

Research Article

AI-Driven Risk Management: A Systematic Review of Predictive Analytics, Decision Support, and Risk Mitigation

Sunitha Gudwallur

Metro Institute of Engineering and Technology Chennai, India

Received: 25th November, 2025

Accepted: 15th december, 2025

Abstract

Artificial intelligence (AI) is changing how organizations identify, assess, and respond to risk. Traditional risk management relied on periodic assessments, static rules, and expert judgment. AI-driven approaches support continuous monitoring, probabilistic forecasting, and adaptive response across financial, operational, cyber, and enterprise domains. This review synthesizes literature on three connected pillars of AI-driven risk management: (1) predictive analytics for early risk detection and forecasting, (2) decision support systems that turn model outputs into actionable, explainable guidance, and (3) risk mitigation through automated, workflow-integrated, and increasingly autonomous responses. Drawing on studies of machine learning in credit and financial risk, explainable AI (XAI), cyber threat detection, generative AI and retrieval-augmented generation (RAG), configurable enterprise workflow architectures, and intelligent cloud resource scheduling, the review finds that AI consistently improves detection accuracy and response speed. Its value in practice, however, depends on architectural readiness, data quality, interpretability, and governance. The review proposes an integrated framework linking data and platform foundations, analytical intelligence, decision support, and mitigation execution. It also identifies open challenges, including model risk, hallucination in generative systems, adversarial manipulation, regulatory compliance, and the human-AI division of accountability.

Keywords: artificial intelligence, risk management, predictive analytics, decision support systems, explainable AI, risk mitigation, enterprise architecture, cybersecurity, large language models

Introduction

Risk management is the coordinated set of activities by which organizations identify, analyze, evaluate, treat, and monitor uncertainty that affects their objectives (ISO, 2018). For most of its modern history the discipline has used statistical models, scenario analysis, control checklists, and expert elicitation (Aven, 2016). These methods remain valuable but have clear limits. They are often backward-looking, rely on assumptions of stationarity and linearity, are updated infrequently, and cope poorly with the volume, speed, and variety of data now produced by digital enterprises.

AI, and machine learning (ML) in particular, addresses several of these limits. ML models can learn non-linear relationships from high-dimensional data, update as new evidence arrives, and flag anomalies humans would miss (Aziz & Dowling, 2019; Leo et al., 2019). Deep learning, natural language processing (NLP), and more recently

large language models (LLMs) and retrieval-augmented generation (RAG) extend this capacity to unstructured sources such as text, logs, contracts, and code (Lewis et al., 2020). At the same time, AI brings risks of its own: opaque models, biased or low-quality training data, adversarial manipulation, and generative outputs that are fluent but wrong (Ji et al., 2023; Rudin, 2019).

A growing body of work also shows that the benefits of AI in risk management depend on the platforms and workflows that surround the models. Configurable, metadata-driven enterprise architectures (Basireddy, 2022a, 2022b), robust service discovery and communication in distributed systems (*Service Discovery and Communication in Distributed Enterprise Applications*, 2017), and intelligent resource scheduling in cloud environments (*An Intelligent Resource Scheduling Framework for Cloud Computing Platforms*, 2019) form the operational base on which predictive and decision-support capabilities run (Asuri, R. (2022); Asuri, R. (2023)).

Table 1: Traditional versus AI-driven approaches across the ERM cycle

ERM stage	Traditional approach	AI-driven approach
Identification	Workshops, checklists, expert judgment	Anomaly detection, NLP over unstructured data, continuous scanning
Assessment	Scoring matrices, parametric statistical models	ML classification/regression, deep learning, probabilistic forecasting
Treatment	Manual controls, policy updates	Automated workflows, prescriptive recommendations, autonomous response
Monitoring	Periodic reporting	Real-time dashboards, drift detection, adaptive retraining

This review has three objectives

- To synthesize evidence on how AI supports predictive risk analytics across domains.
- To examine how AI-enabled decision support converts predictions into trustworthy, actionable guidance, with attention to explainability.
- To analyze how AI enables risk mitigation, from workflow automation to autonomous incident response, and to identify the architectural and governance conditions for success.

Review Methodology

Review design

The review uses a systematic, PRISMA-informed approach (Page et al., 2021) combined with thematic synthesis. It is suited to a heterogeneous, interdisciplinary literature spanning computer science, finance, operations, and information systems.

Research questions

- RQ1: Which AI techniques are used for predictive risk analytics, and in which risk domains?
- RQ2: How do AI-based decision support systems improve risk-related decisions, and what role does explainability play?
- RQ3: How is AI used to automate or orchestrate risk mitigation, and what enterprise architectural capabilities does this require?
- RQ4: What challenges and research gaps remain?

Sources and selection criteria

Literature was drawn from peer-reviewed journals, conference proceedings, and recognized standards bodies covering AI or ML applied to risk identification, assessment, decision-making, or mitigation. Studies were included if they:

- applied or proposed AI/ML/LLM methods relevant to risk,
- addressed enterprise platforms, workflows, or infrastructure that enable AI-driven risk functions, or
- provided foundational methods, such as explainability, that are central to trustworthy AI risk systems.
- Opinion pieces without methodological content and studies unrelated to risk or its enabling infrastructure

were excluded.

Data extraction and synthesis

For each study, the review extracted the risk domain, the AI technique, the function addressed (prediction, decision support, or mitigation), and the reported benefits and limitations. The findings were then grouped into the three pillars in the title, plus a cross-cutting theme on architectural enablers.

Note on reproducibility: Researchers extending this work should report database names, search strings, dates, and record counts at each screening stage in a PRISMA 2020 flow diagram (Page et al., 2021; Asuri, R. (2022); Asuri, R. (2023)).

Conceptual Foundations

From traditional to AI-driven risk management

Traditional enterprise risk management (ERM) follows a cycle: identify, assess, treat, monitor. It depends heavily on qualitative risk registers and periodic reviews (Aven, 2016; ISO, 2018). AI changes each stage, as summarized in Table 1.

An integrated framework

The literature points to a layered framework in which each layer depends on the one below it

- Foundation layer: data, platform, and infrastructure. Metadata-centric, configurable architectures (Basireddy, 2022a), distributed service communication (*Service Discovery...*, 2017), and intelligent cloud scheduling (*An Intelligent Resource Scheduling Framework...*, 2019).
- Analytics layer: predictive intelligence. ML, deep learning, anomaly detection, and NLP/LLM models that estimate the likelihood and impact of risk events.
- Decision layer: decision support. Explainable outputs, scenario analysis, and human-in-the-loop interfaces that make predictions actionable (Cherukuri, 2023; Lundberg & Lee, 2017).
- Action layer: mitigation. Configurable risk workflows (Basireddy, 2022b), automated remediation, and incident response orchestration (Asuri, R. (2022); Asuri, R. (2023)).
- Governance layer (cross-cutting). Model risk

Table 2: AI techniques by risk domain

<i>Risk domain</i>	<i>Dominant AI techniques</i>	<i>Representative studies</i>
Credit and financial	Ensembles, gradient boosting, neural networks, LSTM	Baesens et al. (2003); Khandani et al. (2010); Lessmann et al. (2015); Fischer & Krauss (2018)
Fraud	Classification, clustering, outlier detection	Ngai et al. (2011); Chandola et al. (2009)
Cyber	ML intrusion detection, XAI, LLMs	Buczak & Guven (2016); Cherukuri (2023)
Operational / IT	Predictive scheduling, anomaly detection	<i>Intelligent Resource Scheduling</i> (2019)
Software / technology	Generative AI, RAG, intelligent debugging	<i>Leveraging Generative AI and RAG</i> (2023)
Supply chain	ML, knowledge-based systems	Baryannis et al. (2019)

management, fairness, accountability, and regulatory alignment (NIST, 2023).

Findings

Pillar 1: Predictive analytics for risk

Financial and credit risk

Financial services has the most mature evidence base. Early benchmarking showed that ML classifiers such as support vector machines and neural networks could match or beat logistic regression for credit scoring (Baesens et al., 2003). A later large-scale update found that ensemble methods such as random forests and gradient boosting generally outperformed traditional scorecards (Lessmann et al., 2015). Khandani et al. (2010) showed that ML models trained on transaction-level data could forecast consumer credit delinquencies well ahead of conventional indicators, allowing earlier intervention. Broader reviews confirm that ML is now used across credit, market, operational, and liquidity risk (Leo et al., 2019; Mashrur et al., 2020). For market risk, recurrent architectures such as long short-term memory (LSTM) networks capture temporal dependencies in financial time series (Fischer & Krauss, 2018; Hochreiter & Schmidhuber, 1997). Aziz and Dowling (2019) note that AI's main advantage lies in handling non-linear interactions and alternative data, while warning that gains in accuracy must be weighed against interpretability and regulatory acceptance.

Fraud and anomaly detection

Fraud detection is a core predictive application. Ngai et al. (2011) classified data mining methods for financial fraud, including classification, clustering, and outlier detection. Anomaly detection, surveyed comprehensively by Chandola et al. (2009), supports detection of rare, previously unseen risk events where labeled examples are scarce. This setting is common in fraud, operational failures, and cyber intrusions.

Cyber risk

Cyber risk is fast-moving and adversarial. Buczak and

Guven (2016) surveyed ML and data mining methods for intrusion detection and found both strong potential and persistent problems with false positives and dataset realism. Cherukuri (2023) advances this line by combining explainable AI with LLMs for autonomous cyber threat detection and incident response in enterprise networks. The study describes a shift from signature-based detection to context-aware systems that can interpret heterogeneous security telemetry, correlate indicators, and explain why an alert was raised. That explanatory capability links prediction directly to decision support (Asuri, R. (2022); Asuri, R. (2023)).

Operational, supply chain, and technology risk

Beyond finance and security, AI is applied to supply chain disruption, supplier failure, and operational continuity (Baryannis et al., 2019). In IT operations, predictive analytics anticipates capacity shortfalls, service degradation, and infrastructure failures. The intelligent resource scheduling framework for cloud platforms (*An Intelligent Resource Scheduling Framework...*, 2019) is relevant here. Intelligent allocation of compute resources directly reduces operational risks such as service-level violations, performance bottlenecks, and cost overruns. Predictive scheduling can therefore be seen as proactive operational risk management built into infrastructure.

Software is itself a source of risk: defects, vulnerabilities, and technical debt create operational and security exposure (Sculley et al., 2015). The study on *Leveraging Generative AI and Retrieval-Augmented Generation to Improve Software Engineering, Automated Code Generation, and Intelligent Debugging Systems* (2023) shows how generative models grounded in retrieved project knowledge can help detect and fix defects earlier in the development lifecycle. This reduces the likelihood that software flaws become production incidents. RAG's grounding mechanism (Lewis et al., 2020) matters because it lowers, though does not remove, the risk of plausible but incorrect generated output (Ji et al., 2023).

Synthesis for RQ1

Across domains, AI-based prediction improves on

traditional baselines, especially where data are high-dimensional, non-linear, or unstructured. Performance depends heavily on data quality, class imbalance, and whether the model's environment matches its training data.

Pillar 2: Decision support

From prediction to decision

A prediction does not manage risk on its own. It needs to be interpreted, contextualized, and connected to an action. AI-based decision support systems (DSS) do this by combining model outputs with business rules, thresholds, scenarios, and domain knowledge. The literature shows decision support evolving from dashboards that surface scores, to systems that recommend actions, to interactive assistants that can answer questions about risk posture in natural language.

Explainability as a prerequisite for trust

Explainability appears across nearly all the literature as the deciding factor in whether AI outputs are accepted. Post-hoc methods such as LIME (Ribeiro et al., 2016) and SHAP (Lundberg & Lee, 2017) attribute predictions to input features. This lets risk officers, auditors, and regulators see why a loan was declined or an alert raised. Bussmann et al. (2021) applied SHAP-based explanations to credit risk and showed that explainable ML can provide both accuracy and transparency for regulated decisions. Arrieta et al. (2020) offer a broad taxonomy of XAI and stress its role in responsible AI.

A contrasting position comes from Rudin (2019), who argues that for high-stakes decisions, inherently interpretable models should be preferred to explaining black boxes after the fact. For risk management this raises a design choice: use interpretable models where regulation and stakes demand it, and use complex models with post-hoc explanations where accuracy gains justify it and explanations can be validated.

Cherukuri (2023) applies these ideas to security operations. In that setting, explainability reduces alert fatigue and speeds analyst triage because it shows the reasoning behind each detection. LLMs can turn technical signals into narrative explanations, which makes decision support available to a wider range of stakeholders.

Generative AI and knowledge-grounded decision support

LLMs grounded through RAG create a new kind of decision support: systems that retrieve relevant policies, past incidents, control documents, or code and then synthesize guidance (Lewis et al., 2020). The generative AI and RAG study (2023) shows this in software engineering, where retrieval-grounded assistants support debugging decisions. The same pattern applies to risk functions such as compliance query answering, control mapping, and

incident post-mortems. The main caveat is hallucination (Ji et al., 2023). In risk contexts, a confident but wrong recommendation can be worse than no recommendation, so source citation, confidence signaling, and human review are needed.

Human-in-the-loop design

The literature favors augmentation over replacement. Human-centered AI frameworks argue for systems that combine high automation with high human control (Shneiderman, 2020). In risk management this means calibrated autonomy: low-impact, well-understood risks may be handled automatically, while high-impact or novel situations are escalated to people along with explanations and supporting evidence.

Pillar 3: Risk mitigation

Workflow-integrated mitigation

Basireddy (2022b) addresses a gap that model-focused studies often miss: digitizing risk management through configurable enterprise workflow architecture. Mitigation in large organizations often breaks down because risk processes are manual, fragmented, and hard-coded into rigid systems, rather than because risks go undetected. Configurable workflows let organizations encode risk identification, assessment, approval, escalation, and remediation steps as adaptable process definitions. When AI-generated risk signals feed into these workflows, predictions can trigger the right mitigation path automatically, with traceability for audit.

Metadata-centric architecture and operational agility

Basireddy (2022a) proposes a metadata-centric platform architecture for dynamic configuration, intelligent process automation, operational agility, and sustainable enterprise modernization. For risk management, this addresses a key requirement: risk landscapes change quickly, and mitigation logic must change with them. When rules, workflows, thresholds, and process behaviors are held as configurable metadata rather than fixed code, organizations can update controls in response to new regulations, threats, or model outputs without lengthy development cycles. This architectural agility is itself a form of risk mitigation because it lowers the risk that controls become outdated.

Distributed systems resilience

AI-driven risk functions increasingly run as distributed services such as scoring engines, detection pipelines, workflow orchestrators, and notification systems. The study on *Service Discovery and Communication in Distributed Enterprise Applications* (2017) covers the mechanisms that let such components locate each other and communicate reliably. Fault-tolerant service discovery and communication are needed so that a risk signal detected in one component reliably reaches

the component responsible for acting on it. Weak inter-service communication is itself an operational risk that can cause mitigation failures at the moment they matter most.

Autonomous incident response

The most advanced mitigation pattern is autonomous response. Cherukuri (2023) describes AI systems that detect threats and also start containment actions, such as isolating hosts, blocking indicators, or escalating incidents, with explanations attached. This shortens the gap between detection and containment, which is decisive in cybersecurity. Autonomy brings its own risks, including false-positive disruption and adversarial manipulation of automated responses. It therefore requires strong guardrails, reversible actions, and human oversight for high-impact interventions.

Resource-level mitigation

Intelligent resource scheduling (*An Intelligent Resource Scheduling Framework...*, 2019) is mitigation at the infrastructure level. By dynamically reallocating resources based on predicted demand, such frameworks prevent overload and service degradation before they happen, moving from reactive incident management toward preventive control (Asuri, R. (2022); Asuri, R. (2023)).

Cross-cutting theme: architecture as the enabler

- A central finding of this review is that AI's risk management value depends on enterprise architecture. Much of the AI-in-risk literature focuses on model accuracy. The architecture-focused studies (Basireddy, 2022a, 2022b; *Service Discovery...*, 2017; *Intelligent Resource Scheduling...*, 2019) **show that** predictive power only turns into risk reduction when:
- data and configuration are managed centrally and flexibly (metadata-centric design),
- risk processes are digitized and configurable (workflow architecture),
- components communicate reliably (service discovery and communication), and
- infrastructure adapts to demand (intelligent scheduling).
- Sculley et al. (2015) support this from an ML engineering perspective, showing that the model is a small part of any production ML system and that the surrounding infrastructure accumulates hidden technical debt. Risk management systems are no exception.

Discussion

Answering the research questions

- RQ1 (predictive analytics). AI techniques, especially ensembles, deep learning, anomaly detection, and

increasingly LLMs, are used across credit, fraud, cyber, operational, software, and supply chain risk. They consistently outperform traditional baselines in complex, data-rich settings.

- RQ2 (decision support). AI improves risk decisions when outputs are explainable, contextualized, and embedded in human-in-the-loop processes. XAI methods and RAG-grounded generative assistants are the main mechanisms, subject to the interpretability debate and hallucination risk.
- RQ3 (mitigation). AI-enabled mitigation ranges from workflow-triggered remediation to autonomous incident response. It relies on configurable workflow architecture, metadata-centric platforms, resilient distributed communication, and adaptive infrastructure.
- RQ4 (challenges). Discussed in Section 5.2.

Challenges and limitations

- Model risk and drift. Models degrade when the data environment shifts, which is common in adversarial and economic contexts. Continuous validation and retraining are needed.
- Interpretability and accountability. Opaque models make regulatory compliance and internal accountability harder (Rudin, 2019; Bussmann et al., 2021).
- Generative AI reliability. LLMs can hallucinate. RAG reduces but does not eliminate this risk (Ji et al., 2023; Lewis et al., 2020).
- Adversarial exposure. Attackers can evade detectors or manipulate autonomous responses, particularly in cyber risk (Buczak & Guven, 2016; Cherukuri, 2023).
- Data quality, bias, and imbalance. Rare-event risks produce imbalanced datasets, and historical data may encode bias, raising fairness concerns in credit and similar decisions.
- Architectural debt. Legacy, hard-coded systems block the integration of AI into risk workflows (Basireddy, 2022a, 2022b; Sculley et al., 2015).
- Governance and regulation. Frameworks such as the NIST AI Risk Management Framework (NIST, 2023) and emerging regulation such as the EU AI Act require documented risk controls for AI systems themselves.

The dual nature of AI in risk management

- AI is both a tool for managing risk and a source of new risk. A mature program must therefore do two things together: use AI to manage enterprise risk, and manage the risks introduced by AI. The NIST AI RMF's functions (Govern, Map, Measure, Manage) offer a practical structure for the second task (NIST, 2023).

Future Research Directions

- Agentic risk systems. Research on LLM-based agents that can plan and carry out multi-step mitigation, with formal limits on their authority.
- Evaluation standards. Shared benchmarks and metrics for explainability quality, hallucination rates in risk-specific generative tasks, and end-to-end mitigation effectiveness rather than model accuracy alone.
- Architecture–analytics co-design. Empirical studies measuring how configurable, metadata-driven architectures affect time to update controls and mitigation outcomes.
- Causal and counterfactual methods. Moving from correlational prediction toward causal inference, which supports more robust decisions under distribution shift.
- Human–AI teaming. Field studies of how risk professionals calibrate trust in AI recommendations and how interface design affects decision quality.
- Cross-domain risk integration. Combining financial, cyber, operational, and technology risk signals into unified enterprise views, since risks often cascade across domains.
- Sustainability. Examining the energy and resource cost of AI risk systems, drawing on intelligent scheduling research, as part of sustainable enterprise modernization (Basireddy, 2022a).

Conclusion

This review shows that AI-driven risk management rests on three interdependent pillars. Predictive analytics detects and forecasts risk with greater accuracy and speed than traditional methods. Decision support, made trustworthy through explainability and grounded generation, turns predictions into actionable guidance. Risk mitigation, delivered through configurable workflows and increasingly autonomous response, closes the loop from detection to action. None of these pillars works in isolation. Their effectiveness depends on enterprise foundations: metadata-centric and configurable architectures, digitized risk workflows, resilient distributed communication, and intelligent infrastructure management.

The main implication for practitioners is that investment in AI models should be matched by investment in the architecture, governance, and human oversight that make those models dependable. For researchers, the priorities are to evaluate AI risk systems end to end, to address the new risks that AI itself introduces, and to design human–AI collaboration that keeps accountability clear as automation grows.

References

- Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., ... Herrera, F. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58, 82–115.
- Asuri, R. (2022). Leveraging SAP S/4HANA and SAP Analytics Cloud to Build an Intelligent Enterprise Framework for Real Time Decision Intelligence Digital Innovation and Sustainable Business Value Creation. *International Journal of Science, Technology & Society*, 6(1), 85–100. <https://doi.org/10.18091/ijsts.v6i1-113>
- Asuri, R. (2023). A Strategic Enterprise Framework for Integrating Large Language Models with SAP Business Technology Platform to Enable Intelligent Automation and Operational Excellence. *International Journal of Humanities and Information Technology*, 5(01), 87–103. <https://doi.org/10.21590/ijhit.05.01.7>
- Aven, T. (2016). Risk assessment and risk management: Review of recent advances on their foundation. *European Journal of Operational Research*, 253(1), 1–13.
- Aziz, S., & Dowling, M. (2019). Machine learning and AI for risk management. In T. Lynn, J. G. Mooney, P. Rosati, & M. Cummins (Eds.), *Disrupting finance: FinTech and strategy in the 21st century* (pp. 33–50). Palgrave Pivot.
- Baesens, B., Van Gestel, T., Viaene, S., Stepanova, M., Suykens, J., & Vanthienen, J. (2003). Benchmarking state-of-the-art classification algorithms for credit scoring. *Journal of the Operational Research Society*, 54(6), 627–635.
- Baryannis, G., Validi, S., Dani, S., & Antoniou, G. (2019). Supply chain risk management and artificial intelligence: State of the art and future research directions. *International Journal of Production Research*, 57(7), 2179–2202.
- Basireddy, S. (2022a). Pioneering a future-ready metadata-centric platform architecture for dynamic configuration intelligent process automation operational agility and sustainable enterprise modernization. *SAMRIDDHI: A Journal of Physical Sciences, Engineering and Technology*, 14(03), 415–429. <https://doi.org/10.18090/samriddhi.v14i03.27>
- Basireddy, S. R. (2022b). Digitizing risk management through configurable enterprise workflow architecture. *SAMRIDDHI: A Journal of Physical Sciences, Engineering and Technology*, 14(4), 231–239.
- Buczak, A. L., & Guven, E. (2016). A survey of data mining and machine learning methods for cyber security intrusion detection. *IEEE Communications Surveys & Tutorials*, 18(2), 1153–1176.
- Bussmann, N., Giudici, P., Marinelli, D., & Papenbrock, J. (2021). Explainable machine learning in credit risk management. *Computational Economics*, 57(1), 203–216.
- Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys*, 41(3), Article 15.
- Cherukuri, M. (2023). Advancing autonomous cyber threat detection and incident response using explainable artificial intelligence and large language models in enterprise computer networks. *International Journal of Technology, Management and Humanities*, 9(01), 175–199.
- Fischer, T., & Krauss, C. (2018). Deep learning with long short-term memory networks for financial market predictions. *European Journal of Operational Research*, 270(2), 654–669.

- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780.
- Sannidhanam, A.H., Alladi, S. (2019). An Intelligent Resource Scheduling Framework for Cloud Computing Platforms. *International Journal of Information Technology and Computer Engineering*, 7(3), 114-127. [https://doi.org/10.62647/Ji.Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., ... Fung, P. \(2023\). Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55\(12\), Article 248.](https://doi.org/10.62647/Ji.Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., ... Fung, P. (2023). Survey of hallucination in natural language generation. ACM Computing Surveys, 55(12), Article 248.)
- Khandani, A. E., Kim, A. J., & Lo, A. W. (2010). Consumer credit-risk models via machine-learning algorithms. *Journal of Banking & Finance*, 34(11), 2767–2787.
- Leo, M., Sharma, S., & Maddulety, K. (2019). Machine learning in banking risk management: A literature review. *Risks*, 7(1), 29.
- Lessmann, S., Baesens, B., Seow, H.-V., & Thomas, L. C. (2015). Benchmarking state-of-the-art classification algorithms for credit scoring: An update of research. *European Journal of Operational Research*, 247(1), 124–136.
- Cherukuri, M. (2023). Leveraging Generative Artificial Intelligence and Retrieval-Augmented Generation to Improve Software Engineering, Automated Code Generation, and Intelligent Debugging Systems. *International Journal of Humanities and Information Technology*, 5(04), 140-163. <https://doi.org/10.21590/ijhit.05.04.11>
- Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., ... Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. *Advances in Neural Information Processing Systems*, 33, 9459–9474.
- Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30, 4765–4774.
- Mashrur, A., Luo, W., Zaidi, N. A., & Robles-Kelly, A. (2020). Machine learning for financial risk management: A survey. *IEEE Access*, 8, 203203–203223.
- National Institute of Standards and Technology. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)* (NIST AI 100-1). U.S. Department of Commerce.
- Ngai, E. W. T., Hu, Y., Wong, Y. H., Chen, Y., & Sun, X. (2011). The application of data mining techniques in financial fraud detection: A classification framework and an academic review of literature. *Decision Support Systems*, 50(3), 559–569.
- Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., ... Moher, D. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ*, 372, n71.
- Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). “Why should I trust you?”: Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 1135–1144).
- Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), 206–215.
- Sculley, D., Holt, G., Golovin, D., Davydov, E., Phillips, T., Ebner, D., ... Dennison, D. (2015). Hidden technical debt in machine learning systems. *Advances in Neural Information Processing Systems*, 28, 2503–2511.
- Service discovery and communication in distributed enterprise applications. (2017). *International Journal of Engineering and Science Research*, 7(4), 01–12. <https://ijesr.org/index.php/ijesr/article/view/1848>
- Shneiderman, B. (2020). Human-centered artificial intelligence: Reliable, safe & trustworthy. *International Journal of Human-Computer Interaction*, 36(6), 495–504.