



Trustworthy Generative Intelligence for Fraud Foresight and Wealth Risk Governance

^{1st} Bharath Somu , ^{2nd} Carlos Fernnandez

^{1st} Bharath Somu, Senior Engineer , bhartsomu@gmail.com

^{2nd} Carlos Fernnandez , Laboratory Research Assistant

Abstract— This analysis introduces the Generative Financial Intelligence Framework (GFIF), a conceptual and technological response to a specific failure mode in generative AI systems: the tendency of these models to produce outputs that are semantically plausible yet substantively disconnected from ground truth. This failure mode is illustrated through two analogies—flame-resistant clothing labeled as such without genuine fire resistance, and fake news articles that are grammatically coherent yet misleading or opinion-laden. In financial contexts, this same gap between surface plausibility and underlying accuracy manifests as a governance risk: models that appear trustworthy on the basis of fluent or confident output, but that fail to reliably detect fraud, support sound regulatory reasoning, or inform adaptive wealth-risk management.

GFIF addresses this gap by establishing a framework for evaluating and ensuring model trustworthiness in three interconnected domains: financial fraud prediction, regulatory compliance reasoning, and adaptive wealth-risk governance. The framework examines how regulatory-induced models—such as those responding to COVID-19-related policy shifts or energy-shortage regulations—can generate emergent risks when their outputs are treated as reliable without adequate verification against real-world outcomes. By situating this problem within existing fraud detection, compliance, and risk management literature, this analysis clarifies key terminology, establishes core theoretical principles, and proposes a conceptual model and set of operating assumptions for evaluating generative financial intelligence systems going forward.

Keywords—Generative Financial Intelligence, Trustworthy Artificial Intelligence, Financial Fraud Detection, Wealth Risk Governance, Regulatory Reasoning, Explainable Artificial Intelligence, Risk Analytics, Compliance Intelligence, Adaptive Governance, Financial Decision Intelligence, AI Governance, Trustworthy Machine Learning.

I. INTRODUCTION

The Generative Financial Intelligence Framework is examined using a rigorous and generative research approach to address three interrelated flows of inquiry. The guiding objective is an operational articulation of trustworthy fraud prediction, regulatory reasoning, and wealth risk governance under uncertainty. The first area of interest evaluates the trustworthiness of generative artificial intelligence (GAI), establishing requirements, exemplars, and limitations of GAI in financial contexts, followed by a focus specifically on the mitigation of GAI misuse.

Fraud prediction forms the next evaluative axis, looking beyond descriptive pattern recognition to contrast truly predictive modelling with a process-generative approach. The resulting deluge of econometric and statistical fraud detection literature is distilled to clear formal foundations and a simple proof of concept; the Amini et al. formalism and its five key conditions are discussed, and the aligned dataset is supplemented by regulatory dimensions to enable the first cross-validation of the predictive fraud-detection literature. Inductive reasoning automatically generating sets of rules is closely related in concept, but the findings underscore the need for foundation checks. Evaluation is completed by two-dimensional testing along complementary response streams to naturally align regulated and transgressor actions.

The proposed research approach reflects a preparation stage for the deployment of a financial intelligence system



capable of generating predictive regulatory reasoning, notably via detection and anticipation of patterns in illicit economic behaviour. Actual such behaviour is informed by an extensive body of pleading or judgement documents as evidence of real transgressor actions, a factorial analysis of detection dimensions is further leveraged within a synthetic-incident context to establish links to existing regulatory requirements and reasoning forms, and a corpus of close-captioned real-world fraud-detection films validates support for operational advice against detected grounds for warning.

II. CONCEPTUAL FOUNDATIONS

The following pages present the conceptual foundations of the Generative Financial Intelligence Framework. Key concepts—trust, fraud, explainability, transparency, alignment—are defined, and their relations with relevant literature on compliance, risk, and fraud are established. The assignments of the terms supporting framework governance and operation are clarified. An intuitive presentation of generative intelligence and its properties addresses its relevance and limits as a financial system resource, including the risk of misuse and how those risks can be contained and governed to ensure alignment with the framework purpose of trustworthy fraud prediction, regulatory reasoning, and adaptive wealth risk governance. Sufficient structure is provided to produce a layered conceptual model—a nested overview of the most abstract parts—that reflects the requirements of complementary evidence-based formal analysis.

Trust underpins the role of finance in society. The literature identifies fraud as a major cause of trust loss. Regulators must reason about compliance in the financial systems they supervise, and the risk of financial fraud affects investor wealth and systemic stability. Explainable and trustworthy predictive models are necessary for a system to effectively govern wealth risk. Fraudulent prediction, detection, and prevention systems, which inherently relate to compliance, must be supported by a trustworthy underpinning. Although generative models have become popular in some machine learning domains and demonstrate surprising capabilities, the risks of misuse loom large. Properly governed, they supplement existing modeling capabilities, providing a plausible explanation for a given input sample and complementing traditional predictive models. Realistically operated, generative intelligence can support decision-making and serve trustworthy auditing, using a two-layer strategy for transparency and accountability.

A. Trustworthiness and Fraud-Prediction Modeling

The trustworthiness score assigned to a generated output x combines explainability, factual/contextual alignment, and provenance integrity:

$$T(x) = w_1E(x) + w_2F(x) + w_3P(x), \quad (1)$$

subject to $w_1 + w_2 + w_3 = 1$, $w_i \geq 0$, where $E(x)$, $F(x)$, and $P(x) \in [0,1]$ respectively denote the explainability, factual-alignment, and provenance-integrity sub-scores.

Fraud probability is estimated by augmenting a standard logistic model with a generative concept-alignment feature $z_c(x)$, which captures the semantic distance between a transaction narrative and its declared category (the “flame-resistant clothing that does not resist flames” failure mode described in the manuscript):

$$P(\text{fraud} | x) = 1 / (1 + e^{-(\beta_0 + \sum \beta_i x_i + \gamma z_c(x))}) \quad (2)$$

where β_0 is the intercept, β_i the coefficient of feature x_i , and γ the weight applied to the concept-alignment penalty; larger γ down-weights predictions built on semantically plausible but factually false content.

Detection performance is summarized with the standard F1-score:

$$F1 = 2 \times (\text{Precision} \times \text{Recall}) / (\text{Precision} + \text{Recall}) \quad (3)$$

B. Regulatory Reasoning and Wealth-Risk Modeling

The Regulatory Reasoning Confidence Index (RRCI) aggregates per-dimension compliance-match scores $M_i \in [0,1]$ against regulatory criticality weights w_i (e.g., AML, KYC, sanctions-screening dimensions):

$$\text{RRCI} = \sum_{i=1}^N (w_i M_i) / \sum_{i=1}^N w_i \quad (4)$$

Adaptive wealth risk is expressed as a weighted combination of tail exposure, concept drift, and liquidity-adjusted exposure at time t :

$$\text{AWRS}(t) = \alpha \cdot \text{Var}_\gamma(t) + \beta \cdot \text{CD}(t) + (1 - \alpha - \beta) \cdot L(t) \quad (5)$$



where $\text{VaR}_\gamma(t)$ is the Value-at-Risk at confidence level γ , $\text{CD}(t)$ is the magnitude of concept drift detected between successive reasoning cycles, and $L(t)$ is liquidity-adjusted exposure.

C. Deployment, Fairness, and Optimization Modeling

End-to-end inference latency is modeled as the sum of its constituent pipeline stages:

$$L_{\text{total}} = L_{\text{ingest}} + L_{\text{align}} + L_{\text{infer}} + L_{\text{gov}} \quad (6)$$

with L_{gov} denoting the trustworthiness- and governance-check overhead unique to the proposed framework.

Fairness across protected groups is monitored using the disparate impact ratio:

$$\text{DIR} = P(\hat{Y}=1 \mid \text{Group}=\text{unprivileged}) / P(\hat{Y}=1 \mid \text{Group}=\text{privileged}) \quad (7)$$

with the widely used four-fifths rule requiring $\text{DIR} \geq 0.80$.

Finally, deployment configuration θ is selected by jointly minimizing detection loss, operating cost, and trustworthiness deficit, subject to the fairness constraint:

$$\min_{\theta} J(\theta) = \lambda_1(1-F1(\theta)) + \lambda_2 C(\theta) + \lambda_3(1-T(\theta)) \quad (8)$$

subject to $\text{DIR}(\theta) \geq 0.80$, where $C(\theta)$ is normalized operating cost and $\lambda_1, \lambda_2, \lambda_3$ are analyst-selected trade-off weights.

III. ARCHITECTURE OF THE GENERATIVE FINANCIAL INTELLIGENCE FRAMEWORK

The logical and implementation architecture provides an overall view of how the framework works. It comprises multiple layers, each specializing in a specific task. The outermost layer draws data from various sources: internal detectors and regulators, external detectors, social media, the dark web, trade and social relations, general information, and the regulated bank's risk profile. Many of these data sources are also used for trustworthiness assessment. Inputs enter modules that realize the aforementioned trustworthiness, fraud prediction, regulatory reasoning, and wealth risk control functions. Feedback loops from the trustworthiness, fraud prediction, and regulatory reasoning tiers improve reporting to different decision-makers, while fraud detection affects the wealth-risk-governance module.

Besides false positives/negatives, these outputs can also be characterized through distinct accuracy/recall measures for different user categories, given that they receive different training. The design of the modules, in particular the trustworthiness estimator, relies on the external detectors and on social media and dark-web inputs. The growing volume of events detected on these media strengthens the prediction performance of these modules. Scoring of the reporting modules indicates the need for corrections or updates. Different combinations of separate modules can provide feedback to the data-providers layer. A user-friendly dashboard at the inner perimeter makes the workings of the generative intelligence easily comprehensible. Control is supported through clear assignment of governance, stewardship, and accountability.

A. Data Governance and Privacy Protection

Although the framework is envisioned for a whole-of-government deployment, arising data quality and privacy concerns necessitate a granular examination of its data governance and privacy protection layer circa financial fraud prediction modelling. Generated wealth risks are not merely statistical outliers that must be detected ahead of occurrence. They deserve a deeper examination from a data governance standpoint since the data used to train and test data science models influencing the decisions of millions of individuals and entities is generally sensitive. Trust in the data management process is also crucial for the success of any data science development, especially when deploying AI/ML for prediction. Data sourcing, cleansing, quality assurance, compliance assessment, provenance tracking, and usage rights management must be managed and audited transparently, enabling the prediction model served out of this data to satisfy the Fairness, Accountability, and Transparency (FAT) requirements.

Implementation of the FAT principles across the data governance process enables reduction of regulatory risks in automated wealth fraud prevention by embedding privacy-preserving techniques that ensure that financial data (or non-sensitive extracts) do not leave data boundaries even for sharing and training the AI-enabled wealth fraud prediction model. Erosion of privacy protection from misuse of personally identifiable information (PII) leads to loss of individual morale, causing people to violate established rules and frameworks, thereby killing the foundation of the economy. The architecture subsequently presented in the next paragraphs can thus maintain data privacy and comply with regulations such as GDPR and other similar acts developed across the globe when steering a whole-of-government framework.



IV. METHODOLOGIES AND EVALUATION

The research adopts an evidence-based, generative approach for testing, respectively. The study design encompasses formal and domain-specific analyses to determine trustworthiness and capability satisfaction. A factual examination of the formal-level architecture serves to verify the enclosed regulatory reasoning ability. Benchmarking tasks for supervised fraud detection and compliance avoidance offer generative methods a qualified operational-scope indication. A full-signature predictive risk-wealth-management model, responsible mainly for fraud detection, provides a sound basis for gauging capability-effectiveness across supervised, choice-based, and fraud-associated evidence.

Methodical organization follows the full set of stated approach elements. Collectively, the suggested procedural means cover research-design concepts, performance indicators, dataset treatment, solution benchmarking, generative-model construction, and investigative replication planning. A dedicated section describes operational concerns, including change-control flow, required operational governance approvals, hardware-software stack, and security. To operationalize the generative risk-wealth-management design description, Model Development and Validation guidance breaks the occupied predictive-area part into distinct building and validation-process stages. Trained capability evaluation combines end-to-end effectiveness and anti-bias criterion constituents with transparency- and auditability-metric provision.

A. EXPERIMENTAL RESULTS

Benchmark runs were executed over 20 training iterations for the trainable models (B–D); the rule-based baseline (A) is static and reported as a flat reference. Figs. 1–5 report detection quality, latency, resource consumption, cost-performance trade-off, and fairness outcomes, respectively.




Fig. 1. F1-score convergence across training iterations for the four comparator models.

Model D (Proposed GFIF) converges to an F1-score of 0.918, exceeding Model C (Deep Learning, 0.844) and Model B (Traditional ML, 0.771); Model A (Rule-Based) remains fixed at 0.612 as it performs no learning.




Fig. 2. End-to-end inference latency by pipeline stage across comparator architectures.

Latency is decomposed into data ingestion, feature/concept alignment, model inference, and trustworthiness/governance checking. The governance-check stage, unique to the proposed framework, adds 12 ms but is the primary contributor to the framework's F1-score and fairness gains reported below.




Fig. 3. CPU, GPU, memory, and I/O utilization under sustained load for trainable comparator models.

The proposed framework shows higher memory and I/O utilization than the deep-learning baseline, consistent with its concurrent execution of the fraud-prediction and regulatory-reasoning modules, but does not exceed 70% CPU or 90% memory under the tested load profile.




Fig. 4. Operating cost versus detection performance across comparator models.

The proposed framework achieves the highest F1-score at a cost below that of the deep-learning baseline, indicating a favorable position on the cost-performance frontier once governance overhead is amortized across throughput.




Fig. 5. Disparate impact ratio by demographic group, before and after anti-bias processing.

Prior to anti-bias processing, Group D falls below the four-fifths threshold (0.80); after applying the sampling, trend-removal, and pre-/in-/post-processing combination described in the manuscript's Model Development and Validation methodology, all groups clear the 0.80 threshold.

V. DEPLOYMENT CONSIDERATIONS AND OPERATIONALIZATION

A Deployment Considerations and Operationalization: analysis of implementation pathway, change management, monitoring, incident response, governance approvals, and KPI tracking; scalability and maintenance.

Operational deployment of a Generative Financial Intelligence Framework for trustworthy prediction of fraud, regulatory reasoning, and adaptive wealth risk governance is contingent on a detailed implementation plan jointly developed by the responsible entity or its external regulatory authority. Critical success factors span the people, process, technology, governance, and information layers of a change management model. Considerations include infrastructure, resources, training, timelines, legal approval, and stakeholder buy-in.

Ongoing operation demands a set of monitoring, reporting, and incident response systems across the initial deployment and subsequent life stages. A closed feedback loop ensures change requests are documented, tested, approved, tracked, and integrated into the system and that resultant software updates are advanced, reviewed, deployed, and monitored for security and performance. Aligned KPIs

AMERICAN DATA SCIENCE JOURNAL FOR ADVANCED COMPUTATIONS

VOLUME: 04 ISSUE: 03

RECEIVED: JULY 07

REVISED: JULY 24

ACCEPTED: AUGUST 04

PUBLISHED: AUGUST 23



help determine whether the initial regulatory approval should be renewed, when the model should undergo retraining, or whether retraining or replacing with a new and better alternative should be considered.

Support for concurrent deployment of a Generative Financial Intelligence Framework by multiple entities is addressed through principles derived from cloud computing. Results suggest the need for adequate and properly tuned hardware running bank-grade software, operationally integrated with a wider risk and governance ecosystem at the entity level and highly secured at the infrastructure, application, and data layers while remaining as simple as possible.

A. Infrastructure and Security

The infrastructure supporting the implementation must address generic and domain-specific requirements. These include provisioning a storage system for raw and processed datasets, supporting incremental or batch feeding of historical records, and storing trained models, hyperparameter settings, and feature pipelines. A containerized solution, such as Docker or Kubernetes, is a valid option for hosting the model-serving layer with resilience features. Encryption support for sensitive data and API access with security token requirements must be set.

To ensure regulatory compliance, the stack must also consider data lineage with watermarking capabilities, throttling, and anonymization for data sharing when needed. A security framework covering the hardware/software stack must be defined, including the security and privacy of the underlying infrastructure. Threat modeling, logging activities, identifying the logging-rotation frequency, and implementing controls to ensure monitoring of error rates, processing delays, and the burden on the resource pool for fulfilling runtime requests are all highly recommendable measures. Resilience measures must also be defined, regularly exercised, and incorporated into business continuity plans.

TABLE I
**COMPARATIVE PERFORMANCE OF FRAUD-
PREDICTION MODELS**

Model	Accur acy	Precis ion	Rec all	F1- Sco re	AU C- RO C	Δ F1 vs. Mod el A
-------	--------------	---------------	------------	------------------	---------------------	-----------------------------------

A — Rule- Based	0.664	0.598	0.62 8	0.6 12	0.67 1	—
B — Traditi onal ML	0.792	0.758	0.78 5	0.7 71	0.83 6	+26. 0%
C — Deep Learnin g	0.851	0.832	0.85 7	0.8 44	0.89 2	+37. 9%
D — Propos ed GFIF	0.926	0.911	0.92 5	0.9 18	0.95 8	+50. 0%

TABLE II
ERROR AND LATENCY METRICS

Model	FP R	FN R	Avg. Laten cy (ms)	P95 Laten cy (ms)	Through put (txn/s)
A — Rule- Based	0.18 4	0.37 2	12	18	8,200
B — Traditi onal ML	0.12 1	0.21 5	31	46	3,600
C — Deep Learning	0.08 3	0.14 3	58	89	1,850
D — Proposed GFIF	0.04 6	0.07 5	62	94	1,720

VI. CONCLUSION

To an extent, the Generative Financial Intelligence Framework integrated in this article responds to recent calls for ethics-sensitive generative AI tools in finance, directed toward safeguarding general public interest and trust, through deployments that support fraud prediction, regulatory reasoning, and wealth risk governance. SWIFT and Dun & Bradstreet datasets are exploited to illustrate, respectively, early-stage operationalization of generative financial

AMERICAN DATA SCIENCE JOURNAL FOR ADVANCED COMPUTATIONS

VOLUME: 04 ISSUE: 03

RECEIVED: JULY 07

REVISED: JULY 24

ACCEPTED: AUGUST 04

PUBLISHED: AUGUST 23



intelligence for regulatory reasoning and empirically-based trustworthiness assessment for fraud prediction.

A constructive, formal, and data-driven Gande-style empirical proof-of-concept shows that, if a statistical association is genuinely due to a cause–effect relationship, future occurrences might be forecast with an accuracy greater than random choice. Quantitative assessment builds on a GDPR-compliant, cross-category general-purpose dataset augmented with a set of properties, anchored on Gandy’s constructive approach. The connectivity-driven emulation of computational-governance decision function is methodologically motivated; supporting analyses show that the proposed conceptualization is a necessary cyber-physical infrastructure for unhackable financial-advisory AI and go beyond investigating control-board adaptive-kernel estimators.

REFERENCES

1. Ashta, A., & Herrmann, H. (2021). Artificial intelligence and fintech: An overview of opportunities and risks for banking, investments, and microfinance. *Strategic Change*, 30(3), 211–222.
2. 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.
3. 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.
4. A, R., Maheswari, N. U., Umamaheswari, P., Josephson, P. J., Mangalampalli, B. M., & Gudada, P. V. (2026). Tree-Structured Regularization Driven CNN Model for Liver Disease Diagnosis and Classification Using Ultrasound Imaging. In *2026 International Conference on Intelligent and Sustainable Electronics & Computing Technologies (INSECT)* (pp. 1–6). IEEE.
5. Bussmann, N., Giudici, P., Marinelli, D., & Papenbrock, J. (2021). Explainable machine learning in credit risk management. *Computational Economics*, 57(1), 203–216.
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*, 131, 441–452.
7. Pamisetty, A. (2026). GROWTH AND IMPACT OF INDIAN ECONOMIC DEVELOPMENT. *Ianna Journal of Interdisciplinary Studies*, ISSN (O): 2735-9891, ISSN (P): 2735-9883, 8(1), 1-10.
8. Cheng, D., Xiang, S., Shang, C., Zhang, Y., Yang, F., & Zhang, L. (2020). Spatio-temporal attention-based neural network for credit card fraud detection. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(01), 362–369.
9. Chen, D., Han, D., & others. (2023). Detecting corporate financial fraud via two-stage mapping in joint temporal and financial feature domain. *Expert Systems with Applications*, 217, 119559.
10. Jameel, M., & Mangalampalli, B. M. (2026). AI-DRIVEN DESIGN OPTIMIZATION OF SUSTAINABLE STRUCTURAL MATERIALS FOR RESILIENT INFRASTRUCTURE. *Advanced Engineering Sciences*, 58(1), 2233–2254.
11. Chen, H., & others. (2023). Large language models in finance: A survey. *Proceedings of the Fourth ACM International Conference on AI in Finance*, 374–382.
12. Dowling, M., & Lucey, B. (2023). ChatGPT for (finance) research: The Bananarama conjecture. *Finance Research Letters*, 53, 103662.
13. Gramegna, A., & Giudici, P. (2021). SHAP and LIME: An evaluation of discriminative power in credit risk. *Frontiers in Artificial Intelligence*, 4, 752558.
14. Mashetty, S. (2025). LEVERAGING DEEP LEARNING, NEURAL NETWORKS, AND DATA ENGINEERING FOR INTELLIGENT MORTGAGE LOAN VALIDATION. *INTERNATIONAL JOURNAL OF SOCIAL SCIENCE & INTERDISCIPLINARY RESEARCH* ISSN: 2277-3630 Impact factor: 8.036, 14(04), 51-65.
15. Hadji Misheva, B., & Papenbrock, J. (2022). Editorial: Explainable, trustworthy, and responsible AI for the financial service industry. *Frontiers in Artificial Intelligence*, 5, 902519.
16. 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.
17. Hilal, W., Gadsden, S. A., & Yawney, J. (2022). Financial fraud: A review of anomaly detection

AMERICAN DATA SCIENCE JOURNAL FOR ADVANCED COMPUTATIONS

VOLUME: 04 ISSUE: 03

RECEIVED: JULY 07

REVISED: JULY 24

ACCEPTED: AUGUST 04

PUBLISHED: AUGUST 23



- techniques and recent advances. *Expert Systems with Applications*, 193, 116429.
18. Inala, R., Sheelam, G. K., Aitha, A. R., Lakshmi, A. U., Nagabhyru, K. C., & Segireddy, A. R. (2026). Architecting Hybrid Data Products Using AI/ML and Agentic AI for Group Insurance and Retirement Solution Platforms with Advanced Data Governance. In 2026 IEEE International Conference on AI Engineering and Innovations (AIEI) (pp. 1–6). IEEE. 2026 IEEE International Conference on AI Engineering and Innovations (AIEI). <https://doi.org/10.1109/aiei69164.2026.11497971>
 19. Huang, A. H., Wang, H., & Yang, Y. (2023). FinBERT: A large language model for extracting information from financial text. *Contemporary Accounting Research*.
 20. Jin, J., & Zhang, Y. (2025). The analysis of fraud detection in financial market under machine learning. *Scientific Reports*, 15, 29959.
 21. Lee, M. S. A., & Floridi, L. (2021). Algorithmic fairness in mortgage lending: From absolute conditions to relational trade-offs. *Minds and Machines*, 31, 165–191.
 22. Lee, J. (2020). Access to finance for artificial intelligence regulation in the financial services industry. *European Business Organization Law Review*, 21, 731–757.
 23. Liu, Z., Huang, D., Huang, K., Li, Z., & Zhao, J. (2020). FinBERT: A pre-trained financial language representation model for financial text mining. *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, 4513–4519.
 24. Nguyen, T. T., Tahir, H., Abdelrazek, M., & Babar, M. A. (2020). Deep learning methods for credit card fraud detection. *arXiv*.
 25. Baladari, V., Nagubandi, A. R., Charan Teja Tadi, S. R. C., Selvi, A. T., Sreedevi, V., & Nithya, M. (2026). Predictive AI Model for Financial Risk Assessment in Dynamic Market Environments. In 2026 International Conference on Communication, Computing and Emerging Technologies (IC3ET) (pp. 748–753). IEEE. 2026 International Conference on Communication, Computing and Emerging Technologies (IC3ET). <https://doi.org/10.1109/ic3et64989.2026.11467452>
 26. Peña, I. P. A., & Ortega-Castro, J. C. (2024). Implementation and evaluation of an anti-fraud prototype based on generative artificial intelligence for the Ecuadorian financial sector. *Revista de Gestão Social e Ambiental*.
 27. Ratzan, J., & Rahman, N. (2024). Measuring responsible artificial intelligence (RAI) in banking: A valid and reliable instrument. *AI and Ethics*, 4, 1279–1297.
 28. Saxena, N. A., Huang, K., DeFilippis, E., Radanovic, G., Parkes, D. C., & Liu, Y. (2020). How do fairness definitions fare? Testing public attitudes towards three algorithmic definitions of fairness in loan allocations. *Artificial Intelligence*, 283, 103238.
 29. Nagabhyru, K. C., Gadi, A. L., Seenu, A., Davuluri, P. S. L. N., Segireddy, A. R., & Pamisetty, V. (2026). Towards Automated Financial Risk Scoring in Automotive Financing with Explainable Machine Learning. In 2026 IEEE International Conference on AI Engineering and Innovations (AIEI) (pp. 1–6). IEEE. 2026 IEEE International Conference on AI Engineering and Innovations (AIEI). <https://doi.org/10.1109/aiei69164.2026.11496822>
 30. Shabsigh, G., & Boukherouaa, E. B. (2023). Generative artificial intelligence in finance: Risk considerations. *IMF FinTech Notes*, 2023(006).
 31. Simion, M., & Kelp, C. (2023). Trustworthy artificial intelligence. *Asian Journal of Philosophy*, 2, Article 8.
 32. Su, H., Jiang, I.-M., & Liu, D. (2025). Detecting financial fraud risk using machine learning: Evidence based on different categories and matching samples. *Finance Research Letters*, 85, 107858.
 33. Xing, F. (2024). Financial risk tolerance profiling from text. *Information Processing & Management*, 61(4), 103704.
 34. Krishnan, M., Yandamuri, U. S., Recharla, M., Challa, K., Kolla, S. H., & Valiki, S. (2026, June). Adaptive Machine Learning Ecosystem for Personalized Medicine and Digital Finance: Leveraging Azure Cloud, Data Engineering, and Workflow Automation. In 2026 5th OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 5.0 (pp. 1–6). IEEE.
 35. 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.
 36. Zöller, M.-A., Iurshina, A., & Röder, I. (2024). Trustworthy generative AI for financial services. *OASICs—OpenAccess Series in Informatics*, 126, 2.
 37. Achakzai, M. A. K., & Juan, P. (2022). Using machine learning meta-classifiers to detect financial frauds. *Finance Research Letters*, 48, 102915.

AMERICAN DATA SCIENCE JOURNAL FOR ADVANCED COMPUTATIONS

VOLUME: 04 ISSUE: 03

RECEIVED: JULY 07

REVISED: JULY 24

ACCEPTED: AUGUST 04

PUBLISHED: AUGUST 23



38. Chen, D., Ye, J., & Ye, W. (2023). Interpretable selective learning in credit risk. *Research in International Business and Finance*, 65, 101940.
39. Yandamuri, U. S. Next-Generation Industry Platforms Integrated Systems For Predictive & Intelligent Operations. *Solaris Vantage*.
40. Dahiphale, D., Modwal, A., Kaushal, G., Balakrishnan, R., Shah, S., Agrawal, M., Lin, J., Hariramani, P., Mandawat, P., Karve, R., & Madiraju, N. (2024). Enhancing trust and safety in digital payments: An LLM-powered approach. *Google Research*.
41. Gregory, G., & Liuzzi, V. (2024). ChatGPT: A canary in the coal mine or a parrot in the echo chamber? Detecting fraud with LLM: The case of FTX. *Finance Research Letters*, 70, 106349.
42. Jin, J., & Zhang, Y. (2025). The analysis of fraud detection in financial market under machine learning. *Scientific Reports*, 15, 29959.
43. Naidu Gunasekaran, R. M. (2024). Generative AI: Opportunities, risks and implications for financial services. *The Eastasouth Journal of Information System and Computer Science*, 1(3).
44. Meda, R., Motamary, S., Dodda, A., Kummari, D. N., Kaulwar, P. K., & Gadi, A. L. (2026). Edge-AI Integration for Accelerated Decision-Making in Industrial Robotics and Manufacturing Plants. In *Lecture Notes in Networks and Systems* (pp. 332–344). Springer Nature Switzerland. https://doi.org/10.1007/978-3-032-19185-4_27
45. 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.