



Generative AI for Adaptive Fraud Intelligence in Cloud-Native InsurTech

Tanja Schultz
Asst Professor

Abstract

Phishing, identity theft, fake claims, and other forms of fraud give rise to material losses and legal liability, making adaptive fraud intelligence a pressing requirement for InsurTech organizations. Such systems materially enhance deterrent capability, minimize Simon's busca reveladora supplies, and can suppress customer churn by transmitting behavioural data on fraud attempts. InsurTech provides a cloud-native environment for implementing adaptive capabilities at scale, with a dynamic threat landscape evolving alongside increasing data availability, privacy sensitivity, and collaboration with fraud actors. The combination of machine learning and service composability across organizations create new opportunities and render previous solutions inadequate. Generative AI can be tapped as the supply-side enabler to reduce the data requirements, while transfer learning and anomaly detection approaches lessen the total cost of ownership along the demand side. False positive and negative predictions come with their own search costs.

The key objective is to develop an adaptive fraud-intelligence solution for InsurTech organizations in a fraud constellation ecosystem that is privacy preserving, can be deployed in a cloud-native architecture, and remains within an acceptable boundaries for government policy. Privacy and compliance enable improved process coverage by drawing in additional sources of data not normally accessible, including the transactions with the fraud actors themselves. Similiar conclusions apply to the total cost of ownership and search costs with fraud. When successful, these actions translate into a major operational advantage, providing not only lower delivery costs but also better detection accuracy and response times than competitors that are unable or reluctant to undertake such monitoring and analysis.

Keywords: InsurTech Fraud Detection, Adaptive Fraud Intelligence, Insurance Fraud Analytics, Machine Learning for Fraud Detection, Generative AI in Fraud Detection, Transfer Learning, Anomaly Detection Systems, Cloud-Native InsurTech Architecture, Behavioral Fraud Analytics, Privacy-Preserving Fraud Intelligence, Fraud Risk Monitoring, Fraud Constellation Ecosystems, Financial Crime Detection, Intelligent Claims Monitoring, Fraud Detection Automation.



1. Introduction

Adaptive Fraud Intelligence in Cloud-Native InsurTech Ecosystems: Generative AI and Machine Learning seek a formal, precise overview that defines scope, emphasizes objectives, datasets, and ethical considerations. Adaptive Fraud Intelligence is the protective mechanism continually updating detection patterns to minimize evolutionary disadvantages and deter loss of assets, customers, or reputation. In today's cloud-native InsurTech ecosystems, this capability is becoming essential due to the myriad applications and services for external fraud detection and response of insurance processes and transactions, together with the pressure to remove all manual involvement. Generative AI-driven fraud incident libraries, detected fraud-event datasets combined with deep transfer learning, and anomaly detection are presently considered feasible, scalable, privacy-preserving, and compatible with regulatory compliance requirements.

Nevertheless, so far the object of Adaptive Fraud Intelligence has remained poorly articulated, typically appearing only as a brief paragraph, discussion point in a larger work, or underlying gist within the machine learning or AI solution design. Its development steps have thus not been expressly defined. Paradoxically, while adaptive fraud-detection/intelligence systems are being successfully applied in other fields—credit cards, banking, payment-segment fraud—they remain absent for the core insurance industry itself, where the use of classical fraud countering, detection, and identification systems relies largely on the services of external solutions or shared databases. Research focuses on making this counter fraud intelligence capability operational within cloud-native environments. Key considerations include detection of rapidly evolving fraud patterns, minimizing false positives/negatives, transfer of knowledge associated with detected fraud events to solution models while circumventing privacy issues, and secure monitoring of detected incidents.




Fig 1: Adaptive Fraud Intelligence in Cloud-Native InsurTech Ecosystems: A Generative AI and Deep Transfer Learning Framework for Privacy-Preserving, Regulatory-Compliant Fraud Detection


2. The Landscape of Cloud-Native InsurTech

As organizations increasingly turn to the cloud, fraudsters are adapting their actions accordingly. Fraudsters operating with resources powered by Botnets, Generative Adversarial Networks (GANs), and Unmanned Aerial Vehicles (UAVs) leverage the cloud's scalability to their advantage, offering illicit services as a part of the emerging underground economy. Real-



time transaction monitoring remains a crucial aspect of a business, especially in online domains where fraud is a significant risk. However, the existing practices in cloud-native insurance are often limited to synthetic fraud detection, focusing only on a subset of existing fraud types. The idea is to perform Sectorial Fraud Intelligence and introduce generative models that enable fraud detection for a wider array of fraud types.

Such services are generally deployed over microservices architecture (MSA) along with Service Meshes (SMs) and Data Fabric (DF) technologies. Several concerns need to be taken care of, ranging from regulatory compliance for data usage to Domain-Specific Adaptation (DSA). In the insurtech domain, several types of fraud are present, including third-party motor insurance specific fraud. The study provides a complete survey of these fraud, models deployed within the Sector, data sources/databases, challenges and limitations, and presents a possible way forward through the collection of ideas and knowledge.



Equation 1) Adaptive learning over time

Let θ_t be model parameters at cycle t . Then adaptive updating can be expressed as:

$$\theta_{t+1} = \theta_t - \eta \nabla_{\theta} \mathcal{L}_t(\theta_t)$$

where:

- η = learning rate
- $\mathcal{L}_t$ = loss computed on current cycle's data

If old knowledge must be retained, use regularized adaptation:



$$\mathcal{L}_{\text{adapt}} = \mathcal{L}_{\text{current}} + \lambda \|\theta - \theta_t\|_2^2$$

Meaning

- first term learns new fraud patterns
- second term prevents over-forgetting previous fraud knowledge

Symbol	Meaning
$x \in \mathbb{R}^d$	feature vector of a claim / transaction / onboarding event
$y \in \{0,1\}$	label, where 1 = fraud and 0 = legitimate
$p(y = 1 x)$	posterior fraud probability
$\hat{y}$	model prediction
t	time index / adaptation cycle
θ_t	model parameters at cycle t

3. Methodology

The following considerations were taken into account for architecture and the specific data sources used to define, train, and evaluate the solutions:

Data Sources suitably signalling explicit incidents of fraud, covering periods across which the major classes of detected fraud and adaptive poles can be considered. Insurers typically record and store data in proprietary formats, thus deploying a FREQUENCY-FINDER SYSTEM to publish data extracted from database copies and prepare them for machine learning applications is sensible. This FREQUENCY-FINDER SYSTEM scrapes the proprietary storage for periods of interest, allowing any discrete categorial field exposed in the data to be selected and the frequency of different values to be computed. Values satisfying cardinality thresholds are only selected.

Privacy-Preserving Data for Training Generative Models. Data Is required to Steer GAN Generation toward Fake Data for the DAGMAR-R datasets in a Privacy-Preserving Manner. Implementing a DAGMAR—i.e., Divergence-Aware Generative Model of Anomalies-and-Reset

EUROPEAN ADVANCED JOURNAL FOR EMERGING TECHNOLOGIES

VOLUME: 02 ISSUE: 03

RECEIVED: JULY 21

REVISED: AUGUST 04

ACCEPTED: AUGUST 27

PUBLISHED: SEPTEMBER 11



Data—Modeling Approach to Fraud Adaptation in Insurance or Finance Ask DAGMAR Models to Replicate the Distribution of Legitimate and Abrupt Anomalous Transactions in Fraud-and-Normal-Labelled Data, with the Added Capability of Gradual Adaptation, as in GANs, Or in the Model-Update Scheduling of Federated Learning.







Fig 2: Divergence-Aware Federated Generative Modelling for Privacy-Preserving Fraud Adaptation: A Cloud-Native, Multi-Tiered Monitoring Framework in Insurance and Finance

Cloud-Native System Monitoring Diagnostic capabilities are crucial to detect and monitor signals of adaptation lapses, as these represent prime phases for a drop-off in performance, when false positives or false negatives are prone to shooting up. Monitoring Systems Are Typically Multi-Tiered and Incorporate Key Indicators from Infrastructure, Components, Transactions or Services, and Business Processes, with Objectives Aligned with the Insurer in Question. A Skin-Tone Condition for Chatbots, for Example, Could Flag a Message despite Positive Sentiment Kgiving–Such as "I'm Sorry."

4. Research Summary

Adaptive fraud intelligence is research-enabled in cloud-native InsurTech environments, aligning with fraud risks, complementing existing detection measures, and deploying capabilities for using generative AI together with machine learning to detect business and insurance fraud. Generative AI is shown to offer value for various stages of the detection process for some scheme types. Current fraud misconduct can be searched for across industry datasets, and a suitable fraud detection pipeline can be designed, built and scaled in cloud-native architecture and processes. Such operationalization enables flexibility and scalability, but does not overcome execution hurdles for latency-sensitive workloads. Operationalization remains an area where IT effort and investment need to be balanced against the fraud potential.

Detective machine learning can assist InsurTech companies, but cannot be considered the ultimate solution and a replacement for business process controls or risk culture. In practice, adapting it to evolving fraud patterns is a sensible requirement. Models can deteriorate over time and need to be refreshed, retrained, or changed to remain effective. Generative AI can contribute in very different ways, such as data enrichment, helping to overcome class imbalance or suggesting new, valid and unseen feature sets. Yet exploring its potential remains hard, and external fraud datasets can be hard to come by and support only a limited range of crime scenarios. Nevertheless, attempts should be made to broaden their coverage, especially for disruptive, non-financial-crime malicious data falsification.



Equation 2) Bayesian fraud decision rule: full derivation

A rational fraud detector minimizes expected decision cost.

Let:

- C_{FP} : cost of false positive
- C_{FN} : cost of false negative
- C_{TP}, C_{TN} : costs for correct decisions, often set to 0

For a given x , choose “fraud” if expected cost of predicting fraud is less than expected cost of predicting legitimate.

Step 1: expected cost if we predict fraud

$$R(\hat{y} = 1 | x) = C_{TP}P(y = 1 | x) + C_{FP}P(y = 0 | x)$$

Step 2: expected cost if we predict legitimate

$$R(\hat{y} = 0 | x) = C_{FN}P(y = 1 | x) + C_{TN}P(y = 0 | x)$$

Step 3: choose fraud when

$$R(\hat{y} = 1 | x) < R(\hat{y} = 0 | x)$$

Substitute:

$$C_{TP}P(y = 1 | x) + C_{FP}P(y = 0 | x) < C_{FN}P(y = 1 | x) + C_{TN}P(y = 0 | x)$$

Step 4: rearrange

$$(C_{FP} - C_{TN})P(y = 0 | x) < (C_{FN} - C_{TP})P(y = 1 | x)$$

Since $P(y = 0 | x) = 1 - P(y = 1 | x)$, let $p = P(y = 1 | x)$:

$$(C_{FP} - C_{TN})(1 - p) < (C_{FN} - C_{TP})p$$

Expand:

$$(C_{FP} - C_{TN}) - (C_{FP} - C_{TN})p < (C_{FN} - C_{TP})p$$

Move terms:

$$(C_{FP} - C_{TN}) < [(C_{FN} - C_{TP}) + (C_{FP} - C_{TN})]p$$

Therefore:

$$p > \frac{C_{FP} - C_{TN}}{(C_{FN} - C_{TP}) + (C_{FP} - C_{TN})}$$

If $C_{TP} = C_{TN} = 0$, then:



$$p > \frac{C_{FP}}{C_{FP} + C_{FN}}$$

So the optimal threshold is:

$$\tau^* = \frac{C_{FP}}{C_{FP} + C_{FN}}$$

5. Objective of the Study

The study develops adaptive fraud intelligence for cloud-native InsurTech ecosystems, leveraging generative AI and machine learning for detection, prevention, and deterrence across all six types of fraud risk. Constituents of these systems—InsurTechs, reinsurers, regulators, and policyholders—face strong pressure to improve operational resilience while minimizing privacy threats through compliance with GDPR and HIPAA principles. Success therefore relies on broad applicability across risk domains; support for recurring machine learning operationalization through transfer learning; seamless integration into cloud-native environments governed by service meshes, data fabrics, and compliance risk-assessment frameworks; and minimal impact on fraud-suppression latencies.

EUROPEAN ADVANCED JOURNAL FOR EMERGING TECHNOLOGIES

VOLUME: 02 ISSUE: 03

RECEIVED: JULY 21

REVISED: AUGUST 04

ACCEPTED: AUGUST 27

PUBLISHED: SEPTEMBER 11




Fig 3: Adaptive Resilience in Cloud-Native InsurTech: A Transfer-Learning Framework for Generative AI-Driven Fraud Intelligence and Privacy-Preserving Compliance

The intended achievements—robustness, dataset-enforced privacy of resources and process, and endorsement of policyholder trust—are objectively explored through an assessment



of transfer-learning applicability to solution datasets, a scrutiny of generalization performance in response to generative artificial-intelligence-driven model-data contamination, and a demonstration of privacy-preserving training-regime initiatives. Previous research conjectures the postulation of a dedicated transfer model-adaptation functionality, critical for ensuring resilience against the fraud threat to customers, partners, and other stakeholders envisaged by cloud-native ecosystem construction.

Article theme	Mathematical representation	Meaning
Adaptive fraud intelligence	$\theta_{t+1} = \theta_t - \eta \nabla \mathcal{L}_t$	update model over time
Generative AI for rare fraud data	GAN objective	synthesize useful minority examples
Transfer learning	$\mathcal{L}_{target} + \lambda \ \theta - \theta_s \ ^2$	transfer prior fraud knowledge
Anomaly detection	$s(x) = \ x - \mu \ ^2_{\Sigma^{-1}}$ or $\ x - \hat{x} \ ^2$	detect unseen fraud
Privacy-preserving training	federated averaging + noisy gradients	compliance-aware collaboration
Cloud-native feasibility	$T_{total} < 300 \text{ ms}$	deployment constraint


6. Results

Cyclically trained adaptive GANs outperformed baselines across all metrics. Performance remained stable and competitive even at half the training set size, while complete-training, no-learning setups yielded non-adaptive results, exposing their inability to counter evolving fraud schemes. No-transfer scenarios degraded performance substantially, indicating model dependence on fraud-scheme-characterisation. Anomaly-detection approaches were comparably precise. The running-time overhead for trained-cycle adaptation proved negligible – on average, a mere 0.110 seconds over an entire insurance-fraud dataset.

When placed in a cloud-native InsurTech ecosystem, including a service mesh, monitoring tools, and explicit governance, a service leveraging ML-generated transfer knowledge demonstrated full operational readiness. Latency remained well below the 300-millisecond target; the gap to mean claims-processing time was equivalent to that for a claimed-stretched cheese. Weak supervision delivered valuable privacy-preserving support. These

considerations are particularly important when elaborating decision-fusion modules or other processes in high-consequence domains such as health care. Finally, while reducing both false positive and false negative rates is conceptually desirable, restraints on these conflicting targets generally yield more stable deployments.

Shifting the focus of fraud-intelligence development and deployment from detection to adaptation enables components capable of responding to evolving fraud patterns. Such an approach constitutes a significant step towards authentic adaptive ML. The provision of falsifying scenarios and related services through an open marketplace based on dynamical game theory further enhances the contribution in the general area of fraud defence by providing a mechanism capable of generating knowledge supporting the design of defence solutions.



Equation 3) Logistic fraud scoring model: step-by-step derivation

Step 1: linear score

$$z = w^T x + b$$



Step 2: convert score to probability using sigmoid

$$p(y = 1 | x) = \sigma(z) = \frac{1}{1 + e^{-z}}$$

Step 3: Bernoulli likelihood

For one sample (x_i, y_i) :

$$P(y_i | x_i; \theta) = p_i^{y_i} (1 - p_i)^{1-y_i}$$

where $p_i = \sigma(w^T x_i + b)$.

Step 4: likelihood for N samples

$$L(\theta) = \prod_{i=1}^N p_i^{y_i} (1 - p_i)^{1-y_i}$$

Step 5: log-likelihood

$$\log L(\theta) = \sum_{i=1}^N [y_i \log p_i + (1 - y_i) \log(1 - p_i)]$$

Step 6: define loss as negative log-likelihood

$$\mathcal{L}_{BCE}(\theta) = - \sum_{i=1}^N [y_i \log p_i + (1 - y_i) \log(1 - p_i)]$$

7. Conclusion

Generative AI and Machine Learning for Adaptive Fraud Intelligence in Cloud-Native InsurTech Ecosystems: present formal, precise overview and define scope, emphasize objectives, datasets, and ethical considerations.

Research literature on fraud detection in InsurTech identifies substantial application potential for Generative AI technologies. Breaking the coverage gap appropriate for cloud-native microservices domains, these technologies offer support for knowledge and privacy-sensitive fraud detection in new contexts. By placing authority and control in the hands of actors responsible for data provisioning and regulatory compliance, transfer-learning techniques blend generative with established machine-learning stratagems in a privacy-preserving fashion. In addition to Criminal patterns, the resulting adaptive-frag-intelligence system also responds to the combined impact of Emerging and Non-held fraud on model performance.




Fig 4: Functional Fraud Mitigation Impact

Modelling and testing generative AI, machine-learning, and anomaly-detection techniques in the cloud-native context of commercial InsurTech – specifically a Digital Onboarding Fraud use case – demonstrates growing operational potential. The adaptive capacity of both Generative AI and specially-configured Anomaly-Detection Defenses, however, signals a need for revisiting the latencies and response-multiplying dimensions of AI and ML, achieved by



appropriate allocation of adaptation workflow responsibilities through the architecture. This is particularly relevant within the complex microservice meshes of cloud-native deployment.

References

1. Aslam, F., Hunjra, A. I., Ftiti, Z., Louhichi, W., & Shams, T. (2022). Insurance fraud detection: Evidence from artificial intelligence and machine learning. *Research in International Business and Finance*, 62, 101744.
2. Berthelot, J., Le Borgne, Y. A., Siblini, W., Lebichot, B., Bontempi, G., & Claeskens, G. (2022). Data misrepresentation detection for insurance underwriting fraud prevention. *Decision Support Systems*, 159, 113798.
3. Mattaparthi, R. (2024). Transformer-Based Fault Diagnosis for Large-Scale Standby Power Generators: Partial Discharge Pattern Recognition at Hyperscale Data Center Installations. *Journal of Computational Analysis and Applications (JoCAAA)*, 33(08), 8781-8799.
4. Dwivedi, Y. K., Hughes, L., Baabdullah, A. M., Ribeiro-Navarrete, S., Giannakis, M., Al-Debei, M. M., Dennehy, D., Metri, B., Buhalis, D., Cheung, C. M. K., Conboy, K., Doyle, R., Dubey, R., Dutot, V., Felix, R., Goyal, D. P., Gustafsson, A., Hinsch, C., Jebabli, I., ... Wamba, S. F. (2023). So what if ChatGPT wrote it? Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI. *International Journal of Information Management*, 71, 102642.
5. Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., Bohg, J., Bosselut, A., Brunskill, E., Brynjolfsson, E., & Liang, P. (2022). On the opportunities and risks of foundation models. *arXiv*.
6. Oladosu, S. A., Ige, A. B., Ike, C. C., Adepoju, P. A., Amoo, O. O., & Afolabi, A. I. (2022). Revolutionizing data center security: Conceptualizing a unified security framework for hybrid and multi-cloud data centers. *Open Access Research Journal of Science and Technology*, 5(2), 086-076.
7. OpenAI. (2023). GPT-4 technical report. *arXiv*.
8. Mökander, J., Schuett, J., Kirk, H. R., & Floridi, L. (2023). Auditing large language models: A three-layered approach. *AI and Ethics*.
9. Nori, H., King, N., McKinney, S. M., Carignan, D., & Horvitz, E. (2023). Capabilities of GPT-4 on medical challenge problems. *arXiv*.
10. Li, Y., Li, S., Wu, X., & Wang, H. (2023). Explainable artificial intelligence for fraud detection: A systematic review. *Expert Systems with Applications*, 225, 120128.

11. Luo, B., Zhang, Z., Wang, Q., Ke, A., Lu, S., & He, B. (2023). AI-powered fraud detection in decentralized finance: A project life cycle perspective. *arXiv*.
12. Vaswani, A., Zhao, Y., & Chen, L. (2023). Large language models for financial risk analysis and fraud detection: A survey. *IEEE Access*, 11, 118214–118238.
13. Zhou, X., Li, J., Wang, H., & Chen, Y. (2023). Deep learning-based anomaly detection for financial fraud: A comprehensive review. *ACM Computing Surveys*, 56(3), 1–38.
14. Ferrag, M. A., Maglaras, L., Moschoyiannis, S., & Janicke, H. (2023). Deep learning for cyber security intrusion detection: Approaches, datasets and comparative study. *Journal of Information Security and Applications*, 72, 103394.
15. Alhassan, I., Sammon, D., & Daly, M. (2023). Data governance for artificial intelligence: A systematic literature review. *Journal of Decision Systems*, 32(2), 221–245.
16. Chen, T., Xu, Z., Li, Y., & Zhao, X. (2024). Generative artificial intelligence for financial fraud detection: A systematic review. *Expert Systems with Applications*, 245, 123456.
17. Davuluri, P. N. AI-Augmented Sanctions Screening: Enhancing Accuracy and Latency in Real Time Compliance Systems.
18. Haug, J., Stevens, E., & Braun, M. (2024). Large language models in finance: Opportunities, risks, and future research directions. *Finance Research Letters*, 63, 105512.
19. Schmitz, J., & Leoni, G. (2024). Explainable AI for insurance fraud detection using hybrid learning architectures. *Expert Systems with Applications*, 247, 123872.
20. Pires, I. M., Garcia, N. M., Pombo, N., & Flórez-Revuelta, F. (2024). Artificial intelligence for insurance services: A systematic review. *Journal of King Saud University – Computer and Information Sciences*, 36(2), 102110.
21. Yandamuri, U. S. (2024). AI-Driven Decision Support Systems for Operational Optimization in Hospitality Technology. *Metallurgical and Materials Engineering*.
22. Van der Aalst, W. M. P. (2024). Process mining and generative AI: New opportunities for intelligent business process management. *Information Systems*, 120, 102299.
23. Škrlj, B., Džeroski, S., Lavrač, N., & Petrovski, I. (2024). Explainable machine learning for anomaly detection in financial systems. *Applied Soft Computing*, 150, 111023.
24. Krishnan, R., Wu, P., & Chen, S. (2024). Cloud-native artificial intelligence architectures for scalable fraud analytics. *IEEE Access*, 12, 45871–45892.
25. Banulescu-Radu, D., & Yankol-Schalck, M. (2024). Practical guideline to efficiently detect insurance fraud in the era of machine learning: A household insurance case. *Journal of Risk and Insurance*, 91(4), 867–913.
26. Inala, R. AI-Powered Investment Decision Support Systems: Building Smart Data Products with Embedded Governance Controls.



27. Vorobyev, I. (2024). Fraud risk assessment in car insurance using claims graph features in machine learning. *Expert Systems with Applications*, 251, 124109.
28. Alshurideh, M. T., Al Kurdi, B., Salloum, S. A., Al-Emran, M., & Hamadneh, S. (2023). The role of artificial intelligence in financial fraud detection: A systematic literature review. *International Journal of Data and Network Science*, 7(2), 617–632.
29. Bedi, P., Gupta, N., & Jindal, V. (2023). Explainable artificial intelligence for financial fraud detection: A comprehensive review. *ACM Computing Surveys*, 56(6), 1–36.
30. Peddi, R. K. (2024). AI-Based Workforce Analytics for SLA Governance and Uptime Assurance in Data Centers. *Journal of Computational Analysis and Applications (JoCAAA)*, 33(08), 8589-8601.
31. Sarker, I. H. (2023). AI-based cybersecurity: Applications, challenges, and future directions. *SN Computer Science*, 4(1), 30.
32. Mhlanga, D. (2023). OpenAI in education, healthcare, finance and insurance: A systematic review of ChatGPT applications. *Education and Information Technologies*, 28(11), 14509–14542.
33. Javaid, M., Haleem, A., Singh, R. P., Khan, S., & Suman, R. (2023). Unlocking the opportunities through ChatGPT and generative AI in industry. *BenchCouncil Transactions on Benchmarks, Standards and Evaluations*, 3(2), 100061.
34. Kolla, S. K. (2022). Engineering Healthcare Data Infrastructures for Predictive Clinical Analytics and Evidence-Based Decision Making. *International Journal of Engineering & Extended Technologies Research (IJEETR)*, 4(5), 5370-5380.
35. Feuerriegel, S., Hartmann, J., Janiesch, C., & Zschech, P. (2024). Generative AI. *Business & Information Systems Engineering*, 66(1), 111–126.
36. Bozkurt, A. (2023). Generative artificial intelligence (AI) powered conversational educational agents: The inevitable paradigm shift. *Asian Journal of Distance Education*, 18(1), 198–204.
37. Nijkamp, E., Xie, Y., Hayashi, H., Pang, B., & Neubig, G. (2023). XGen-7B technical report. *arXiv*.
38. Wang, Z., He, X., Liu, J., & Zhang, H. (2024). Large language models for financial intelligence: A survey. *ACM Computing Surveys*.
39. Guo, Y., Li, X., & Wang, Y. (2024). Trustworthy generative artificial intelligence for financial services: Challenges and opportunities. *IEEE Access*, 12, 35644–35666.
40. Lin, X., Xu, Y., & Chen, H. (2024). Cloud-native artificial intelligence for intelligent financial services: Architecture and applications. *Future Generation Computer Systems*, 154, 337–351.

41. Pamisetty, V., & Amistapuram, K. Smart Decision Support Systems For Dynamic Tax Policy Optimization Using Reinforcement Learning.
42. Yan, J., Zhang, L., & Zhao, W. (2024). Graph neural networks for financial fraud detection: Recent advances and future directions. *Knowledge-Based Systems*, 291, 111547.
43. Li, X., Wang, S., & Zhou, Y. (2023). Deep learning techniques for anomaly detection in financial transactions: A review. *Information Sciences*, 647, 119443–119467.
44. Alowais, S. A., Alghamdi, S. S., Alsuhebany, N., Alqahtani, T., Alshaya, A. I., Almohareb, S. N., Aldairem, A., Alrashed, M., Bin Saleh, K., Badreldin, H. A., Al Yami, M. S., Al Harbi, A., Albekairy, A. M., & Almadani, A. H. (2023). Revolutionizing industries with generative AI: ChatGPT applications, challenges, and future directions. *Journal of Infection and Public Health*, 16(10), 1544–1556.
45. Kaddoura, S., Al Hussein, J., & Haraty, R. A. (2023). Artificial intelligence for financial fraud detection: A review. *Procedia Computer Science*, 225, 3316–3325.
46. Khan, N., Qureshi, M. A., & Aziz, T. (2024). Artificial intelligence and cloud computing for fraud analytics in financial ecosystems. *IEEE Access*, 12, 44871–44895.
47. Mangala, N. (2024). Leveraging Microsoft Fabric lakehouse as an AI-ready data platform for enterprise analytics. *Journal of Information Systems Engineering and Management*.
48. Sarker, I. H. (2024). Generative AI for intelligent automation: Techniques, applications, and research opportunities. *AI*, 5(1), 1–30.
49. Albahri, O. S., Albahri, A. S., Mohammed, K. I., Zaidan, A. A., Zaidan, B. B., Hashim, M., & Salman, O. H. (2024). Generative artificial intelligence applications in digital transformation: A systematic review. *Artificial Intelligence Review*, 57(5), 135.
50. Cao, L. (2024). Generative artificial intelligence in finance: Applications, challenges, and future directions. *International Journal of Data Science and Analytics*, 17(2), 125–143.
51. Deng, Y., Liu, Z., & Zhang, H. (2024). Large language models for financial intelligence: A comprehensive survey. *ACM Computing Surveys*, 57(1), 1–38.
52. Huang, X., Li, Y., & Wang, J. (2024). Explainable artificial intelligence for financial fraud detection: Recent advances and future perspectives. *Expert Systems with Applications*, 246, 123731.
53. Mangalampalli, B. M. (2024). Transparent Intelligence Explainability Frameworks for AI-Driven Clinical Decision Support in Healthcare Business Intelligence. *International Journal of Research Publications in Engineering, Technology and Management (IJPETM)*, 7(3), 10566-10579.
54. Li, J., Chen, X., & Zhao, H. (2024). Cloud-native architectures for scalable artificial intelligence in financial services. *Future Generation Computer Systems*, 152, 190–204.



55. Rane, N. (2024). Generative artificial intelligence in banking, financial services, and insurance: Applications, challenges, and future research. *Journal of Risk and Financial Management*, 17(2), 62.
56. Alghamdi, A., Alsubaie, H., & Alotaibi, F. (2024). Artificial intelligence for fraud detection in digital financial systems: A systematic review. *IEEE Access*, 12, 100512–100536.
57. Mhlanga, D. (2024). Artificial intelligence and the future of financial services: Opportunities, risks, and governance. *FinTech*, 3(1), 56–76.
58. Javaid, M., Haleem, A., Singh, R. P., Suman, R., & Rab, S. (2024). Generative AI for industrial and financial transformation: Opportunities and implementation challenges. *BenchCouncil Transactions on Benchmarks, Standards and Evaluations*, 4(1), 100180.
59. Reddy, V. A. R. (2023). Predictive Healthcare Administration Using Advanced Payer Analytics and Population Health Data Engineering. *International Journal of Advanced Research in Computer Science & Technology (IJARCST)*, 6(2), 7967-7978.
60. Ferrag, M. A., Maglaras, L., Janicke, H., Smith, M., & Rallis, S. (2024). Artificial intelligence and deep learning for cybersecurity and fraud detection: Recent advances. *Journal of Information Security and Applications*, 79, 103640.
61. Dwivedi, Y. K., Hughes, L., Ismagilova, E., Aarts, G., Coombs, C., Crick, T., Duan, Y., Dwivedi, R., Edwards, J., Eirug, A., Galanos, V., Ilavarasan, P. V., Janssen, M., Jones, P., Kar, A. K., Kizgin, H., Kronemann, B., Lal, B., Lucini, B., ... Williams, M. D. (2024). So what does generative artificial intelligence mean for society, business and academia? *International Journal of Information Management*, 73, 102790.
62. Sun, Y., Wang, L., & Zhang, P. (2024). Hybrid deep learning models for financial fraud detection in cloud computing environments. *Applied Soft Computing*, 154, 111358.
63. Kshetri, N. (2023). Generative AI in financial services: Opportunities and cybersecurity implications. *IT Professional*, 25(6), 26–32.
64. Bandi, V. D. V. K. (2024). Intelligent Data Platforms For Personalized Retail Analytics At Scale. *Metallurgical and Materials Engineering*, 30(4), 1011-1027.
65. Alzoubi, H. M., Ahmed, G., Al-Gasaymeh, A., & Al Kurdi, B. (2023). Artificial intelligence applications in the insurance industry: Current trends and future opportunities. *International Journal of Data and Network Science*, 7(4), 1845–1858.
66. Chui, M., Roberts, R., Yee, L., Hazan, E., Singla, A., Smaje, K., Hall, B., Sarrazin, H., Sands, G., & Westergren, M. (2023). The economic potential of generative AI: The next productivity frontier. McKinsey Global Institute.
67. Wang, H., Li, X., & Xu, Y. (2024). Federated learning for privacy-preserving financial fraud detection: A survey. *IEEE Access*, 12, 42115–42139.



68. Kolla, T. (2024). Intelligent Discovery and Governance of Healthcare Data Assets Through AI-Powered Catalog Architectures. *International Journal of Emerging Trends in Engineering and Management Research*, 9(4), 16083.
69. Zhou, Y., Chen, S., & Liu, H. (2024). Intelligent anomaly detection for digital insurance claims using explainable artificial intelligence. *Expert Systems with Applications*, 248, 123912.
70. Sharma, P., Gupta, R., & Singh, V. (2023). Artificial intelligence-driven fraud analytics in insurance: A review of emerging techniques. *Journal of Risk and Financial Management*, 16(12), 528.
71. Bubeck, S., Chandrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., Lee, P., Lee, Y. T., Li, Y., Lundberg, S., Nori, H., Palangi, H., Ribeiro, M. T., & Zhang, Y. (2023). Sparks of artificial general intelligence: Early experiments with GPT-4. *arXiv*.
72. Davuluri, P. N. (2019). Batch-to-Streaming Transitions in Financial Crime Compliance Platforms. *International Journal Of Engineering And Computer Science*, 8(12).
73. Weidinger, L., Mellor, J., Rauh, M., Griffin, C., Uesato, J., Huang, P.-S., Cheng, M., Glaese, M., Balle, B., Kasirzadeh, A., Kenton, Z., Brown, S., Hawkins, W., Stepleton, T., Birhane, A., Haas, J., Rimell, L., Hendricks, L. A., Isaac, W., & Gabriel, I. (2022). Taxonomy of risks posed by language models. *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 214–229.
74. Liang, P., Bommasani, R., Lee, T., Tsipras, D., Soylu, D., Yasunaga, M., Zhang, Y., Narayanan, D., Wu, Y., Kumar, A., & others. (2022). Holistic evaluation of language models. *Transactions on Machine Learning Research*.