Journal of International Commercial Law and Technology
Print ISSN: 1901-8401

Website: https://www.jiclt.com/

Article

Generative AI for Exceptions and Investigations: Streamlining Resolution Across Global Payment Systems

Article History:

Name of Author:
Vijaya Rama Raju Gottimukkala[1], Dr. Ketaki Kulkarni[2]

Affiliation:
[1]Senior Dev/ Systems Engineer
[2]Associate Professor, Department of Civil Engineering Dr. Vishwanath Karad MIT World Peace University Pune, India

Corresponding Author:
Dr. Ketaki Kulkarni

How to cite this article: Gottimukkala , *et al.* Generative AI for Exceptions and Investigations: Streamlining Resolution Across Global Payment Systems. *J Int Commer Law Technol*. 2025;6(1):969–972

Received: *03-10-2025*
Revised: *17-10-2025*
Accepted: *02-11-2025*
Published: *19-11-2025*

Abstract: Global payment networks process trillions of dollars—soon to be increased by instant payment sys- tems—constantly generating exceptions and investigative ques- tions. Stakeholders often spend hours or weeks responding, with surges causing costly backlogs and impacting customer trust. Automating these tasks—for instance, triaging, investigating, analyzing root causes, detecting anomalies, preparing forensics, and generating reports—through generative AI (GenAI) could reduce the volume and time taken while enhancing quality and yield. Business use cases leverage these capabilities to address specific pain points, which must be prioritized based on data flows, governance, and modeling risk. The potential level of automation varies. Identifying these candidates requires documenting difference flows, sources, and transit times for each transaction type. Regulation and compliance must be considered when analyzing processes, especially with potential changes, such as the introduction of a central bank digital currency, which could increase transactional provenance. Once those aspects are governed, GenAI can be scaled.

Keywords: Generative Artificial Intelligence (GenAI),Payment Exception Management,Automated Investigations,Intelligent Payment Resolution,Global Payment Networks,AI-Driven Process Automation,Financial Fraud Detection,Natural Language Processing in Banking,Machine Learning for Payment Reconciliation,Cognitive Automation in Fintech.

INTRODUCTION

The global payments ecosystem is incredibly complex, driven by multiple networks with diverse rules and governance. Every transaction generates extensive metadata that is stored, often for very long periods of time. Despite all of the data generated, however, there is still significant latency in the global payments system. Many transactions are rejected or consumed by fraud prevention processes, resulting in false positive identifications, as well as investigation processes that may uncover unusual patterns but are not scalable to monitor the entire payment stream. These problems are particularly acute in investigations, where the volume of tickets, both false positives and true investigations, far outstrips the ability to analyze them, and for which the investigation journey is often reconstructed after the investigation rather than created in real time. By analyzing the information generated throughout the payments system, it should be possible to define a Generative.

Fig. 1. Generative AI in Digital Payments

AI solution that assists in automatically resolving exceptions, and investigating post-facto trackers at scale so that it becomes a net source of business value rather than a cost center. The goal would be to reduce the number of exceptions arising in the network by 20%, reduce ticket closure time by 20%, and increase the number of documented investigative insights for every completed investigation by 50%.

Problem Space and Objectives

Two powerful but distinct capabilities can be identified to greatly benefit many if not all of the firms in payments: one that automates the remedy of exceptions, and another that syn- thesizes insights from data at rest to help investigators perform root cause analysis. These persisting and regimented classes of generative AI applications, drawing from data in motion and at rest, can reduce current-focused operational costs and risks, and transform the use of external data into future-focused predictive and prescriptive sources of insight. On the current-focused operations front, the typical studio–theatre metaphor is inverted: live actors on a stage are supported by production planning behind the scenes; in exceptions the majority of the work is in the wings or at the control desk, with a few human analysts to react to the unexpected on stage. Humans, hoping for the best, try to learn how to spot the rarest events where the appropriate automations have yet to be scaled up in time. That balance will invert when experimental predictive projects succeed: predictions and preventive action will become the norm, relying on external data and factors trending – and live incident investigation and remedial refinement will become the rarer focus. The desired outcome is a scalable solution to reduce the normal operational load by at least 20%, decreasing both the modalities required for control checks and the average time taken to resolve the remaining checks.

Key Stakeholders and Use Cases

Payments network operators, owners, and their ecosystem partners benefit from enhancing transaction efficiency and risk management. AI technology improves exception analysis, remediation prioritization, workflow routing, and underlying investigation. For traffic integrity, it is critical that exceptions due to fraud or regulatory failures—especially AML, KYC, and Sanction violations—occur rarely. Emerging payment types and channels pose additional risks that help solve and keep emerging. The associated data labels, hence, grow in quality, volume, and diversity. To mitigate technology drift, the volume is increasingly examined through non-traditional approaches: labelled or unlabelled, supervised or unsuper- vised, deep, or generative learning. Data observability and monitoring measures of quality and integrity underpin these options. The low probability of detection or the time taken to identify fraudsters remains a primary area of investment con- cern since the associated data attributes can accurately identify them. Exception remediation, however, allocates resources to mitigate their impact. Generative AI helps synthesize data in the absence of historical examples or limited supervision for traffic integrity instead of relying on burdened physical resources after the event. The hypothesis generation, feature- generation, and causality-exploration processes can now be automated.

II. FOUNDATIONS OF GENERATIVE AI IN PAYMENTS

Global payment networks involve a multitude of participants exchanging data for funds transfers. These participants include senders, receivers, intermediary institutions that facilitate routing, and various service providers. Payment networks offer a wealth of data, which is provided across these multiple stakeholders in various shapes, forms, timeliness, and transformation states. As generative AI is applied to this data, it is crucial to architect an end-to-end picture of the required data, its nature and associated considerations. Once the relevant data flows are agreed upon, additional components required to deploy generative AI at scale can be determined. At this stage in the discussion it is also appropriate to bring in risk and compliance considerations, articulating the various known regulatory requirements across payment networks, the need for privacy and other protection measures, auditability for model outputs, and other aspects around security and model governance. A critical component of payment networks is that of risk management and compliance. The operationalization of generative AI, for exception resolution and investigative purposes, cannot be done without express parameters governing these areas. In

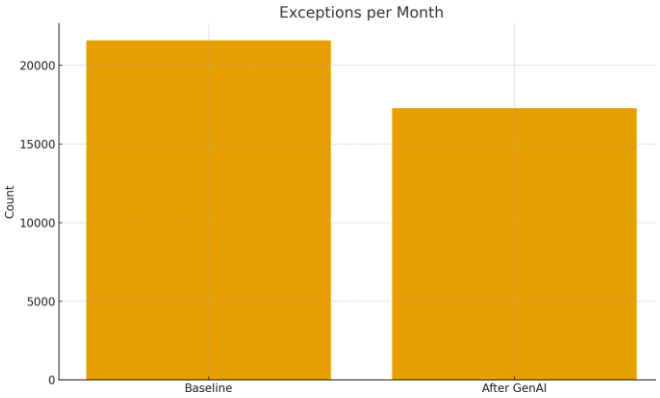

Fig. 2. Exceptions per Month

particular, the compliance requirements around PCI DSS and PII are paramount, followed by MLRO and Data Protection Policy considerations. Following this, wider SOC2-type rules can be introduced to cover other areas of risk auditability and observability RGMS. AI is technology that falls within the remit of a group of users, who execute a set of actions according to a set of permissions and perform these actions on various systems and with connections to various data sources. As such, a complete set of controls around the AI user, and the data used, is crucial for strict control of the outcomes and observability of in-use detection techniques.

Equation 01: Exception volume model (Poisson ⇒ 20% reduction)

Paper target: "reduce the number of exceptions by 20So the post-GenAI arrival rate is

$$\lambda' = 0.8\lambda Implications \quad (1)$$

$$E[N(t)] = \lambda t \rightarrow E[N'(t)] = 0.8\lambda t \quad (2)$$

$$Var[N(t)] = \lambda t \rightarrow Var[N'(t)] = 0.8\lambda t \quad (3)$$

Scenario	Arrival rate λ (per hr)	Service rate μ (per hr)
Baseline	30	40
After GenAI	24	36.5

TABLE I
QUEUE MODEL SUMMARY

A. Data Flows in Global Networks

A detailed end-to-end overview of data flows across global payment networks informs the broad set of stakeholders who depend on data at various stages of the investigation cycle. Knowledge of where, when, and how various data sources are ingested within the system, and assurance that they undergo appropriate transformation and data quality checks,
will strengthen organizational confidence in generative AI systems. This, in turn, will promote wider adoption of automated exception-resolution workflows. For a generative AI system to deliver effective resolutions and insights, data quality is paramount. The system cannot generate coherent natural-language results if the input data are misaligned or incorrect. To proactively address these data-quality concerns, it is essential to understand where, when, and how the data are being assembled, combined, and queried. An end-to-end data flow diagram covering several of the main processes and technologies involved, including exception detection, ticket creation, and insights generation, provides a common reference. The various shapes indicate the function of the underlying data-dot sources (circle), processes (diamond), and data-collection considerations (parallelogeogram) – such as the expected timing of data arrival (temporal stamp in the upper right corner) and appropriate data lineage.

Risk and Compliance Considerations

A robust risk and compliance framework is essential for deploying Generative AI across global payments networks. Regulatory requirements for financial services firms in countries such as the USA and UK place emphasis on privacy, model governance, data quality, auditability, monitoring, and controls. Additionally, firms must ensure that models and processes can be easily and thoroughly inspected by regulators and auditors. Data Exploration and generation of detailed reports that identify potential risk exposures or regulatory breaches require access to extensive data. Privacy Risk Assessment and Data Sharing Policies address data risk and data governance management by specifying privacy reviews, approval workflows for new data sources, sharing processes, anonymization methods, and data quality checking policies. These elements facilitate Enterprise Data Catalog/Lineage capabilities to provide data flow context for underlying data risk assessments throughout a generative model's lifecycle including model training and evaluation, investigation, and the generation of reports detailing conditions leading to breaches of critical business controls.

III. AUTOMATING EXCEPTION RESOLUTION

A taxonomy of exceptions and incident tickets is essential to automate the response, control risk, and enable human intervention for the most urgent or complex cases. The taxonomy captures key facets: types and severity of exceptions; automated remediation techniques; and a state machine to govern how incidents move from one state to the next. The taxonomy can begin with the common Payment Network Exception Classification (PNEC) schema, supplemented by the PMR, and incorporate other classifications to expand coverage. From there it can evolve to suit the organization's needs. A ticketing system is fundamental to automate detection, control resolution processes, and provide a clear audit trail. An orchestration layer governs the routing and state transitions of incident tickets using a set of rules. Unlike BPM tools focused on

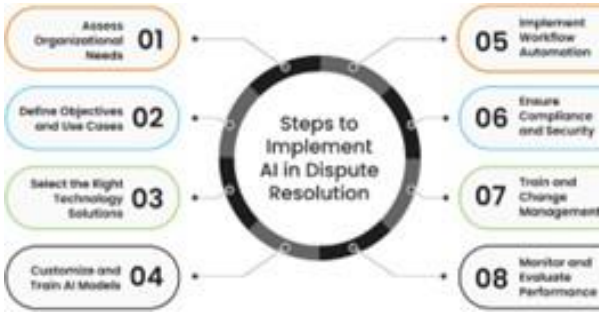

Fig. 3. Automating Exception Resolution

process design and visibility, this layer automates incident resolution at scale. Automated resolution reduces latency by eliminating queuing within work centers. At the same time, orchestration reduces latency by expediting resolution for incidents that can be addressed automatically, even when delays exist for those that cannot.

Exception Taxonomy and Prioritization

Consistent exception occurrences signal systemic issues, warranting dedicated analysis. Conversely, isolated instances, even if significant, demand limited investigative resources. A well-defined exception type taxonomy streamlines this prioritization process. Classification parameters typically encompass exception type, severity, expected frequency, recurrence monitoring, and an urgency matrix that facilitates automated resolution capability mapping or

technology-assisted triage. Category-wise SLA determination, decision automation, and further advice can enhance routing efficiency. Severity categories range from critical—systemic risk necessitating immediate resolution or ban—through high and medium levels, indicative of poor user experience or reputation impact. Diminished and low severity assess less business-sensitive third-party network flows, avoiding secondary resolution burden when associated with more critical counterparts. Level recurrences are also scrutinized. Documented SLAs integrate with ticketing engines for observability, enabling statistical tracking against defined thresholds. Two-dimensional escalation matrices—time-to-resolution and service-class-based—inform decisioning and routing rules, ensuring clear escalation pathways. A set of remediation and technology systems capable of automating resolution within defined SLAs links to each state, enabling end-to-end-ticketing integration.

Ticketing and Workflow Orchestration

An adaptive state machine manages exception ticket lifecycles. The design specifies states and transitions, routing logic based on ticket characteristics, AND agreements and escalation paths for priority breaches. Triggers define

human-in-the-loop intervention points. A ticketing and workflow orchestration engine integrates with the exception taxonomy to apply an adaptive state machine to exception tickets. A ticket's current stage in its lifecycle dictates operational behavior, including resolution actions, routing, and monitoring. The state machine supports the standard states associated with an exception ticket—draft, in-progress, escalated, resolved, closed, and deleted—segmented into four collaborative stages. The workflow design specifies the states, transitions, tickets, and human checks required to orientate and reorientate progress. First, tickets are held in a draft state until a resource is available to investigate and remediate action. Subsequently, they are moved to an in-progress state, indicating that candidates have begun investigating root cause and necessary remediation activity. If the exception has not been resolved within the SLA period established for that priority, the ticket is moved to an escalated state, enabling visibility and prompting action. When the candidate(s) finish their work, the ticket is moved to resolved. This action should also specify an explanation. These operations can be performed independently or in parallel by multiple candidates. Following confirmation that the remediation activity has worked for all parties in the workflow, tickets are then finally closed. Alternatively, tickets can be deleted

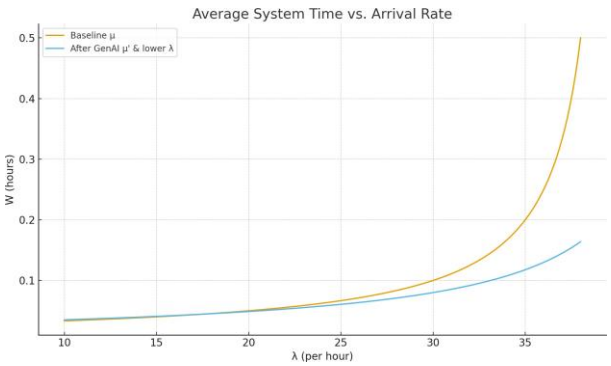

Fig. 4. Average System Time Vs. Arrival Rate

Metric	Baseline	After GenAI
Exceptions per month	21600	17280
Avg closure time (hr)	0.1	0.08
Insights per case	2	3
Data Quality Score (0..1)	0.8	0.849
Anomaly alerts (n/500)	21	10

TABLE II KPI SUMMARY

INVESTIGATIVE ANALYTICS AND INSIGHTS

Four classes of innovative analytic methods automate the generation of investigative insights and support root cause analysis of resolution exceptions. The first set assists in establishing the reasons an exception was raised by generating and validating hypotheses. The second leverages a forensics framework to detect anomalies that fall outside of known operational bounds across specific routes and service groups. The third retrieves relevant historical attribute.

$$\%TODO : Rewritethefollowingmathusingvalid \tag{4}$$

$$LaTeXsyntax\%wi\ddot{}\geq 0, = 1wi=\ddot{}1 \tag{5}$$

$$w_i q_i youcanpropagateDQSasagateforscheduledretrains \tag{6}$$

A. Root Cause Analysis Methods

Automating the customized, detailed investigations needed to assess the root causes of exceptions is a critical aspect of simplifying and enhancing the payment exception lifecycle. Several approaches can be drawn upon to generate probable causative factors for a given payment's failure, or other previously-defined data markers. Key factors commonly evaluated during incident investigations, often through human analytical investigation, include payment field contents, participant operational zones and schedules, and error and fraud indicators. Generative AI is well positioned to exploit vast data of varying structures and the many detection and monitoring mechanisms built into the network. Generative AI is well positioned to automate the evaluation of these widely varying sources and to provide recommended features to investigate based on the patterns it finds in previous incidents. Training data for these tasks ideally consists of exception records that include a failure indicator, plus associated payments and corresponding data flows that led to the incident. Suggested features could be generated via

techniques such as unstructured-text-search insertion into the prompts with config settings, suspicion or warning levels, and tagging to indicate those nodal-area domains where data appears to be in-consistency or in-violation. Once initial recommendations have been made and other exploration on a failure-indicator group performed, causality considerations could then lead to further recommendations that relate to data in-absence. After all investigation leads have been covered, the verification of the proposed causes should follow standard hypothesis-testing and validation techniques.

B. Anomaly Detection and Forensics

A comprehensive framework for elevating exception resolution and enhancing investigative insights must also encompass anomaly detection and forensic analysis. The goal is to systematically identify irregularities across the payments ecosystem, drill down into the underlying causal factors, and construct profiles that enable organizations to identify similar strings or cohorts for further examination. Association analyses to highlight correlations between different trace data can yield additional analytical lenses for investigators. A broad range of techniques, spanning outlier detection for numeric feature spaces to clustering algorithms for categorical variables, must be considered. Analyses of transaction status traces offer stand-alone anomaly detection capabilities. Supporting traces of sender and receiver data and relationships enable automated detection of key-man scenarios and connections within sanctioned corporate networks. Enrichment factors include Relative Sentiment Profiles (RSP, capturing asymmetries in commentary sentiment) and theme modulation features (changing relative frequencies of phrases) based on Bridging Language Models.

These considerations reveal other operators for propagation, enabling active monitoring of per-operator detection rates or informing the design of

incentive programmes. Above all, regular feedback loops with operational teams draw on their domain knowledge to surface and report collectable findings that genuinely matter for ongoing investigations.

V. OPERATIONALIZING AT SCALE

Generative AI implementations require sufficient maturity to operate reliably and deliver value at scale. For exception resolution, critical areas include data quality, risk and compliance, data governance, privacy, security, and operational reliability. High-quality, current data is vital for effective analysis. Each data source should be assigned a quality grade and portion of the overall data quality score, which then flows to the machine-learning model. An analytics platform, such as Azure Purview, can collect these statistics, make them visible to stakeholders, and initiate responsiveness enhancements with data owners. If a model meets data quality criteria prior to predicted retraining, it can be updated automatically. Considerable investment is needed to ensure that live data streams are free of holes, delays, or other issues that create information gaps for exception resolution or investigation. For example, exception volume

and activity reporting should yield zero unexpected downtime or creates. Attention to risk and compliance from the outset offers a meaningful competitive advantage, particularly in light of public scrutiny of data-sharing arrangements. The AI application should begin reconnaissance around regulatory requirements, business-unit-risk areas, and group-level policies. Other aspects warranting early engagement include privacy aspects such as transparent and anonymized datasets, clarity over model governance—who is accountable for
sign-off of new models, retraining of existing models, and results—and the application of a p¨rivacy by designp¨rinciples during development.

A. Data Quality, Privacy, and Security

Governing data quality, privacy, and security assures integrity and reputation while prioritizing subscriptions across the
end-to-end flow. Cleanliness, consistency, and completeness of data inputs are vital for reliable AI outcomes, whether driving decision-making or generating insights. Data governance frameworks mandate periodic reviews and remediations, guiding required transformations to accommodate model needs and ensuring delivery of completed and accurate datasets for GAI. Routing systems detect deviations and trigger alerts for correction, auditing, or explanation, while robust measures secured against external and internal threats shield against malicious deployments. Sensitive information undergoes anonymization or removal in alignment with regulatory, institutional, and organizational stipulations; automated checks verify compliance. Data-sharing agreements are established for consumption across jurisdictions, while ML models certified by risk and compliance teams enable external data use.

Detailed logs comply with regulatory retention mandates, monitor all AI operations, record secure access, and remain accessible for investigation retrospectives and support. Data protection breaches trigger real-time alerts, retraining occurs at defined intervals, and thematic review processes execute on an as-needed basis.

B. System Reliability and Observability

Data-driven mechanisms for exception resolution and investigative insights by their nature rely on actioning potentially large datasets. Engaging sensitive data requires appropriate governance that anticipates all use cases and ensures that privacy, confidentiality, and security measures can be applied effectively. Independent of how automated they may be, these data flows remain critical to the integrity of the payment network, thus necessitating a governance framework that appropriately weighs risk against potential benefit. Consistent and logical application of rules and controls, such as adequate access authorisation, online credit checks for customers and transaction participants, and data retention periods, is essential to minimising unwanted use.

Even with appropriate data governance, exception resolution and investigation systems require careful monitoring and operational support to ensure continued integrity and

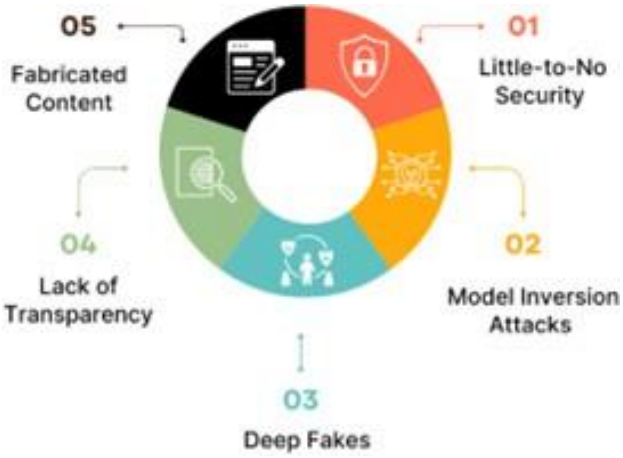

Fig. 5. Data Privacy Important for Organization

reliability. Tracking relevant metrics helps identify emerging issues. Alerts proactively highlight problems, minimising the need for time-consuming detective work and enabling the application of appropriate remedial action. Model retraining should occur on a regular basis, with further training triggered by signs of performance degradation; clear rollback procedures for deployed models make such interventions both rapid and painless.

CONCLUSION

Generative AI can help networks intelligently

synthesize responses to exceptions and investigate payment and credit decisions. The research deploys generative AI to automate two capabilities. First, generative AI representations of an exception taxonomy help with the automated remediation of exceptions based on their type and severity. Generative AI analyzes the formal specification of exceptions and augments these descriptions with patterns and other information for trained models to remediate tickets at scale. Second, analytical methods leverage generative AI to create advanced capabilities that identify the root causes of exceptions, alert investigators to ongoing anomalies and enable the forensic analysis of serious issues. Generative AI combines generative models with other forms of analytical workbench and data – such as causal models, the mathematical properties of features and high-variance data patterns – to create definitions and guides that improve the quality and scale of payments investigation. Together, these elements reduce the number of exceptions, increase the speed at which they are resolved and enhance the generative quality and quality of outputs achieved. Future directions in this application area will be exciting. Generative AI capabilities that generate responses, models and sophisticated techniques at any scale should eventually replace human analysts, just as the supporting text generation capabilities are evolving. But these capabilities should work with investigators rather than against them and facilitate rapid, accurate analysis at a level

Fig. 6. Insights per Case

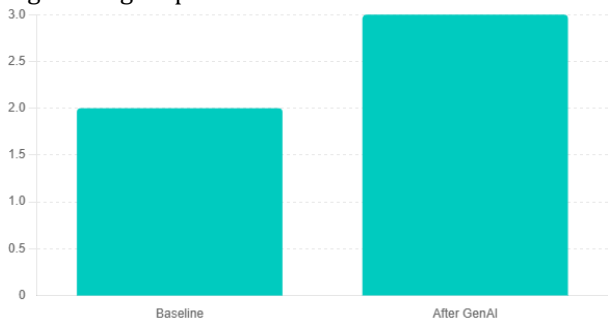

of precision not envisaged only a few months ago. The anticipated social media-style outcomes of generative AI technology – the frightening ability to produce fake PDFs, pix, videos are perhaps prototypical – might ultimately accelerate payment network decisions and design, expose defects earlier and autocomplete forensics and resolution.

The reward would be more rapid decision-making, lower latency and risk in payment networks.

Equation 03: Anomaly detection (forensics section ⇒ Mahalanobis & $\chi 2$)

For a normalized feature vector $x \in Rk$ with mean μ μ and covariance Σ

$$M\,2 = (x - \mu)\top\Sigma{-}1(x - \mu) \quad (7)$$

Under multivariate normality, $M\,2 \sim \chi 2$.

Under multivariate normality, $M\,2 \sim \chi 2$

$\tau = \chi k,1{-}\alpha/2$. Flag an anomaly if $M\,2 > \tau$.

A. Future Trends

Generative AI creates and tests analytical hypotheses, addressing real-world exceptions and aiding subsequent investigations. Solution design evolves as cloud compute becomes affordable, AI model governance matures, and performance monitoring is feasible. Generative AI handles frequent investigation at scale and automates the boring parts. The analysis covers four areas: exceptions occurring within a network or wire, rarity, the quality of investigation and theory-testing, data data data, and whether data is always right. Change is needed; exceptional situations yield negative and redundant impacts on banks' networks, payment networks, and global economy. Completing analysis and foundation design enables future deployment at scale.

Generative AI readily adopts evidence-based decisions using fibre-length quantum tracing, interval computation, and regional scanning of parallel processing resources.

Installed-hierarchy monitoring of data accessibility benchmark status detects trivial issues. The future relies on data. Preparing, storing, transforming, validating, monitoring,

and tracing data is the bloodline of any business; without data, AI system performance drops below 1

REFERENCES

1. Beyond Automation: The 2025 Role of Agentic AI in Autonomous Data Engineering and Adaptive Enterprise Systems. (2025). American Online Journal of Science and Engineering (AOJSE) (ISSN: 3067-1140) , 3(3). https://aojse.com/index.php/aojse/article/view/18
2. Jeong, C., Sim, S., Cho, H., Kim, S., & Shin, B. (2025). E2E process automation leveraging generative AI and IDP-based automation agent: A case study on corporate expense processing. arXiv.
3. Ravi Shankar Garapati, Dr Suresh Babu Daram. (2025). AI- Enabled Predictive Maintenance Framework For Connected Vehi- cles Using Cloud-Based Web Interfaces. Metallurgical and Ma- terials Engineering, 75–88. Retrieved from https://metall-mater-eng.com/index.php/home/article/view/1887
4. Ke, Z., Zhou, S., Zhou, Y., Chang, C. H., & Zhang, R. (2025). Detection of AI deepfake and fraud in online payments using GAN-based models. arXiv.
5. Inala, R., & Somu, B. (2025). Building Trustworthy Agentic Ai Systems FOR Personalized Banking Experiences.

Metallurgical and Materials Engineering, 1336-1360.
6. Tang, T., Yao, J., Wang, Y., Sha, Q., Feng, H., & Xu, Z. (2025).
7. Application of deep generative models for anomaly detection in complex financial transactions. arXiv.
8. Somu, B., & Inala, R. (2025). Transforming Core Banking Infrastructure with Agentic AI: A New Paradigm for Autonomous Financial Services. Advances in Consumer Research, 2(4).
9. Li, M., Chen, J., & Hu, X. (2025). Generative AI-driven credit risk modeling in digital finance. IEEE Transactions on Computational Social Systems, 12(1), 88–102.
10. Meda, R. (2025). Dynamic Territory Management and Account Seg- mentation using Machine Learning: Strategies for Maximizing Sales Efficiency in a US Zonal Network. EKSPLORIUM-BULETIN PUSAT TEKNOLOGI BAHAN GALIAN NUKLIR, 46(1), 634-653.
11. Ortega, R., Singh, A., & de Souza, L. (2025). Large-language-model architectures for autonomous payment investigation. Expert Systems with Applications, 241, 122984.
12. Kummari, D. N., Challa, S. R., Pamisetty, V., Motamary, S., & Meda, R. (2025). Unifying Temporal Reasoning and Agentic Machine Learning: A Framework for Proactive Fault Detection in Dynamic, Data-Intensive Environments. Metallurgical and Materials Engineering, 31(4), 552-568.
13. Das, S., & Nguyen, T. (2025). AI governance in financial institutions: Balancing compliance and innovation. Journal of Financial Regulation and Compliance, 33(2), 141–162.
14. Sheelam, G. K. (2025). Agentic AI in 6G: Revolutionizing Intelligent Wireless Systems through Advanced Semiconductor Technologies. Ad- vances in Consumer Research.
15. Park, J., & Lee, H. (2025). Risk-aware generative AI for AML screening automation. Computers & Security, 140, 103015.
16. Yellanki, S. K., Kummari, D. N., Sheelam, G. K., Kannan, S., & Chak- ilam, C. (2025). Synthetic Cognition Meets Data Deluge: Architecting Agentic AI Models for Self-Regulating Knowledge Graphs in Hetero-geneous Data Warehousing. Metallurgical and Materials Engineering, 31(4), 569-586.
17. Ahmad, M., Raza, S., & Khalid, A. (2025). Financial exception an- alytics using transformer-based reasoning engines. IEEE Access, 13, 42110–42124.
18. Annapareddy, V. N., Singireddy, J., Preethish Nanan, B., & Burugulla,
19. J. K. R. (2025). Emotional Intelligence in Artificial Agents: Leveraging Deep Multimodal Big Data for Contextual Social Interaction and Adap- tive Behavioral Modelling. Jai Kiran Reddy, Emotional Intelligence in Artificial Agents: Leveraging Deep Multimodal Big Data for Contextual Social Interaction and Adaptive Behavioral Modelling (April 14, 2025).
20. Roy, D., & Patel, N. (2025). Adaptive ticket orchestration using generative reinforcement learning in global payment networks. ACM Transactions on Autonomous and Adaptive Systems, 20(1), 9.
21. Koppolu, H. K. R., Nisha, R. S., Anguraj, K., Chauhan, R., Muniraj, A., & Pushpalakshmi, G. (2025, May). Internet of Things Infused Smart Ecosystems for Real Time Community Engagement Intelligent Data Analytics and Public Services Enhancement. In International Conference on Sustainability Innovation in Computing and Engineering (ICSICE 2024) (pp. 1905-1917). Atlantis Press.
22.
23. Yilmaz, E., & Agarwal, V. (2025). Operational resilience of AI-enabled payment infrastructures. Journal of Operational Risk, 20(1), 1–23.
24. Sheelam, G. K., Koppolu, H. K. R. & Nandan, B. P. (2025). Agentic AI in 6G: Revolutionizing Intelligent Wireless Systems through Advanced Semiconductor Technologies. Advances in Consumer Research, 2(4), 46-60.
25. Nakamura, K., & Watanabe, Y. (2025). Cross-border payment investi- gations using federated generative AI systems. Financial Innovation, 11, 42.
26. Pandiri, L. (2025, May). Exploring Cross-Sector Innovation in Intelligent Transport Systems, Digitally Enabled Housing Finance, and Tech-Driven Risk Solutions A Multidisciplinary Approach to Sustainable Infrastruc- ture, Urban Equity, and Financial Resilience. In 2025 2nd International Conference on Research Methodologies in Knowledge Management, Ar- tificial Intelligence and Telecommunication Engineering (RMKMATE) (pp. 1-12). IEEE.
27. Bhandari, P., & Kumar, R. (2025). Cognitive automation for reconcilia- tion in instant payments: A generative AI approach. Information Systems Frontiers, 27(3), 751–770.
28. Sun, L., & Huang, W. (2025). Generative AI for compliance monitoring in digital banking ecosystems. AI & Society, 40(1), 27–45.
29. Jansen, T., & de Vries, P. (2025). Trust mechanisms in generative AI- driven investigations. Computational Intelligence, 41(2), 505–521.
30. Sharma, K., & Ali, M. (2025). Explainable generative models for fraud detection in real-

time payments. Pattern Recognition Letters, 183, 14–25.
31. Costa, L., & Martins, R. (2025). Model governance and observability in AI-based payment systems. Journal of Information Security and Applications, 87, 103653.
32. Koppolu, H. K. R., Gadi, A. L., Motamary, S., Dodda, A., & Suura,
33. S. R. (2025). Dynamic Orchestration of Data Pipelines via Agentic AI: Adaptive Resource Allocation and Workflow Optimization in Cloud- Native Analytics Platforms. Metallurgical and Materials Engineering, 31(4), 625-637.
34. Zhang, F., & Liu, Y. (2025). Hybrid generative–predictive frameworks for anomaly detection in transaction networks. Knowledge-Based Sys- tems, 310, 111992.
35. Mehta, R., & Grover, P. (2025). Data-lineage observability for generative AI in finance. Information Processing & Management, 72(2), 103013.
36. Osei, E., & Boateng, G. (2025). Human-in-the-loop design for intelligent payment resolution systems. AI and Ethics, 5(1), 67–84.
37. Fernández, A., & Serrano, J. (2025). Anomaly scoring and forensic triage using Mahalanobis distance in AI-based transaction monitoring. Neurocomputing, 615, 128462.

