Vector Machine	29	High-dimensional data	Effective for outlier/anomaly detection.
Decision Tree (DT)	29	Behavioural profiling	Visual; easy to explain to stakeholders.
Naive Bayes (NB)	19	Text/document fraud	Probabilistic classification approach.

Unsupervised Learning and Anomaly Detection

Supervised models depend heavily on labelled training data. Labels are costly to produce and go stale quickly. Fraudsters innovate; labelled datasets cannot always keep up. Unsupervised learning requires no labels at all. Models, like autoencoders and Isolation Forests, find statistical outliers. This covers roughly 18.27% of current research [10] and is especially useful for novel, lone-wolf fraud schemes.

Performance Evaluation Metrics

Accuracy alone misleads in fraud detection. Legitimate transactions vastly outnumber fraudulent ones. A model labelling everything legitimate can still appear highly accurate [3]. This is called class imbalance, and it is pervasive [10]. Four metrics provide a more honest picture:

- *Recall:* identifies all genuine fraud cases. Minimises missed financial losses.
- *Precision:* confirms flagged cases are truly fraudulent. Reduces customer friction.
- *F1 Score:* balances precision and recall in one figure.
- *AUC-PR:* focuses on minority-class performance. Most reliable for fraud models [3].

DEEP LEARNING AND ADVANCED NEURAL ARCHITECTURES

Traditional ML handles structured data well. Modern fraud is neither structured nor simple. Digital payments create high-dimensional, relational data. Deep learning extracts features that simpler models miss entirely [3]. It captures subtle, non-linear relationships within vast datasets.

Long Short-Term Memory Networks and Sequential Analysis

LSTMs are Recurrent Neural Networks built for sequences. In fraud, the order of transactions often reveals intent [3]. A single transaction rarely tells the whole story. Consider smurfing: large sums are broken into small deposits. Each deposit looks harmless in isolation. LSTMs maintain memory across time and catch the full pattern [10].

Transformers and Behavioural Context

Transformers power today's large language models. They are now applied to fraud detection as well. Standard models process transactions in strict sequence. Transformers process all transaction history simultaneously. Self-attention weighs each signal's importance regardless of timing [3]. This captures context that sequential models routinely overlook. Their use in fraud detection grew sharply through 2024 and 2025 [8].

Graph Neural Networks and Network Analysis

Fraud rarely occurs in isolation. Criminals operate through networks of accounts, merchants, and devices. Graph Neural Networks model these relationships directly. Entities become vertices;

connections become edges. GNNs trace funds through intermediaries and expose laundering cycles. Graph Convolutional Networks outperform traditional models in blockchain fraud detection [3].

Generative Adversarial Networks and the Class Imbalance Challenge

Training data for fraud models is almost always imbalanced. Ninety-two percent of studies confirm this problem [10]. Fraud is rare; clean transactions dominate every dataset. Generative Adversarial Networks directly address this gap. One network generates synthetic, realistic fraud examples. A second network tries to distinguish them from real data. This competition produces high-quality synthetic training samples [3]. Detection accuracy and model robustness both improve as a result. ML-based predictive models generalise well beyond financial data. Analogous GAN and regression frameworks now drive emissions forecasting. Process yield prediction and property screening use the same core methods. Cross-sector validation strengthens confidence in these detection architectures [11].

THE DLSG FRAMEWORK: A STRATEGIC BLUEPRINT FOR GOVERNANCE

Technical capability alone is not enough for safe deployment. AI systems must be lawful and contextually appropriate. The DLSG framework was developed to bridge this gap. It has three interlocking layers: Deep Learning, Sector Context, and Governance [3].

The Deep Learning layer focuses on model performance. The Sector Context layer recognises that sectors differ sharply. Credit card fraud requires sub-second decisions on a massive scale. Insurance fraud calls for federated learning to protect patient data. GDPR and HIPAA compliance must be built in from the start.

The Governance layer integrates all legal mandates. GDPR's Right to Explanation requires AI decisions to be interpretable [2]. Table 3 summarizes all three layers and their compliance outcomes [3].

Table 3. The DLSG framework: layers, technologies, and compliance outcomes [3].

DLSG layer	Primary focus	Key technologies	Compliance / outcome
Deep Learning	Technical innovation	CNNs, LSTMs, GNNs, Transformers	Accuracy, automation, scale.
Sector Context	Operational needs	Real-time APIs, Federated Learning	Low latency, contextual logic.
Governance	Ethical compliance	XAI (SHAP), PCA anonymization	GDPR/CCPA alignment, fairness.

Figure 1 illustrates this end-to-end detection workflow. Each DLSG layer maps to a concrete processing stage. The diagram shows data ingestion, model scoring, and governance output. Together these stages form a closed, auditable feedback loop.

CASE STUDIES: REAL-WORLD AI IMPLEMENTATION

Global payments leaders have set the benchmark for AI deployment [5, 6]. Their results quantify what responsible implementation can achieve.

Mastercard: Decision Intelligence and Network Insight

Mastercard moved away from rigid rule-based filters. Those rules flagged every transaction over a set threshold. False-positive rates were high, and customer experience suffered. Decision Intelligence replaced that approach entirely [5, 12].

The system analyses over 150 billion transactions each year [5]. Smart Agents technology scores each one for risk [12]. Scoring takes between 100 and 120 milliseconds in the cloud. Behavioural context is central to each scoring decision. A luxury purchase during a seasonal sale by a known customer is treated differently from the same purchase on an unfamiliar device. Mastercard reported a 20% increase in detection rates. False positives fell substantially at the same time [5].

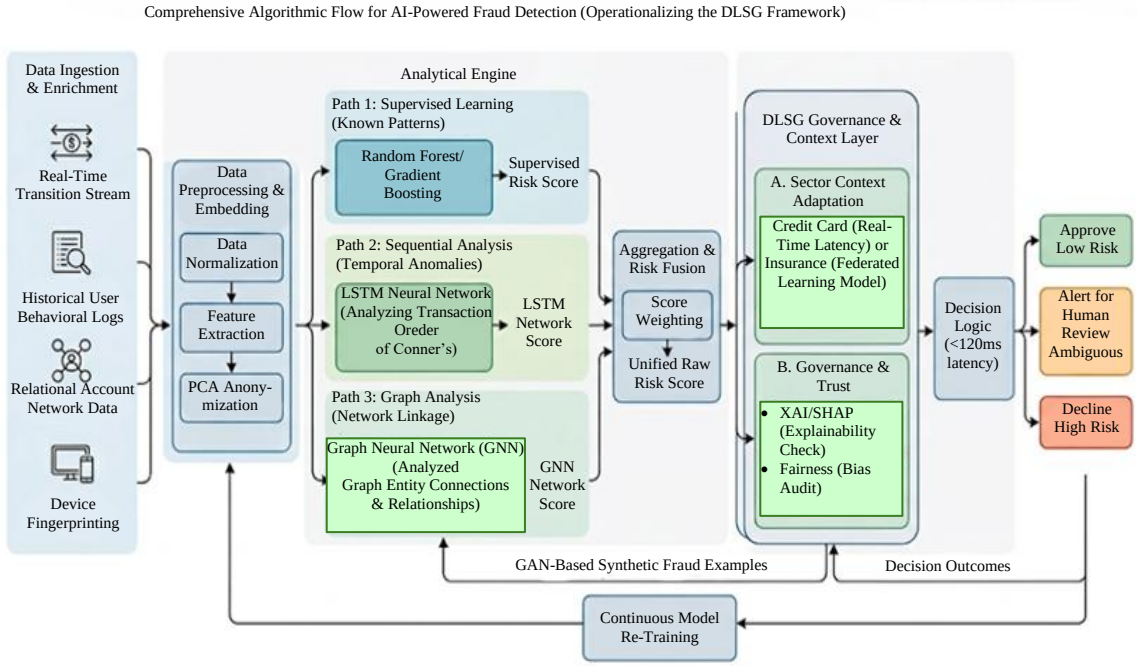

Figure 1. Integrated algorithmic workflow for real-time fraud detection: operationalizing the DLSG framework.

Visa: Fraud Management Essentials and Identity Behaviour Analysis

Visa trains its ML models on the VisaNet dataset. That dataset spans over 269 billion global transactions. Risk scores are generated in under one second per transaction. The Fraud Management Essentials tool targets small businesses. It provides out-of-the-box fraud filters and automated risk scoring. Identity Behaviour Analysis is Visa's most recent innovation. It builds a positive model for each customer's normal behaviour. A known customer buying from a new device is still recognised [6]. This significantly reduces friction for legitimate users.

Results and ROI Comparisons

AI adoption has produced measurable gains across the industry. Eighty-five percent of institutions report direct returns from AI. Those returns come through faster pattern recognition and investigation. Some merchants cut fraud loss rates by 30 to 50%. That reduction came within six months of AI deployment. Table 4 contrasts rule-based and AI-enhanced systems directly [5, 6, 12].

Table 4. Performance comparison: rule-based systems vs. AI-enhanced systems [5, 6, 12].

Implementation metric	Rule-based systems	AI-enhanced systems
Detection Lag	High (manual reviews)	Near real-time (<120 ms).
False Positive Rate	High (static rules)	Reduced by up to 83%.
ROI / Savings	N/A	>$5 million for 42% of issuers.
Manual Review Volume	High	80% reported reduction.

THE HUMAN ELEMENT: TACIT KNOWLEDGE AND STRATEGIC INSIGHT

AI as Enhancement for Human Analysis

AI does not replace financial professionals – it elevates them. Management accountants are moving from data custodians to strategic advisors [2]. AI handles repetitive work that once consumed hours. Data reconciliation that took a day now takes seconds. Optical Character Recognition processes invoices automatically at scale [9]. Generative AI lets CFOs run financial scenarios on demand. Freed from routine tasks, professionals can focus on strategic insight.

Tacit Knowledge and Behavioural Red Flags

Experience builds knowledge that no algorithm can replicate [9]. Identifying a discrepancy is not enough on its own [4]. Investigators must also assess means, motive, and opportunity. That judgment comes from years of practice, not training data. Skilled accountants notice things with no model flags: resistance to new controls, living beyond one's means [4], and unusually close vendor relationships. The strongest fraud defence combines AI speed with human wisdom [9].

ETHICAL CONSIDERATIONS, BIAS, AND GOVERNANCE RISKS

Managing Algorithmic Bias

Bias enters AI models through the data they learn from. ZIP codes used in mortgage analysis can become proxies for race. Models trained on biased data may discriminate unintentionally. Audit history creates a parallel problem. Departments monitored more closely produce more training fraud signals. Models then inherit the same blind spots. Regular output auditing for disparate impact is essential [2, 3]. The awareness-practice gap here is striking. Seventy-five percent of respondents say bias matters. Yet only 18% actually test their models for fairness [8]. That gap is both an ethical failure and a regulatory risk [2].

Transparency and Explainability

Only 6% of anti-fraud professionals trust their AI's reasoning fully. That lack of confidence is itself a barrier to adoption. Eighty-two percent of practitioners say explainability matters [8]. Yet most cannot explain how their systems reach decisions. GDPR's Right to Explanation makes this a legal obligation. SHAP values identify which features drove each fraud flag. These tools give auditors the evidence trail they need [2, 3].

Privacy-Preserving Technologies

Crime prevention and privacy rights exist in genuine tension [2, 3]. Several technologies help institutions navigate it responsibly. Principal Component Analysis anonymises transactional data. It reduces dimensions while retaining the variance needed for detection. Federated Learning trains models without centralising raw data. Each server learns locally; data never leaves its source. Blockchain adds an immutable audit trail through smart contracts. Data integrity is enhanced across the entire transaction lifecycle [1].

FUTURE FRONTIERS: AGENTIC AI AND QUANTUM SECURITY

The Rise of Agentic AI

A new class of threat is now emerging. Agentic AI can plan and execute complete fraud campaigns. No human operator is needed at any stage. These agents perform reconnaissance, craft social engineering scripts, and issue ransom demands fully autonomously. AI-enhanced fraud is 4.5 times more profitable than traditional methods [7]. In response, 31% of organisations plan agentic AI deployment for fraud defence by 2028 [8].

Quantum Computing and Enhanced Security

Sixty-two percent of anti-fraud professionals expect quantum impact by 2030. The challenge is dual-sided. Quantum processing could break current encryption standards [8]. Yet it also offers Quantum AI as a detection tool. Quantum Graph Neural Networks can process relationship graphs at extraordinary speed. Classical computers cannot match that computational capacity [3].

Identity verification is also evolving rapidly. Biometric orchestration combines physical and behavioural signals. Ninety-eight percent of organisations now want this capability. Physical biometrics adoption rose from 34% in 2022 to 45% in 2026 [8].

STRATEGIC RECOMMENDATIONS FOR CFOS AND FINANCE LEADERS

Fraud has been industrialised, and the response must match that scale [1, 7]. Management accountants must evolve into orchestrators of ethical AI systems [9].

Implementing Robust Internal Controls

Eighty-two percent of victim organisations changed controls after the fact. Leaders should not wait for losses before acting. Proactive analytics cut median fraud losses by 55%. Surprise audits and strict authorisation processes remain primary defences. Fraud awareness training accelerates detection by roughly 50%. Organisations without such programmes detect fraud far later [4].

Adopting a Governance-First AI Strategy

CFOs deploying ML models should apply the DLSG framework. Technical innovation must stay balanced with legal requirements. Seventy-five percent acknowledge bias as a concern, yet only 18% test their models for it regularly [3, 8]. That inconsistency is both an ethical and regulatory exposure. SHAP-based XAI tools help boards defend algorithmic decisions [2, 3]. Regular model auditing is non-negotiable.

Cultivating Professional Vigilance

Technology never replaces human vigilance entirely. Forty-three percent of fraud is still found through tips. The culture of reporting is as valuable as any algorithm [4]. Employees, vendors, and customers need a safe reporting channel. Fear of retaliation silences the most important witnesses. Leaders must fund anti-fraud initiatives and upskilling consistently. CFOs and CIOs must collaborate far more closely than they do today. AI is only as reliable as the governance surrounding it.

CONCLUSION

The $4.4 trillion illicit finance figure is not an abstraction. It represents stolen assets, damaged institutions, and human harm. Financial crime is now central to global economic stability [1]. AI and ML offer a credible and urgent answer. They move the industry from reactive detection to proactive defence [3].

The DLSG framework ensures this transition is governed properly. Neural architectures from LSTMs to GNNs provide the analytical power [3]. Human judgment provides the wisdom to deploy it responsibly [9]. The coming era of Agentic AI will test every institution [7]. Quantum-enhanced fraud will raise the stakes further [8]. Organisations that build governance-first AI strategies will be ready [2, 3]. Those who wait for a breach will define the losses to come [4].

REFERENCES

1. Nasdaq Verafin. 2026 Global Financial Crime Report. 2026. Available from: https://verafin.com/wp-content/uploads/2026/03/global-financial-crime-report-2026-nasdaq-verafin-20260316.pdf
2. Garcia-Segura LA. The role of artificial intelligence in preventing corporate crime. J Econ Criminol. 2024;5:100091. Available from: https://doi.org/10.1016/j.jeconc.2024.100091
3. Chen Y, Zhao C, Xu Y, Nie C, Zhang Y. Deep learning in financial fraud detection: Innovations, challenges, and applications. Data Science and Management. 2025. Advance online publication. Available from: https://doi.org/10.1016/j.dsm.2025.08.002
4. Association of Certified Fraud Examiners (ACFE). Occupational Fraud 2024: A Report to the Nations. 2024. Available from: https://www.anchin.com/wp-content/uploads/2024/08/2024-ACFE-Occupational-Fraud-Report.pdf
5. Mastercard. AI is helping banks save millions by transforming payment fraud prevention. 2026. Available from: https://www.mastercard.com/global/en/news-and-trends/Insights/2026/ai-is-helping-banks-save-millions-by-transforming-payment-fraud-prevention.html
6. Bankcard International Group. Inside AI fraud detection in payments 2026. 2026. Available from: https://bankcardinternationalgroup.com/inside-ai-fraud-detection-in-payments-2026/
7. INTERPOL. INTERPOL report warns of increasingly sophisticated global financial fraud threat. 2026. Available from: https://www.interpol.int/News-and-Events/News/2026/INTERPOL-report-warns-of-increasingly-sophisticated-global-financial-fraud-threat
8. TransUnion. H1-2026 update to the top fraud trends report. 2026. Available from: https://newsroom.transunion.com/h1-2026-update-to-the-top-fraud-trends-report/

9. IMA—Strategic Finance. Using algorithms to root out fraud. 2023 Feb. Available from: https://www.sfmagazine.com/articles/2023/february/using-algorithms-to-root-out-fraud
10. Hernandez Aros L, Bustamante Molano LX, Gutierrez-Portela F, Moreno Hernandez JJ, Rodríguez Barrero MS. Financial fraud detection through the application of machine learning techniques: A literature review. Humanit Soc Sci Commun. 2024;11:1130. Available from: https://doi.org/10.1057/s41599-024-03606-0
11. Aggarwal M. Sustainable aviation fuel: Powering the future of clean air travel. Hydrocarbon Processing. 2025 Oct. Available from: https://read.nxtbook.com/gulf_energy_information/hydrocarbon_processing/october_2025/biofuels_alternative_renewable_fuels_aggarwal_emerson_s_aspen_technology_business.html
12. Mastercard. Advanced AI & machine learning technology for risk decisioning—Mastercard Risk Decisioning Platform. [cited n.d.]. Available from: https://www.mastercard.com/content/mccom/eeu-language-masters/en/business/cybersecurity-fraud-prevention/risk-decisioning/mastercard-risk-decisioning-platform.html

