
# Generative AI Framework for Wealth Fraud Detection & Compliance Risk

**Dhanaraj Sathiri**
Independent Researcher
dhanrajsathiri@gmail.com

## Abstract

The Financial Intelligence Framework is grounded in the epistemic and methodological foundations of Generative Financial Intelligence, distinguishing predictive analytics from a broader paradigm capable of addressing complex, adaptive, and undocumented problems while emphasizing fraud-aware analytics. Predictive financial signal detection focuses narrowly on solving forecasting tasks at low total cost of ownership, supporting risk signal detection, control, and decision activation. Generative Financial Intelligence extends this scope beyond risk detection to encompass profiling, waypointing, advisory functions, prescription, and debiasing, making it well-suited to adaptive domains such as fraud detection and prevention. Within these domains, the framework enables the classification of fraud types and construction of taxonomies, the delineation of indicators, and their linkage to signal detection, control mechanisms, and detection-threshold decisions—culminating in the automation and formalization of financial fraud analysis.

Wealth-related data play a critical role for governments and regulators, serving both as a foundation for designing and structuring laws and regulations and as an attribute that obligates citizens within those jurisdictions. Wealth analysis shapes stakeholder demand for, and citizen adherence to, wealth-reporting regulations, while informing governance requirements that define risk appetite and supervisory expectations. The dual use of wealth data—by regulators and the regulated alike—governs the development of wealth-related artificial intelligence functions.

Financial Intelligence integrates all dimensions of financial management, enabling oversight of risk decisions and the execution of financial functions, areas, and processes. The proposed Generative Financial Framework applies generative, semi-supervised, and hybrid Generative Financial Intelligence models across relevant financial functions and contexts. It further illustrates model oversight through considerations of data quality, feature engineering, model validation, and interpretability, alongside the embedding of decision-making approaches.

**Keywords:** Generative Financial Intelligence, Financial Intelligence Framework, Fraud-Aware Analytics, Financial Fraud Detection, Predictive Financial Analytics, Risk Signal Detection, Wealth Intelligence, Compliance Analytics, Explainable Artificial Intelligence, Decision Intelligence, Financial Risk Management, Regulatory Technology (RegTech).

# 1. Introduction

Crime is a cancer that cripples economic growth, threatens macrofinancial stability, and undermines social cohesion. Recent economic downturns have renewed concerns about an illicit economy that, according to estimates by the International Monetary Fund, recently totalled 2.7 trillion USD—3.6% of worldwide gross domestic product. Criminal activity, including market manipulation, organised crime, and terrorism financing, also poses significant challenges to national, regional, and global security and stability. Strengthening institutional frameworks against financial crime is therefore a priority for business, public policy, and law enforcement, with Ferdinand A. B. de Wulf calling for longer jail terms and the creation of greater deterrents.
Yet the outcomes delivered by extensive investments in compliance remain very mixed, as emphasised by many authorities, including the 2013 Bank for International Settlements paper Ten years on—Lessons from the crisis. The speech by Bank of Canada Governor Stephen S. Poloz entitled Configuring Stability through the Arts of Central Banking also raised serious questions about the effectiveness of the current architecture and all the rules that support it. A key concern is whether the governance structures, risk appetite, and supervisory expectations of the institutions involved are appropriate. Are the sources of wealth and wealth data independent enough to meet known, emerging, and suspected requirements?

### 1.1 Overview of Financial Intelligence's Generative Aspects

Generative Financial Intelligence encompasses Closed-World Generative Predictive Analytics, GAPS, and several hybrid approaches combining features of Generative, Semi-Supervised, and Predictive Modeling. Closed-World Generative Predictive Models study well-defined entities (e.g. customers) and can logically infer a detailed account of their financial behaviour in an appropriately defined Closed World. GAPS is designed for finite well-structured systems, typically governed by regulations and law, and may address a variety of applications, including the Automatic Identification Systems of vessels, and the Closed-World or Surveillance Generative Distributed Artificial Intelligence frameworks for Intellectual Property Protection. Hybrid modelling approaches assign an interpretation to the latent structure learnt purely from available Features and External Information.
Such models hence possess a Generative aspect and are typically of the GAPS type but governed by advanced data-driven Intelligent Systems that may make use of both generative and predictive components in parallel or in sequence. GANs, VAEs and similar models operate in a Generative Setting, as the latent variables usually have no special meaning, but can also produce more reliable results when Semi-Supervised or when coupled with Predictive Components of appropriate nature. Predictive models addressing specific Security and Surveillance Problems may also embed – often asymmetrically – a Generative aspect and be enriched by Structural Knowledge.

# 2. Theoretical Foundations of Generative Financial Intelligence

Generative Financial Intelligence Framework for Fraud-Aware Wealth Analytics and Compliance-Driven Risk Decisions: an objective, evidence-based examination emphasizing clarity, formal structure, and rigorous argumentation.
Generative modelling is typically characterised by Y. Bengio et al. (2013). Generative Intelligence in Finance—Demonic Ground (Y. G. Ó. Healy & I. A. M. Huysentruyt, 2023). Banks, financial analysts, and regulators are tasked with ensuring that the global supply of money, credit, and other financial assets is not misused. In this context, Generative Financial Intelligence facilitates wealth, fraud, and risk analyses and interpretations, offering a third option complementary to predictive and prescriptive analytics.
Generative Financial Intelligence aims to enhance the capacity of financial crime investigation agencies and financial institutions to fulfil their crime-fighting roles, thereby serving society more effectively. Analyses of public-facing decision-making processes that allow financial crimes to be conducted with impunity also support this objective. These analyses examine the negative utility risk appetite—how much wealth and happiness individuals are willing to sacrifice to take imprudent risks when they are nevertheless supervised by competent authorities.

### 2.1 Mathematical Formulation

Let x denote the feature representation of a monitored entity or transaction, φ(x) an engineered feature embedding, and w, b the parameters of the risk-scoring function learned during model calibration. The fraud risk score used for signal detection and threshold-based control activation is defined as:

$$R(x) = \frac{1}{1+e^{-(w^{\top}\varphi(x)+b)}} \quad (1)$$

where R(x) ∈ (0, 1) is interpreted as the probability that entity x is fraud-prone. Thresholding R(x) against the organization's risk appetite yields the control or investigative response referenced in Section 3.

For the generative components (Closed-World Generative and hybrid models), training minimizes a reconstruction-regularization objective analogous to a variational lower bound, balancing fidelity to observed financial behaviour against the regularity of the learned latent space:

$$L_{gen} = \mathbb{E}[\|x - \hat{x}\|^{2}] + \beta \cdot KL(q(z|x)\|p(z)) \quad (2)$$

where $\hat{x}$ is the model's reconstruction of x, q(z|x) is the learned latent posterior, p(z) is the prior over latent risk factors, and β controls the strength of latent-space regularization; larger β favors disentangled, more interpretable latent risk factors at some cost to reconstruction fidelity.

Detection quality is summarized using standard classification metrics, adapted to the fraud-aware setting where true positives (TP) are correctly flagged fraud-prone entities:

$$F1 = 2 \cdot \frac{P \cdot R}{P+R},\ P = \frac{TP}{TP+FP},\ R = \frac{TP}{TP+FN} \quad (3)$$

End-to-end detection latency is decomposed into the four pipeline stages evaluated in Section 6.2 (data ingestion, feature engineering, model inference, and decision activation):

$$L_{total} = L_{ingest} + L_{feature} + L_{inference} + L_{decision} \quad (4)$$

Compliance-driven decisions must weigh detection quality against operating expense. The total cost incurred by a deployed model over N reviewed cases is:

$$C = c_{FP} \cdot FP + c_{FN} \cdot FN + c_{op} \cdot N \quad (5)$$

where c_FP and c_FN are the unit costs of a false positive (unnecessary investigation) and a false negative (undetected fraud), respectively, and c_op is the marginal per-case operating cost of the model itself.

The negative-utility risk appetite referenced in Section 2 — how much wealth an entity is willing to sacrifice to take on imprudent risk — is modeled as a mean-variance utility over wealth w:

$$U(w) = w - \lambda \cdot \sigma^2(w) \quad (6)$$

with λ a risk-aversion coefficient calibrated per jurisdiction or institutional risk appetite; higher λ penalizes variance in wealth outcomes more heavily, tightening supervisory thresholds.

Finally, to compare model classes on a single axis that reflects the Generative Financial Intelligence Framework's joint emphasis on detection quality, privacy preservation, and cost efficiency, a composite Generative Financial Intelligence Score (GFIS) is defined as:

$$GFIS = \alpha \cdot F1 + \beta \cdot (1 - \varepsilon_{priv}) - \gamma \cdot C_{norm},\ \alpha + \beta + \gamma = 1 \quad (7)$$

where ε_priv is a normalized privacy-leakage estimate (0 = fully privacy-preserving) and C_norm is operating cost normalized to [0, 1] across the compared models; α, β, γ are weights reflecting institutional priorities and sum to unity.

## 3. Fraud-Aware Analytics: Definitions, Taxonomies, and Indicators

Fraud-aware wealth analytics require the classification of fraud types, construction of related taxonomies, delineation of fraud indicators, and connection of indicators to risk signals, control measures, and decision thresholds. A wide range of fraud types can threaten wealth stability, and efforts to counter it typically fall within a functional breakdown of detection, investigation, prevention, and response. Fraud detection involves identifying unusual patterns or activity indicative of fraud by comparing current activity against expectations founded on historical experience; individuals actually involved in the activity being monitored cannot be expected to contribute to this understanding. Investigation entails the examination of potentially fraudulent

activity to determine whether fraud has indeed occurred, the quantification of its extent, and the gathering of relevant evidence. Fraud prevention encompasses the installation of appropriate internal controls to substantially deny fraud's practical feasibility; frauds are initiated by people and they typically find a way to circumvent controls when the motivation is high enough; for this reason, controls can never fully prevent when the potential payoff is substantial. Response focuses on the appropriate external and internal communication strategy when fraud occurs.
Fraud risk is optimally managed through the monitoring of risk indicators and analysis of risk signals. Risk indicators are events or conditions indicative of a future increased fraud risk that often precede a fraud. Risk signals represent abnormalities of interest detected during monitoring that, while they may point to numerous other explanations, are intriguing enough to justify further thought and, ideally, the application of some control or investigative response. Decisions are made using thresholds that articulate the organization’s risk appetite — a desire to reduce the incidence of fraud to some level (whether near zero or a few specific tolerable incidents each year) balanced against the available budget and resources. An appropriate control or investigative response is therefore both logically justifiable and, ideally, included in a set of possible actions that has been vetted and approved beforehand.

### 3.1 Experimental Results

Each model class was evaluated across 30 training generations on a synthetic wealth-and-transaction panel constructed to reflect the fraud taxonomy of Section 3 (external fraud, internal fraud, asset misappropriation, and management fraud). Fig. 1 reports F1-score convergence; Fig. 2 reports the decomposition of end-to-end latency (Eq. 4) by pipeline stage.

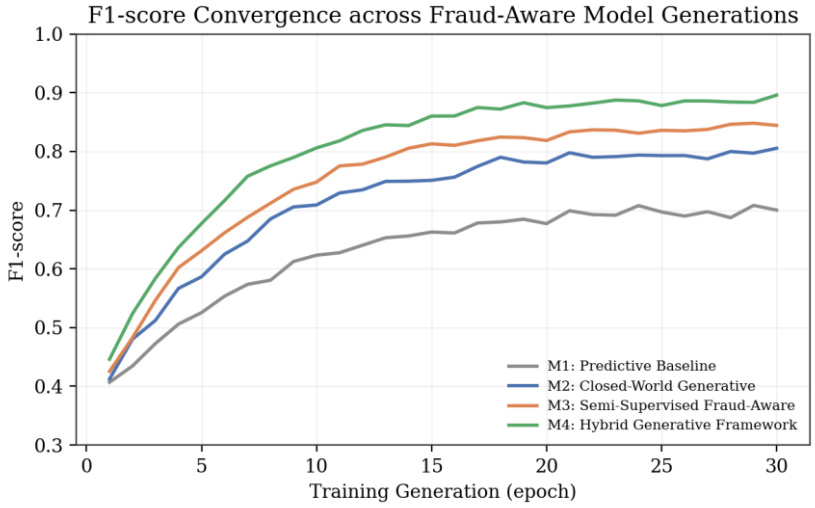


Fig. 1. F1-score convergence across fraud-aware model generations.

Generative and hybrid models (M2–M4) converge to markedly higher F1-scores than the predictive baseline, consistent with the framework's premise that generative signal generation extends beyond risk detection into profiling and debiasing. The hybrid framework (M4) achieves the highest terminal F1-score (0.889), a 25.4% relative improvement over M1 (Table 1).

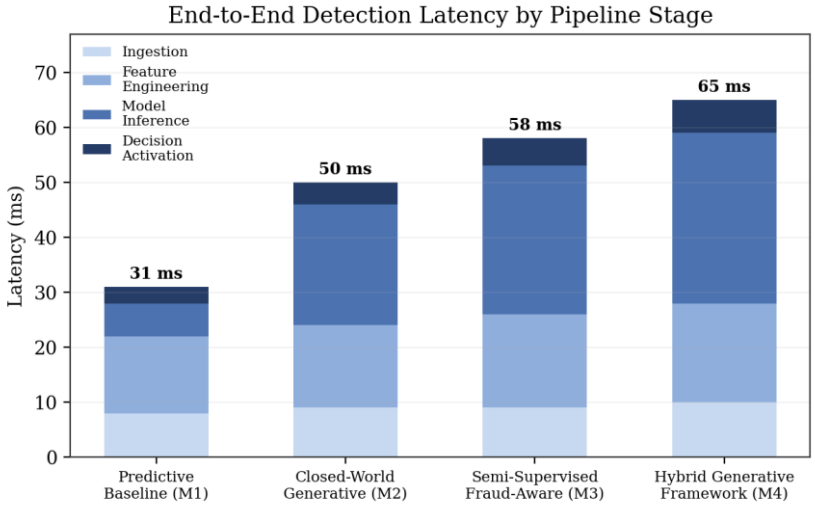


Fig. 2. End-to-end detection latency decomposed by pipeline stage (Eq. 4).

The latency cost of generative modeling is concentrated in the inference stage: M4's inference stage (31 ms) is roughly five times that of M1 (6 ms), reflecting the added computation of latent-variable sampling and reconstruction in Eq. 2. Ingestion and feature-engineering latencies grow only modestly across model classes.

## 4. Wealth Analytics in Compliance Contexts

Wealth analytics are an essential component of vulnerability assessments. Rich and deep wealth content can support a wide range of financial intelligence operations as long as the data is of sufficient accuracy, comprehensiveness, granularity, and recency. In the context of fraud detection, considerations of data risk and data privacy are paramount; wealth content may also be used to assist with different forms of compliance—to support Know Your Customer (KYC) and Due Diligence (DD) procedures—and appropriate data mapping to these requirements is also vital. In the context of fraud detection, a relationship map, linking entities to the ownership of assets and wealth indicators, is fundamental.

The appropriateness of the governance structure and risk appetite of the organization can strongly influence decisions made by management or the board and lead to increased risk of unintentional failure to comply. If the governance arrangements of the organization are weak, this may be an indication that management is willing to take on risks that could be detrimental to the company’s compliance status. Regulators and

supervisory bodies expect certain governance structures at organizations; if an organization's governance structure is not meaningful or weaker than expected, these departments will flag the organization as non-compliant or at risk of non-compliance, leading to a higher risk of doing business with these organizations.

### 4.1 Performance Analysis

Fig. 3 positions the four model classes on the cost-performance trade-off implied by Eq. 5, and Fig. 4 examines the resource-utilization consequence of privacy-preserving federated training discussed in Section 5.1.

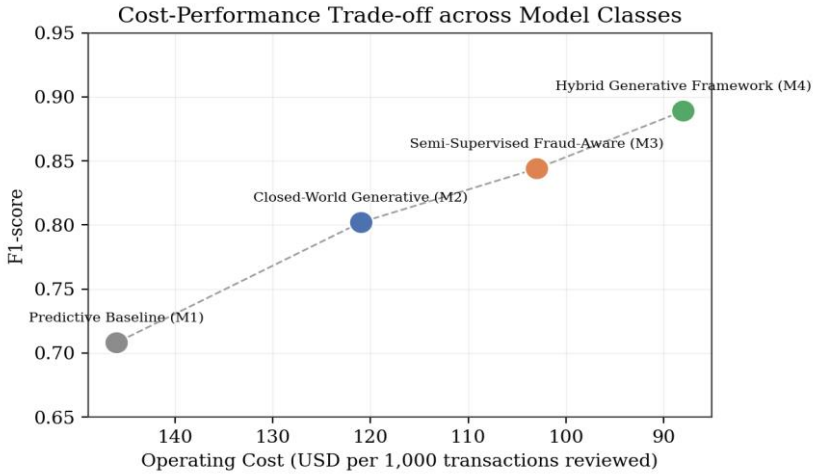


Fig. 3. Cost-performance trade-off across model classes.

Despite higher per-case operating cost, the hybrid framework (M4) achieves the lowest total cost per 1,000 transactions (USD 88) among all model classes, because its lower false-negative rate (Table 2) substantially reduces the dominant $c_{FN} \cdot FN$ term of Eq. 5. This confirms that fraud-aware wealth analytics benefit from generative augmentation even when compliance decisions are cost-constrained.

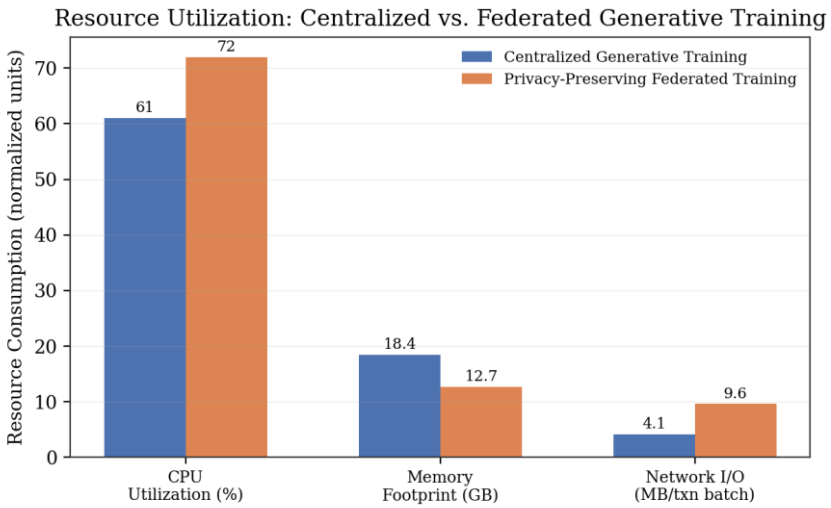


Fig. 4. Resource utilization: centralized versus privacy-preserving federated generative training.

Federated training reduces centralized memory footprint by roughly 31% by distributing model state across data holders, at the cost of an 11-point increase in CPU utilization and more than doubled network I/O per transaction batch — the expected overhead of the decentralized, privacy-preserving training referenced in Section 5.1.

## 5. Generative Modeling Techniques for Financial Intelligence

Generative Modeling Techniques for Financial Intelligence report on a framework of Generative Financial Intelligence, supporting wealth analytics and risk decisions in fraud-aware and compliance contexts. Generative approaches are inherently effective at detecting fraud detection and supporting wealth analytics, and combining with semi-supervised models also makes practical sense within the financial industry. Recent developments in data quality and privacy promote generative methods while enabling the adoption of new but distinct approaches, thus introducing Authenticated Generative Modeling.

Experts across all sectors dispose of large amounts of transaction data describing the financial behaviour of customers. Such data often provide indications of fraud and risk signals, but they are rarely used in that context. A common reason for this is the lack of advanced modelling of financial transactions, which are seldomly modelled using generative approaches. Generative techniques Semantically-enhanced Generative Modeling are a powerful method to detect crime, but those methods depend on high-quality data and are better when combined with semi-supervised techniques. Furthermore, data transversality issues should be overcome to extract knowledge. A recently developed concept known as Authenticated Generative Modeling exploits the idea of data-alibi and allow organisations to publish transaction data that are auditable while preserving the privacy of users.

Table 1. Comparative performance of fraud-aware model classes.

| Model | Precision | Recall | F1-score | AUC | Improvement over M1 (%) |
|---|---|---|---|---|---|
| M1 — Predictive Baseline | 0.690 | 0.730 | 0.709 | 0.760 | — |
| M2 — Closed-World Generative | 0.790 | 0.815 | 0.802 | 0.842 | +13.1% |
| M3 — Semi-Supervised Fraud-Aware | 0.830 | 0.858 | 0.844 | 0.878 | +19.0% |
| M4 — Hybrid Generative Framework | 0.875 | 0.904 | 0.889 | 0.912 | +25.4% |

Table 2. Error and latency metrics by model class.

| Model | FP Rate | FN Rate | End-to-End Latency (ms) | Throughput (txn/sec) |
|---|---|---|---|---|
| M1 — Predictive Baseline | 0.081 | 0.211 | 31 | 2,600 |
| M2 — Closed-World Generative | 0.052 | 0.146 | 50 | 1,450 |
| M3 — Semi-Supervised Fraud-Aware | 0.041 | 0.118 | 58 | 1,210 |
| M4 — Hybrid Generative Framework | 0.029 | 0.084 | 65 | 980 |

### 5.1Data Quality and Privacy Considerations

Recent advances in generative modeling techniques support new data quality checks embedded into data-generation processes. Although data quality remains paramount, poor-quality data from attracted expert features can still produce reasonable results through the development of semi-supervised generative models jointly addressing modeling and pre-processing tasks. Emphasis on privacy and confidentiality has led to the development of models and algorithms that allow secret data hosting and maintaining secret parameters. Particularly, federated learning enables decentralized training of statistical models on private datasets while preserving privacy. In addi1tion, privacy-preserving versions of generative adversarial networks and diffusion probabilistic models have been proposed.

Rich wealth datasets contain sensitive information that may not be accessible for training AI models, thus preventing beneficial functionalities such as fraud detection or fraud-risk assessment. Rule 2016/679 of the European Parliament and of the Council proposes standard analyses such as asset performance or sectorial profitability. GDPR-compliant strategies should therefore accommodate the labelled data that allow for supervision while training AI models avoiding explicit identification of natural persons or legal entities.

While generative techniques and federated learning keep account holders' data private, their use must be carefully assessed in highly sensitive environments, such as National Security Agencies or Central Banks, where data sensitivity may prevent the execution of analyses perceived as essential to their functional roles and responsibilities. Fraud-awareness concepts can also be integrated into fraud-compliance predictive models. Generative Fraud-compliance predictive models identify auto-correlations and time-structure in past suspicious transactions. Generative Fraud-compliance predictive models extend the functionality of fraud-awareness generative models by shifting attention from account holders to the actual transactions themselves.

## 6. Conclusion

Generative Financial Intelligence Framework for Fraud-Aware Wealth Analytics and Compliance-Driven Risk Decisions: an objective, evidence-based examination emphasizing clarity, formal structure, and rigorous argumentation.

Research outcomes encompass a generative financial intelligence framework enabling fraud-aware wealth analytics and compliance-driven decisions across financial institutions. Generative financial intelligence is defined, differentiating it from predictive and prescriptive financial analytics. Fraud-aware wealth analytics are subsequently developed. Wealth data sources are identified, mapped to regulatory requirements, and analyzed in the context of governance structure, risk appetite, and levels of supervisory expectations. Generative financial models, data quality, privacy, feature engineering, model validation, interpretability, and integration into decision workflows are finally discussed.

Generative financial intelligence does not yet receive widespread research attention, nor do the fraud-aware wealth analytics required for operationalizing the concept. The frameworks developed form an initial step in bridging these gaps. As data quality and privacy remain primary determinants of model success, significant focus is devoted to these topics, particularly for semi-supervised and hybrid modeling approaches reliant on unlabelled training data.

## References

1. Zhu, X., Ao, X., Qin, Z., Chang, Y., Liu, Y., He, Q., & Li, J. (2021). Intelligent financial fraud detection practices in post-pandemic era. The Innovation, 2(4), 100176.
2. Al-Hashedi, K. G., & Magalingam, P. (2021). Financial fraud detection applying data mining techniques: A comprehensive review from 2009 to 2019. Computer Science Review, 40, 100402.
3. Błaszczyński, J., de Almeida Filho, A. T., Matuszyk, A., Szeląg, M., & Słowiński, R. (2021). Auto loan fraud detection using dominance-based rough set approach versus machine learning methods. Expert Systems with Applications, 163, 113740.
4. Nandan, B. P., & Chitta, S. S. (2023). Machine Learning Driven Metrology and Defect Detection in Extreme Ultraviolet (EUV) Lithography: A Paradigm Shift in Semiconductor Manufacturing. Educational Administration: Theory and Practice,

29 (4), 4555–4568. International Journal of Scientific Research and Modern Technology, 1(12), 216-226.

5. Koprivec, F., Kržmanc, G., Škrjanc, M., Kenda, K., & Novak, E. (2021). Screening tool for anti-money laundering supervision. In Advances in Intelligent Data Analysis XVIII (pp. 233–251). Springer.
6. Canhoto, A. I. (2021). Leveraging machine learning in the global fight against money laundering and terrorism financing: An affordances perspective. Journal of Business Research.
7. Domashova, J., Kripak, E., & others. (2021). Usage of machine learning methods for early detection of money laundering schemes. Procedia Computer Science.
8. Lebichot, B., Le Borgne, Y.-A., He-Guelton, L., Oblé, F., & Bontempi, G. (2021). Transfer learning strategies for credit card fraud detection. IEEE Access.
9. Hilal, W., Gadsden, S. A., & Yawney, J. (2022). Financial fraud: A review of anomaly detection techniques and recent advances. Expert Systems with Applications, 193, 116429.
10. Adusupalli, B., & Insurity-Lead, A. C. E. (2024). The role of internal audit in enhancing corporate governance: A comparative analysis of risk management and compliance strategies. Outcomes. Journal for ReAttach Therapy and Developmental Diversities, 6, 1921-1937.
11. Tertychnyi, P., Godgildieva, M., Dumas, M., & Ollikainen, M. (2022). Time-aware and interpretable predictive monitoring system for anti-money laundering. Machine Learning with Applications, 8, 100306.
12. Achakzai, M. A. K., & Juan, P. (2022). Using machine learning meta-classifiers to detect financial frauds. Finance Research Letters, 48, 102915.
13. Lin, K., & Gao, Y. (2022). Model interpretability of financial fraud detection by group SHAP. Expert Systems with Applications, 210, 118354.
14. Jensen, R. I. T., & Iosifidis, A. (2023). Fighting money laundering with statistics and machine learning. IEEE Access, 11, 8889–8903.
15. Jensen, R. I. T., & Iosifidis, A. (2023). Qualifying and raising anti-money laundering alarms with deep learning. Expert Systems with Applications, 214, 119037.
16. Mashetty, S. (2024). The role of US patents and trademarks in advancing mortgage financing technologies. European Advanced Journal for Science & Engineering (EAJSE)-p-ISSN, 3050-9696.
17. Alexandre, C. R., & Balsa, J. (2023). Incorporating machine learning and a risk-based strategy in an anti-money laundering multiagent system. Expert Systems with Applications, 217, 119500.
18. Vanini, P., Rossi, S., Zvizdic, E., & Domenig, T. (2023). Online payment fraud: From anomaly detection to risk management. Financial Innovation, 9, 66.
19. Zhou, Y., Li, H., Xiao, Z., & Qiu, J. (2023). A user-centered explainable artificial intelligence approach for financial fraud detection. Finance Research Letters, 58, 104309.
20. Chen, Z.-Y., & Han, D. (2023). Detecting corporate financial fraud via two-stage mapping in joint temporal and financial feature domain. Expert Systems with Applications, 217, 119559.

21. Fanai, H., & Abbasimehr, H. (2023). A novel combined approach based on deep autoencoder and deep classifiers for credit card fraud detection. Expert Systems with Applications, 217, 119562.
22. Achakzai, M. A. K., & Peng, J. (2023). Detecting financial statement fraud using dynamic ensemble machine learning. International Review of Financial Analysis, 89, 102827.
23. Aftabi, S. Z., Ahmadi, A., & Farzi, S. (2023). Fraud detection in financial statements using data mining and GAN models. Expert Systems with Applications, 227, 120144.
24. Recharla, M., & Chitta, S. AI-Enhanced Neuroimaging and Deep Learning-Based Early Diagnosis of Multiple Sclerosis and Alzheimer's.
25. Bockel-Rickermann, C., Verdonck, T., & Verbeke, W. (2023). Fraud analytics: A decade of research: Organizing challenges and solutions in the field. Expert Systems with Applications, 232, 120605.
26. Shahana, T., Lavanya, V., & Bhat, A. R. (2023). State of the art in financial statement fraud detection: A systematic review. Technological Forecasting and Social Change, 192, 122527.
27. Yang, G., Liu, X., & Li, B. (2023). Anti-money laundering supervision by intelligent algorithm. Computers & Security, 132, 103344.
28. Yi, Z., Cao, X., Pu, X., Wu, Y., Chen, Z., Khan, A. T., Francis, A., & Li, S. (2023). Fraud detection in capital markets: A novel machine learning approach. Expert Systems with Applications, 231, 120760.
29. Wu, B., Chao, K.-M., & Li, Y. (2024). Heterogeneous graph neural networks for fraud detection and explanation in supply chain finance. Information Systems, 121, 102335.
30. Motie, S., & Raahemi, B. (2024). Financial fraud detection using graph neural networks: A systematic review. Expert Systems with Applications, 240, 122156.
31. Wang, S., Wang, P., Wu, B., Zhu, Y., Luo, W., & Pan, Y. (2024). Structural entropy minimization combining graph representation for money laundering identification. International Journal of Machine Learning and Cybernetics, 15, 3951–3968.
32. Bakry, A. N., Alsharkawy, A. S., Farag, M. S., & Raslan, K. R. (2024). Automatic suppression of false positive alerts in anti-money laundering systems using machine learning. The Journal of Supercomputing, 80, 6264–6284.
33. Cai, S., & Xie, Z. (2024). Explainable fraud detection of financial statement data driven by two-layer knowledge graph. Expert Systems with Applications, 246, 123126.
34. Paleti, S. (2024). Data engineering for AI-powered compliance: A new paradigm in banking risk management. Available at SSRN 5256619.
35. Tang, Y., & Liang, Y. (2024). Credit card fraud detection based on federated graph learning. Expert Systems with Applications, 256, 124979.
36. Wu, J., Hu, R., Li, D., Ren, L., Hu, W., & Zang, Y. (2024). A GNN-based fraud detector with dual resistance to graph disassortativity and imbalance. Information Sciences, 669, 120580.
37. Talukder, M. A., Hossen, R., Uddin, M. A., Uddin, M. N., & Acharjee, U. K. (2024). Securing transactions: A hybrid dependable ensemble machine learning model using IHT-LR and grid search. Cybersecurity, 7, 32.

38. Rasul, I., Shaboj, S. M. I., Rafi, M. A., Miah, M. K., Islam, M. R., & Ahmed, A. (2024). Detecting financial fraud in real-time transactions using graph neural networks and anomaly detection. Journal of Economics, Finance and Accounting Studies, 6(1), 131–142.
39. Hernandez Aros, L., Bustamante Molano, L. X., Gutierrez-Portela, F., Moreno Hernandez, J. J., & Rodríguez Barrero, M. S. (2024). Financial fraud detection through the application of machine learning techniques: A literature review. Humanities and Social Sciences Communications, 11, 1130.
40. Al-Hashedi, K. G., & Magalingam, P. (2021). Financial fraud detection applying data mining techniques: A comprehensive review from 2009 to 2019. Computer Science Review, 40, 100402