Proceedings of CONF-BPS 2026 Symposium: GenAI, Labour Markets, and the Economics of Human and Financial Capital
DOI: 10.54254/2754-1169/2026.LD31794

The Impact of Artificial Intelligence on Enterprise Risk Management

Qingyang Long

School of Accounting and Finance, The Hong Kong Polytechnic University, Hong Kong, China
sherry.long@connect.polyu.hk

Abstract. The rapid development of Artificial Intelligence (AI) has profoundly influenced various aspects of enterprise operations, particularly in the field of risk management. As organizations face increasingly complex, dynamic, and interconnected risk environments, the integration of cutting-edge AI technologies offers new opportunities as well as challenges for identifying, assessing, and mitigating multifaceted risks. This paper explores the impact of AI on enterprise risk management (ERM) by reviewing relevant theories, analyzing practical applications, and discussing associated risks and challenges. Through case studies in financial services, supply chain management, and cybersecurity, the research demonstrates that AI enhances risk detection, improves decision-making, and increases operational efficiency. However, the widespread adoption of AI also introduces new risks, such as algorithmic bias, data privacy concerns, and model transparency issues. The study concludes that while AI significantly advances ERM, organizations must adopt robust, proactive governance frameworks to address these emerging challenges and ensure responsible deployment of AI.

Keywords: Artificial Intelligence, Enterprise Risk Management, Algorithmic Bias, Data Privacy, Decision-Making

1. Introduction

In the era of digital transformation, Artificial Intelligence (AI) has emerged as a pivotal technology reshaping business landscapes across industries. AI refers to the simulation of human intelligence processes by machines, especially computer systems, which include learning, reasoning, and self-correction [1]. The integration of AI into enterprise operations has accelerated in recent years, driven by advancements in machine learning, natural language processing, and big data analytics. One of the most significant areas affected by this technological revolution is enterprise risk management (ERM). ERM is a structured and systematic approach to identifying, assessing, and managing risks that may affect an organization's ability to achieve its objectives [2]. Traditional risk management methods often rely on historical data, expert judgment, and manual processes, which can be time-consuming and prone to human error. The increasing complexity and interconnectivity of modern business environments have exposed the limitations of these conventional approaches, necessitating more agile and data-driven solutions. AI offers transformative potential for ERM by enabling real-time risk detection, predictive analytics, and automated decision-making. For instance, AI-powered algorithms can analyze vast amounts of structured and unstructured data to identify emerging risks, detect anomalies, and forecast potential threats with greater accuracy than traditional methods [3]. However, the adoption of AI in risk management introduces new challenges, like algorithmic bias, data privacy concerns, and the need for transparent and explainable models.

This paper aims to systematically examine the impact of AI on enterprise risk management. The study employs a qualitative approach, drawing on academic literature, industry reports, and real-world examples to provide a comprehensive analysis. The findings contribute to a deeper understanding of how AI is reshaping risk management practices and highlight the importance of responsible AI governance in the enterprise context.

2. Theoretical foundations

Modern AI introduces several key technical capabilities, such as real-time data analysis, automation, unstructured data processing, and pattern recognition, which collectively overcome the limits of traditional ERM. For example, AI systems can continuously monitor and analyze data streams in real time, detecting emerging threats instantly rather than relying on periodic reviews [4,5]. Automation technologies relieve employees of repetitive compliance tasks, enabling instant policy updates and streamlined reporting. AI also excels at processing unstructured inputs: natural language processing (NLP) allows the system to scan legal texts, news and social media, while computer vision can interpret images for risk signal. Crucially, machine learning (ML) algorithms identify hidden patterns and anomalies across vast datasets—spotting fraud or market shifts that human teams might miss [4]. These features break through legacy ERM constraints: AI's continuous, automated analysis replaces quarterly or manual snapshots, broadens inputs beyond single historical data sources, and improves the timeliness and scope of risk detection [5].

Enterprise Risk Management is a holistic approach to managing risks across an organization. The Committee of Sponsoring Organizations of the Treadway Commission (COSO) defines ERM as "the culture, capabilities, and practices, integrated with strategy-setting and performance, that organizations rely on to manage risk in creating, preserving, and realizing value" [6]. Key components of ERM frameworks include risk identification, risk assessment, risk response, risk monitoring, and communication. Traditional ERM relies heavily on historical data, scenario analysis, and expert judgment. However, the increasing volume, velocity, and variety of data in modern enterprises have challenged the effectiveness of these methods. AI-driven ERM seeks to augment traditional approaches by leveraging advanced analytics and automation to enhance risk management capabilities [7].

3. The impact of AI on enterprise risk management

3.1. AI for risk identification

AI combines NLP and ML to enable advanced risk detection. NLP models can automatically interpret unstructured text (such as contracts, news articles, or regulatory filings) to extract relevant risk terms. At the same time, ML algorithms learn from both structured and unstructured data to recognize risk patterns and anomalies that indicate fraud, credit issues, or operational failures. In practice, this means algorithms can flag unusual transaction patterns or hidden correlations across datasets far beyond human capacity.

By leveraging these AI capabilities, ERM processes overcome the traditional reliance on historical data and periodic reviews. Instead of sampling a few incidents, AI systems ingest diverse data sources (historical records, live feeds, external news, etc.) and continuously update risk signals. This breaks the old "once-per-quarter" assessment cycle and reduces blind spots. Real-time analytics give decision makers immediate alerts—transforming risk identification from a delayed, manual task into a dynamic, data-driven workflow. As a result, organizations can detect emerging threats early and allocate mitigation resources proactively.

A prominent example is JPMorgan's COiN (Contract Intelligence) platform. COiN uses NLP and ML to automatically parse loan agreements and legal documents. It scans each contract for key clauses (e.g. default terms or covenants) and highlights potential risk points. In one report, COiN processed 12,000 commercial loan agreements in seconds, a task that formerly took human lawyers 360,000 hours per year. This demonstrates how AI can rapidly extract critical risk insights from text at scale. The practical impact is dramatic. In JPMorgan's case, COiN's automation yielded a roughly 360,000-man-hour reduction in contract review time, with much higher consistency and fewer errors than manual analysis [8]. Generally, applying NLP/ML to risk identification greatly boosts efficiency (processing thousands of data points in seconds) and accuracy (catching easily-overlooked clauses or anomalies). Pilot implementations often deliver multifold speedups and higher detection rates.

3.2. Improved decision-making and risk mitigation

AI enables more informed and timely decision-making by providing predictive insights and scenario analysis. Predictive analytics models can forecast potential risks based on historical data and current trends, allowing organizations to proactively implement mitigation strategies [7]. In supply chain management, AI-driven tools optimize inventory levels, predict disruptions, and recommend alternative suppliers. Global firms leverage IBM Watson's AI platform to monitor supply chain risks. By analyzing data from weather reports, geopolitical events, and supplier performance, Watson provides real-time risk alerts and recommendations. For instance, during the pandemic, Watson helped companies identify at-risk suppliers and reroute shipments, minimizing operational disruptions [9].

3.3. Automation of risk monitoring and reporting

AI automates routine risk monitoring and reporting tasks, freeing up human resources for more strategic activities. Robotic process automation (RPA) and AI-powered dashboards provide real-time visibility into risk exposures and compliance status [3]. Automated systems can scan for regulatory changes, flag compliance breaches, and generate audit-ready reports.

Table 1. Comparison of traditional vs. AI-driven risk monitoring

Feature	Traditional Risk Monitoring	AI-Driven Risk Monitoring
Data Processing Speed	Manual, slow	Automated, real-time
Data Sources	Limited, structured	Diverse, unstructured
Anomaly Detection	Rule-based	Machine learning-based
Reporting	Periodic, manual	Continuous, automated
Scalability	Limited	High

As shown in Table 1, AI-driven risk monitoring offers significant advantages in speed, data diversity, and scalability compared to traditional methods.

3.4. Strengthening cybersecurity risk management

Cybersecurity is a critical area where AI has made substantial contributions to risk management. AI-powered systems detect and respond to cyber threats by analyzing network traffic, identifying suspicious patterns, and automating incident response [10]. Machine learning models can adapt to evolving attack vectors, providing a dynamic defense against cyber risks.

Darktrace, a cybersecurity company, uses AI algorithms to detect and respond to cyber threats in real time. Its Enterprise Immune System learns the normal "pattern of life" for every user and device, enabling it to identify subtle deviations indicative of cyberattacks. In 2017, Darktrace's AI detected a ransomware attack at a manufacturing firm within seconds, allowing the company to contain the threat before significant damage occurred [11].

4. Risks and challenges of AI in enterprise risk management

4.1. Algorithmic bias and fairness

Algorithmic bias arises when AI models reproduce or amplify prejudices present in training data, leading to systematic unfairness in outcomes. In an ERM context, biased algorithms might under- or over-rate certain risks for particular groups of stakeholders or assets, producing discriminatory effects. For example, a credit-risk model that unfavorably scores applicants from a protected demographic. Such biased outcomes can violate anti-discrimination laws, erode customer and employee trust, and expose the organization to legal and reputational damagep [12,13]. Unchecked bias in risk models can distort decision-making and undermine the ERM integrity.

Mitigating algorithmic bias involves proactive governance of data and models. Organizations should utilize diverse training datasets to minimize sample-skew bias and apply fairness-aware techniques during model development. Regular bias audits are essential, evaluating model outputs across demographic groups to identify disparate impacts. Transparent models enable visibility into decision-making processes, aiding in bias identification. Cross-functional oversight teams are necessary to review AI decisions for blind spots. Overall, incorporating fairness checks throughout the AI lifecycle promotes equity and compliance in ERM systems [14,15].

4.2. Data privacy and security

AI-based ERM often involves handling substantial amounts of sensitive corporate and personal data, leading to privacy and security risks such as unauthorized data exposure and misuse of personal information. The use of personally identifiable information (PII) may result in data breaches, attracting regulatory penalties and harming reputations [16]. Indeed, any leak of AI training data may violate data protection laws and result in fines. In ERM specifically, compromised data can lead to incorrect risk assessments or leaking of strategic information, further compounding business risk. Organizations must therefore anticipate that AI introduces novel attack vectors (e.g. poisoned training sets, model inversion attacks) alongside traditional cybersecurity threats.

To reduce privacy and security risk, enterprises must enforce robust data governance and technical measures. Key strategies include minimizing and anonymizing datasets to prevent AI algorithms from accessing raw PII, employing encryption for data at rest and in transit, and enforcing strict access controls [16]. Utilizing privacy-preserving AI methods like federated learning can further protect sensitive information. Organizations must also adhere to comprehensive policies such as GDPR/CCPA compliance, conduct regular security audits, and establish incident-response plans for AI-related breaches. Additionally, training employees on data-handling practices is essential to minimize human error [15]. In short, combining encryption, strict access management, privacy-by-design in AI development, and continuous monitoring enables firms to safeguard sensitive data while leveraging AI in ERM.

4.3. Model transparency and explainability

Many ERM AI models are often "black boxes" with opaque internal logic, which hinders stakeholders from understanding or verifying how the system reaches specific conclusions. Such opaqueness can undermine trust in the AI-driven risk process and complicate regulatory compliance. For example, EU regulations require that decisions affecting individuals be explainable and that personal data processing be transparent. When a model's reasoning is hidden, it becomes difficult to provide the required explanations or to audit decisions. Moreover, if model outputs are questioned, lack of insight into the model makes error analysis and correction hard. Thus, black-box AI can generate regulatory, legal, and operational risk: without interpretability, organizations may fail to demonstrate due diligence and may make uncorrectable mistakes [12,16].

Enterprises can address this by adopting explainable AI (XAI) techniques and practices. This may involve choosing simpler, inherently interpretable models for high-stakes decisions, or applying post-hoc interpretability tools to complex models. Producing model documentation – including clear model cards, data lineage reports, and decision-logic summaries – increases transparency for reviewers. Continuous monitoring of model behavior (tracking drift, validating outputs on known cases) also helps catch problems. As IBM observes, XAI "helps promote end user trust, model auditability and productive use of AI" and "mitigates compliance, legal, security and reputational risks". Organizations should also provide human oversight: incorporate human review checkpoints and explain the system's logic to users. Zscaler notes that transparency is "critical under GDPR," and explainable AI data processing and decisions are key to compliance and stakeholder trust. Clear AI decision paths allow firms to validate risk assessments, ensure accountability and maintain ERM process confidence [12,16].

4.4. Integration and change management

Integrating AI into an existing ERM framework is as much a people and process challenge as a technical one. Organizations often struggle with aligning new AI tools to current workflows and with preparing staff for the shift. Common issues include technology misalignment and human resistance. Employees may fear job displacement, distrust AI outputs, or lack the skills to work with AI-enhanced tools. Without proper change management, this can lead to low adoption, wasted projects, and unment benefits. For example, one study finds that 63% of organizations cite human factors as a primary AI challenge, with 38% of AI initiatives failing due to insufficient employee training. Additionally, poor integration planning can isolate AI pilots from core systems, preventing models from scaling and delivering lasting value.

To overcome these challenges, enterprises need a structured change management and integration strategy. First, hands-on workshops and education bridge the skill gap – 38% of adoption failures stem from lack of training. Additionally, executives should sponsor AI projects and communicate their vision, since 43% of AI failures are tied to insufficient executive support [13]. Organizations should conduct stakeholder analysis, develop clear communication plans, set success metrics, and designate change champions within teams. For instance, EPAM recommends role-based training programs and "AI champions" in each business unit to build capacity and trust [17]. Pilots should be user-centric, enhancing rather than disrupting existing processes, and involve end-users early so that AI tools fit real needs. Finally, reliable data pipelines and merging AI outputs into risk reporting ensure insights feed seamlessly into the ERM workflow. By combining training, executive sponsorship, clear communication, and iterative deployment, organizations can surmount integration hurdles and embed AI effectively into their risk management culture [13,17].

5. Conclusion

This paper has provided a comprehensive analysis of the impact of AI on enterprise risk management, drawing on academic literature, industry reports, and real-world examples. The findings highlight the dual nature of AI as both an enabler and a source of risk in the enterprise context. Organizations must balance the benefits of AI-driven risk management with the need to address emerging ethical, legal, and operational challenges.

Future research should focus on developing standardized frameworks for AI governance in risk management, exploring the implications of AI for regulatory compliance, and investigating the long-term effects of AI adoption on organizational resilience. Additionally, interdisciplinary collaboration between technologists, risk managers, and policymakers will be essential to harness the full potential of AI while safeguarding against unintended consequences.

In conclusion, AI has the potential to revolutionize enterprise risk management, but its successful implementation requires a holistic approach that integrates technological innovation with ethical responsibility and organizational change management. By embracing AI thoughtfully and proactively addressing associated risks, enterprises can enhance their risk management capabilities and achieve sustainable competitive advantage in an increasingly uncertain world.

References

[1] Russell, S., & Norvig, P. (2021). Artificial Intelligence: A Modern Approach (4th ed.). Pearson.

[2] Frigo, M. L., & Anderson, R. J. (2011). Strategic Risk Management: A Foundation for Improving Enterprise Risk Management and Governance. Journal of Corporate Accounting & Finance, 22(3), 81-88.

[3] Deloitte. (2020). The future of risk: New game, new rules. https: //www2.deloitte.com/global/en/pages/risk/articles/the-future-of-risk.html

[4] Metricstream. (2025). The Ultimate Guide to AI in Risk Management https: //www.metricstream.com/learn/ai-risk-management.html#: ~: text=AI, resilience%20and%20reduces%20financial%20losses

[5] Diligent. (2025). AI in enterprise risk management (ERM): Transforming risk intelligence for strategic advantage. https: //www.diligent.com/resources/blog/ai-in-enterprise-risk-management

[6] Committee of Sponsoring Organizations of the Treadway Commission (COSO). (2017). Enterprise Risk Management: Integrating with Strategy and Performance. https: //www.coso.org/Documents/2017-COSO-ERM-Integrating-with-Strategy-and-Performance-Executive-Summary.pdf

[7] Power, M. (2021). Riskwork: Essays on the Organizational Life of Risk Management. Oxford University Press.

[8] Tearsheet. (2025). Tearsheet Report: The AI Reality Check. https: //tearsheet.co/artificial-intelligence/tearsheet-report-the-ai-reality-check-q1-2025-edition/#: ~: text=, powered%20execution

[9] IBM. (2020). How AI is transforming supply chain management. https: //www.ibm.com/watson/supply-chain/

[10] 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. https: //doi.org/10.1109/COMST.2015.2494502

[11] Darktrace. (2017). Darktrace AI Stops Ransomware Attack at Manufacturing Firm. https: //www.darktrace.com/en/resources/

[12] IBM. (2025). What is explainable AI? https: //www.ibm.com/think/topics/explainable-ai

[13] Prosci. (2025). AI Adoption: Driving Change With a People-First Approach. https: //www.prosci.com/blog/ai-adoption

[14] Di Palma, G., Scendoni, R., Tambone, V., Alloni, R., & De Micco, F. (2025). Integrating enterprise risk management to address AI-related risks in healthcare: Strategies for effective risk mitigation and implementation. Journal of healthcare risk management : the journal of the American Society for Healthcare Risk Management, 44(4), 25–33.

[15] Das, I. (2025). AI and Data Privacy: Mitigating Risks in the Age of Generative AI Tools. https: //blog.qualys.com/product-tech/2025/02/07/ai-and-data-privacy-mitigating-risks-in-the-age-of-generative-ai-tools

[16] Mccabe, M. (2025). AI in Cybersecurity: Navigating GDPR, Privacy Laws, and Risk Management. https: //www.zscaler.com/blogs/product-insights/ai-cybersecurity-navigating-gdpr-privacy-laws-and-risk-management

[17] Monnette, J., & Chaudhary, A. (2025). Why Do 80% of AI Pilots Fail to Scale? Unpacking the Top Enterprise AI Deployment Challenges. https: //www.epam.com/insights/ai/blogs/enterprise-ai-deployment-challenges