

Agentic AI for IT and Beyond: A Qualitative Analysis of Capabilities, Challenges, and Governance

Hesham Allam ^{*,1}, Juan Dempere²

¹ Center of Information and Communication Science (CICS), College of Communication, Information & Media, Ball State University, Muncie, IN, 47304, hesham.allam@bsu.edu

² Business Division, Higher Colleges of Technology, RAK, UAE; jdempere@hct.ac.ae

* Correspondence: hesham.allam@bsu.edu

Abstract: Agentic AI represents a leap forward in AI, characterized by autonomous decision-making, adaptive reasoning, and innovative collaboration in dynamic environments. In their shift away from mere automation towards reflective, goal-oriented behavior, these promises are significant: in IT operations, real-time analytics, strategic decision-making, and more. Nevertheless, and notwithstanding its increasing importance in industry, there is no coherent framework within the academic literature that captures the technological, ethical, and governance aspects of Agentic AI. This study employs a qualitative approach, incorporating thematic analysis and comparative case studies, to interpret the results from academic sources, industrial documents, and regulatory publications from 2023 and 2024. The paper integrates technical with interdisciplinary literature and considers four key areas: (1) the functional architecture and mechanisms of Agentic AI, (2) operational value via AIOps platforms including Moogsoft and Dyna-trace, (3) evolving risks such as bias, data abuse, and autonomy misalignment, and (4) regulatory and ethical lacunae in existing oversight statutes. Furthermore, the work reveals recurring themes, including explainability, human-AI collaboration, and fairness, which are essential for the design and deployment of these systems in the future. The work surfaces recurring themes, such as explainability, human-AI partnership, and fairness, that are crucial to the way these systems are designed and used in the future. We are still at the phase of approximate common knowledge in AI. To address this and other pressing matters, the paper advocates for a novel methodology, called Agentic AI in-the-making, that centers on an "eye-on-eye" interaction between human and AI agencies. By merging theoretical models with practical instances, the paper establishes a holistic frame for deploying and constraining potential Agentic AI. It also provides an initial slate of recommendations to policymakers, innovators, and industry leaders on how to encourage responsible innovation that focuses on transparency, accountability, and interdisciplinary collaboration in the development of new intelligent systems.

Academic Editor(s): Zahoor Khan

Received: 15 March 2025

Revised: 29 April 2025

Accepted: 6 June 2025

Published: date: 10 June 2025

Citation: Allam, Alomar,

Dempere(2025). Agentic AI for IT and Beyond: A Qualitative Analysis of Capabilities, Challenges, and Governance

AI Business Review, 1(1), x.

[https://doi.org/10.64044/aibr1\(tobe-assigned\)](https://doi.org/10.64044/aibr1(tobe-assigned))

Copyright: © 2025 by the authors.

Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

Keywords: Agentic AI, autonomous systems, predictive analytics, AIOps, multi-agent systems, ethical governance, explainable AI, IT operations, AI accountability

1. Introduction

The field of artificial intelligence (AI) is experiencing a substantial transformation with the advent of Agentic AI—innovative systems engineered not just for task automation but also for the capacity to make independent decisions, perform adaptive reasoning, and interact collaboratively with other intelligent systems. Unlike traditional rule-based or supervised machine learning systems, agentic AI systems operate autonomously, continually learning from their environment, making informed decisions, and utilizing tools and resources that extend beyond their original programming (Allam, 2025a; Allam, 2025b; Allam et al., 2025; Portugal et al., 2024). This evolution positions Agentic AI as a powerful enabler of operational transformation across various sectors, particularly in IT operations, cybersecurity, healthcare, and intelligent automation (Sapkota et al., 2025; Hosseini & Seilani, 2025).

Agentic AI systems exhibit capabilities such as reflection, planning, and multi-agent coordination. These features enable them to solve complex problems that traditionally require significant human oversight. As technology matures, the deployment of Agentic AI is expected to scale rapidly, offering benefits such as increased efficiency, enhanced predictive accuracy, and improved adaptive system performance (Archarya et al., 2025; Shavit et al., 2023). However, this advancement also introduces a new wave of ethical, regulatory, and technical challenges, especially concerning data security, bias, transparency, and the erosion of human oversight.

Given these dual realities, the purpose of this paper is threefold. First, it aims to establish a conceptual and functional understanding of Agentic AI by synthesizing key features and frameworks that define its operation. Second, it explores the practical applications and benefits of Agentic AI in IT ecosystems, particularly within AIOps platforms that rely on autonomy and adaptability. Third, it critically examines the risks, vulnerabilities, and governance gaps associated with deploying such systems in high-stakes environments. By addressing these objectives, this paper contributes to a more structured and responsible discourse on the development, deployment, and regulation of Agentic AI systems.

2. Method

This study employs a qualitative approach and is grounded in an extensive review of the literature. The objective was to investigate the evolving concepts of Agentic AI and to distil and narrate the significant technology and ethical trends. The review was limited to papers from 2023 and 2024 to ensure relevance to the latest developments and discussions.

Reputable quality journal papers, conference publications, and preprints were retrieved from IEEE Xplore, ScienceDirect, SpringerLink, and arXiv. We conducted keyword searches for composite terms, such as “Agentic AI,” “autonomous system,” “multi-agent collaboration,” “AIOps,” “ethical AI,” and “AI governance.” Articles were included that centered on real-world applications, theoretical models, and considerations involving ethics or regulations.

To ensure diversity and depth, both technical and interdisciplinary sources were included. Technical papers helped outline the architecture and functionality of Agentic AI systems, while papers from social sciences and policy studies were used to capture ethical considerations and governance challenges. Industry white papers and policy documents were also incorporated to illustrate the practical application and challenges faced in enterprise environments.

Thematic coding was used to systematically analyze insights into four themes: (1) essential capabilities and mechanisms in Agentic AI, (2) operational advantage in IT settings, (3) vulnerabilities and risks, and (4) policy and governance solutions. This methodology followed the phased process suggested by Braun and Clarke (2006), which provides guidelines for analyzing and reporting qualitative patterns. Common themes, including autonomy, explainability, human-AI collaboration, bias detection, and regulatory gaps, were compared and contrasted to crystallize convergences and divergences in the sources. Table 5: Summary Table: Case (study) design method. Due to the nature of the research question and to provide a deeper analysis, a secondary comparative case study analysis was employed, following the methodological steps outlined by Eisenhardt (1989) and Yin (2018). This rendering was empirically substantiated with concrete use cases from AIOps platforms (e.g., Moogsoft, Dynatrace) that illustrate how claims about operational efficiency, anomaly detection, and the enhancement of autonomous decision-making in Agentic AI systems are justified by real-world applications.

We safeguarded methodological quality through the triangulation of information, integrating findings from the academic literature, industry reports, and regulatory documentation. This methodology adhered to the triangulation rules described by Denzin (1978) and later refined by Patton (2015), ensuring a balanced and robust synthesis of the results. Moreover, the incorporation of various viewpoints, as outlined in Creswell and Poth (2018), increased the trustworthiness of the study's findings and recommendations.

3. Definition and Framework

3.1 Definition of Agentic AI

Agentic AI is a new form of AI based on the paradigm of autonomous and collaborative agents. This is particularly powerful in domains such as recommender systems and data analytics, as well as IT operations, where AI systems can act autonomously and human-in-the-loop supervision is limited. The Background will discuss the basics of the building

blocks and operational implications of Agentic AI (Sapkota et al., 2025; Portugal et al., 2024).

3.2 Framework of Agentic AI

3.2.1 Autonomous Decision-Making

Agentic AI refers to AI systems that make decisions based on assessing real-time data independently, without relying on external guidance. This is also vital in IT operations, where predictive algorithms lower downtime and streamline processes (Sivakumar, 2024). Their ability to self-govern reduces human dependence and enhances operational agility.

3.2.2 Multi-Agent Collaboration

These systems are often formed by multiple AI agents collaborating to solve complex tasks. They pool resources, conform to one another's actions, and together improve decisional accuracy. Such models are beneficial in Recommender Systems and Cloud Orchestration (Portugal et al., 2024).

3.2.3 Reflection and Strategic Planning

The agentic AI includes reflection mechanisms, enabling the system to reflect on its past actions, correct mistakes, and improve its future performance. Planning modules facilitate breaking high-level goals into smaller subtasks, mirroring human executive functions (Chawla et al., 2024).

3.2.4 Tool Use and External Integration

Agentic agents extend their capabilities by integrating with APIs, datasets, and computational tools. This external access empowers them to perform simulations, validate assumptions, and refine outputs (Chawla et al., 2024).

4. Operational Benefits in IT Ecosystems

4.1 Predictive Analytics for System Resilience

AI is an agentic platform driving resilience through predictive analytics. Such systems leverage historical large-scale data and real-time operational patterns to predict potential system failures in the form of signals, doing so before the failures occur. By monitoring early warning signs, Agentic AI enables companies to act proactively, minimizing potential outages. This "look-ahead" approach avoids interruption, improves service availability, and ensures IT service provision is more sustained (Cheng et al., 2023).

4.2 Automation through AIOps Platforms

Reduction of AIOps (Artificial Intelligence for IT Operations). In AIOps environments, Agentic AI offers comprehensive automation for various operational activities. It auto-detects incidents, root causes, and resolution playbooks. With much of the operational legwork taken off the plates of IT practitioners, they can now focus on innovation rather than infrastructure. The first enabler is increased IT efficiency, which enables a lean and flexible organizational structure (Parab, 2024; Joseph, 2023).

4.3 Enhanced Decision Intelligence

Intelligent machines may also facilitate better and faster decision-making by routinely processing data and conflicting interpretations. Such systems employ sophisticated analytical models to interpret the situation, predict the outcome, and recommend the most effective courses of action. This richer decision intelligence supports businesses that need to operate in dynamic environments with a need for quick response. Competitive advantage is achieved through both a reduction in response time lag and adaptability, as decisions are based on near real-time insights (Parab, 2024). Organizations that can make decisions in closer alignment with fresher data can obtain a competitive advantage.

4.4 Example Use Case: Moogsoft and Dynatrace

Real-world solutions, such as Moogsoft and Dynatrace, showcase the operational capabilities of Agentic AI when implemented... Agentic AI products feature intelligent agent-based capabilities, including anomaly detection, issue correlation, and automatic remediation (Moogsoft, 2023; Dynatrace, 2023; Dynatrace Docs, 2022).

Moogsoft utilizes AI-driven anomaly detection to correlate and group alerts together, minimizing noise to identify the “needle in the haystack” that Site Reliability Engineers (SREs) and DevOps teams worry about in the flood of data. It is smart, threshold-driven contextual processors can automatically detect anomalies, reduce noise, and enrich incidents with relevant data, correlating those incidents to identify likely root causes, meaning it is faster to detect (MTTD) and resolve (MTTR) incidents (Moogsoft, 2023; DrDroid, 2024). Automated Incident Response: The platform also enables automated incident response by coordinating workflows across third-party tooling (Moogsoft, 2023), such as routing tickets or executing remediation scripts.

Leveraging its unique Davis® AI engine, Dynatrace does this all day, every day, for thousands of customer environments – and it does it automatically. Davis consumes metrics, logs, traces, and topology data in real-time, facilitating a level of automated root-cause analysis that reduces alert storms and directs remediation efforts. It performs automatic anomaly prioritization and correlation across dependencies, all the way to the code level, resulting in frequent MTTR reductions of up to 90% (Dynatrace, 2023; Dynatrace Docs, 2022; OpenObserve, 2025).

When they work together, Moogsoft and Dynatrace can demonstrate how Agentic AI helps reshape IT operations, providing an environment of continuous infrastructure observability, autonomous anomaly detection, accurate root-cause resolution, and proactive incident prevention. This demonstrates the statement that Agentic AI systems profoundly increase the efficiency of processes and decision-making in complex IT systems.

5. Risks, Vulnerabilities, and Ethical Implications

5.1 Over-Reliance and Skill Erosion

In any case, as bodies rely on Agentic AI systems, the erosion of human expertise is likely to become an increasingly concerning issue. Long-term use of AI to make decisions and fix problems may dull the analytical skills of IT personnel. This loss of skill is a potential long-term risk, as it could limit the capacity of initiatives to react appropriately when human oversight is required or AI systems do not perform as anticipated (Sivakumar, 2024).

5.2 Data Privacy and Security Threats

As Agentic AI systems autonomously access a sensitive set of enterprise data, the risk of data leaks and cyber manipulation grows. These systems are networked and interconnected with other devices, and they may receive and transmit information with remote data sources; hence, they are exposed to security threats. Attacks can be launched over these channels, or the system's learning models can be fooled, leading to unintended or harmful behavior. Therefore, strong cybersecurity policies and safe data management should be incorporated from the beginning (Khan et al., 2024).

5.3 Objective Misalignment and Accountability Gaps

Robotic agents are often liable to misunderstanding user intentions or functioning with contradicting objectives. Without accurate alignment mechanisms, Agentic AI may pursue courses of action that, while logically accurate, are inconsistent with ethical, social, or business norms. Also, blame is harder to find when things go wrong, or when there is damage, particularly when humans are doing less watching. This tension strains domestic models of AI accountability and responsibility to their limit (Clatterbuck et al., 2024).

5.4 Bias and Discriminatory Outputs

Agentive AI systems, when asked to learn from biased datasets, can unintentionally reflect or magnify pre-existing inequities. If these biases are then programmed into artificially intelligent systems used for hiring, finance, healthcare, or criminal justice, we could end up with profoundly unjust results. Regular audits, fair-aware training algorithms, and transparent data curation are essential to better addressing and mitigating these risks (Sutherns & Fanta, 2024).

5.5 Regulatory Oversight and Governance Gaps

AI technologies are being deployed at a breakneck pace, yet their regulatory frameworks, especially when it comes to autonomous and agentic systems, remain lacking. In the absence of clear standards for algorithmic transparency, explainability, and accountability, governance of Agentic AI is erratic. Adaptive laws, such as accelerators for technological innovation, pose serious challenges to existing legal instruments designed to safeguard public interest. Industry and governments must therefore pool resources and collaborate to create adaptive laws (Kappel, 2024).

6. Discussion

Agentic AI represents a pivotal milestone in the evolution of AI. Combining autonomy, adaptation, and collaboration, it fills up the space between reactive automation and proactive intelligence. The reading itself clearly shows that when effectively applied, Agentic AI can deliver significant operational improvements, particularly in dynamic areas such as IT operations and real-time analytics.

However, its implementation prompts important theoretical and practical considerations. Transitions from human-in-the-loop to human-on-the-loop setups necessitate new oversight mechanisms that ensure the system remains controllable, interpretable, and continually improves toward human goals. The bar for 'behavior validation'-traditional testing is too low; and new bar(s) should be set for evaluating the safety and reliability of ongoing agents.

Another significant takeaway from the study is the necessity of explainability. The more complex AI systems become, including those with agentic capabilities, the more users (and decision-makers and operators, in particular) require assistance in understanding the reasoning behind their decisions and the variables that affect them. Explainable AI (XAI) is not only a technical necessity but also a governance requirement to preserve trust and accountability in autonomous systems.

Adding to this is the problem that systems with agency also have autonomy, and in particular, can learn from biased data, thereby increasing the risk of systemic bias and ethical misalignment. Fairness-aware design and continual auditing need to be a de facto practice. Regulations like those in the EU AI Act and the U.S. NIST AI Risk Management Framework must be modified to be specifically applicable to Agentic AI.

Interdisciplinarity is also key. Agentic AI is too important to be the sole concern of scientists or policy-makers alone.' Technologists, ethicists, lawyers, and sector experts must all contribute to shaping the design and governance of Agentic AI. These would be truly interdisciplinary efforts to co-develop guidelines that are technically sound, socially responsible, and responsive to global norms and standards.

The conversation also exposes holes in empirical knowledge. There is work that seems much of it is still theoretical and does not yet exist in the real world. Longitudinal empirical studies in enterprises, healthcare, finance, and the public sector should be prioritized to assess the practical impact (positive or negative) of Agentic AI systems in high-stakes contexts.

Finally, the only way forward is to design Agentic AI that is intelligent and autonomous, but also transparent, inclusive, and safe. When these values are embedded in our machines, they can help ensure that the technologies we build serve society, rather than becoming its unwitting victims.

7. Conclusion and Future Directions

Enabling AI to act as an independent agent has the potential to revolutionize AI with autonomy, adaptability, and intelligent cooperation. However, it can only be safely done if its many risks are responsibly addressed in advance. By establishing an equilibrium that

supports innovation and ethics, interests can then utilise Agentic AI for advancement, without sacrificing our social values or security.

Future Research Directions

- Development of explainable Agentic AI models (X-Agentic AI)
- Resilience testing frameworks for autonomous systems
- Human-agent collaboration and shared decision-making
- Cross-cultural AI ethics and inclusivity
- Legal models for AI accountability and redress

Acknowledgement

The authors acknowledge that all research, analysis, and writing were conducted independently without external assistance. Grammarly, an AI-powered writing assistant, was used solely during the revision stage to support minor improvements in grammar, clarity, and readability, without contributing to the intellectual content or structure of the manuscript.

References

- Acharya, D. B., Kuppan, K., & Divya, B. (2025). Agentic AI: Autonomous Intelligence for Complex Goals–A Comprehensive Survey. *IEEE Access*.
- Allam, H. (2025a). Real-Time Competency Mapping of Student Reflections Using Rule-Based AI and Visual Analytics. *The Artificial Intelligence Business Review*, 1(1), 15–15.:
<https://theaibr.com/index.php/aibr/article/view/NLP>
- Allam, H. (2025b). Prescribing the Future: The Role of Artificial Intelligence in Pharmacy. *Information*, 16(2), 131.
<https://doi.org/10.3390/info16020131>
- Allam, H. M., Gyamfi, B., & AlOmar, B. (2025). Sustainable Innovation: Harnessing AI and Living Intelligence to Transform Higher Education. *Education Sciences*, 15(4), 398. <https://doi.org/10.3390/educsci15040398>
- Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101. <https://doi.org/10.1191/1478088706qp063oa>
- Chawla, C., Chatterjee, S. G., Gadadinni, S. S., Verma, P., & Banerjee, S. (2024). Agentic AI: The building blocks of sophisticated AI business applications. <https://doi.org/10.69554/xehz1946>
- Cheng, Q., Sahoo, D., Saha, A., Yang, W., Liu, C., Woo, G., Singh, M., & Hoi, S. C. H. (2023). AI for IT Operations (AIOps) on Cloud Platforms: Reviews, Opportunities and Challenges. <https://doi.org/10.48550/arXiv.2304.04661>

- Clatterbuck, H., Castro, C., & Morán, A. M. (2024). Risk Alignment in Agentic AI Systems. <https://doi.org/10.48550/arxiv.2410.01927>
- Creswell, J. W., & Poth, C. N. (2018). *Qualitative inquiry and research design: Choosing among five approaches* (4th ed.). SAGE Publications.
- Denzin, N. K. (1978). *The research act: A theoretical introduction to sociological methods* (2nd ed.). McGraw-Hill.
- DrDroid. (2024, April 4). Moogsoft anomaly detection and correlation capabilities. DrDroid.io. <https://drdroid.io/moogsoft-anomaly-detection>. Accessed: June 11, 2025.
- Dynatrace. (2023). Automatic root-cause analysis. Dynatrace. <https://www.dynatrace.com/news/blog/automatic-root-cause-analysis-ai/>. Accessed: June 11, 2025.
- Dynatrace Docs. (2022). Root cause analysis concepts. Dynatrace Documentation. <https://docs.dynatrace.com/docs/shortlink/root-cause-analysis>. Accessed: June 11, 2025.
- Eisenhardt, K. M. (1989). Building theories from case study research. *Academy of Management Review*, 14(4), 532–550. <https://doi.org/10.5465/amr.1989.4308385>
- Hosseini, S., & Seilani, H. (2025). The role of agentic AI in shaping a bright future: A systematic review. *Array*, 100399.
- Joseph, C. (2023). Artificial Intelligence for IT Operations (AIOps). <https://doi.org/10.58532/v2bs17p3ch1>
- Kappel, G. (2024). Mesterséges intelligencia alapú döntéstámogató és döntéshozó rendszerek kockázatai a vállalatok vezetői szintű döntéshozatalában. *Pro Futuro*, 14(1). <https://doi.org/10.26521/profuturo/2024/1/15166>
- Khan, R., Sarkar, S., Mahata, S.K., & Jose, E. (2024). Security Threats in Agentic AI System. <https://doi.org/10.48550/arxiv.2410.14728>
- Moogsoft. (2023). AIOps platform features. Moogsoft. <https://www.moogsoft.com/platform/aiops-platform>. Accessed: June 11, 2025.
- Moogsoft. (2023). Anomaly detection. Moogsoft. <https://www.moogsoft.com/features/anomaly-detection>. Accessed: June 11, 2025.
- OpenObserve. (2025). How Dynatrace works and what it does? OpenObserve.ai. <https://openobserve.ai/blog/how-dynatrace-works>. Accessed: June 11, 2025.
- Parab, G. U. (2024). Agentic AI in Data Analytics: Transforming Autonomous Insights and Decision-Making. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*. <https://doi.org/10.32628/cseit241061220>
- Patton, M. Q. (2015). *Qualitative research & evaluation methods: Integrating theory and practice* (4th ed.). SAGE Publications.
- Portugal, I., Alencar, P., & Cowan, D. (2024). An Agentic AI-based Multi-Agent Framework for Recommender Systems. <https://doi.org/10.1109/bigdata62323.2024.10825765>
- Sivakumar, S. (2024). Agentic AI in Predictive AIOps: Enhancing IT Autonomy and Performance. *International Journal of Scientific Research and Management*. <https://doi.org/10.18535/ijsrcm/v12i11.ec01>
- Sapkota, R., Roumeliotis, K. I., & Karkee, M. (2025). Ai agents vs. agentic ai: A conceptual taxonomy, applications and challenge. *arXiv preprint arXiv:2505.10468*.

- Shavit, Y., Agarwal, S., Brundage, M., Adler, S., O’Keefe, C., Campbell, R., ... & Robinson, D. G. (2023). Practices for governing agentic AI systems. *Research Paper, OpenAI*.
- Sutherns, J., & Fanta, G. B. (2024). The implications of integrating artificial intelligence into data-driven decision-making. *South African Journal of Industrial Engineering*. <https://doi.org/10.7166/35-3-3096>
- Veluru, S. P. (2021). Leveraging AI and ML for Automated Incident Resolution in Cloud Infrastructure. *International Journal of Artificial Intelligence, Data Science, and Machine Learning*, 2(2), 51-61.
- Yin, R. K. (2018). *Case study research and applications: Design and methods* (6th ed.). SAGE Publications.