




LIGHTS - Technology

Agentic AI: the next frontier in autonomous enterprise systems

ESCP Impact Paper No.2025-32-EN

Daniel VÖLKER
Valantic

Marc OBERHAUSER
ESCP Business School

[**Agentic AI: the next frontier in autonomous enterprise systems**

Daniel Völker*
Marc Oberhauser**

ESCP Business School

Abstract

Agentic AI refers to artificial intelligence systems capable of autonomous decision-making, contextual understanding, and long-term planning. Unlike traditional or generative AI, which rely on human prompts, agentic systems operate proactively, managing complex, multi-step tasks across dynamic environments. This ESCP Impact Paper explores the technical foundations and practical applications of agentic AI, offering a clear contrast with generative models and highlighting its business potential. Drawing on IBM's watsonx architecture, the paper outlines the core components required for agentic systems, including orchestration, memory, cognition, and governance. Three feasible and impactful use cases are examined in detail: intelligent copilots, customer service automation, and end-to-end process automation. Together, these examples demonstrate how agentic AI is already reshaping business processes, improving decision quality, and enabling new forms of human-AI collaboration. The paper concludes with strategic insights for organizations aiming to harness this emerging capability.

Keywords: agentic AI, human-AI collaboration, AI architecture

*Head of AI Solutions Department at valantic

**Assistant Professor of Management, ESCP Business School

ESCP Impact Papers are in draft form. This paper is circulated for the purposes of comment and discussion only. Hence, it does not preclude simultaneous or subsequent publication elsewhere. ESCP Impact Papers are not refereed. The form and content of papers are the responsibility of individual authors. ESCP Business School does not bear any responsibility for the views expressed in the articles. Copyright for the paper is held by the individual authors.

Agentic AI: the next frontier in autonomous enterprise systems

Introduction

In May 2025, China unveiled the world's first fully AI-operated hospital, "Agent Hospital," developed by researchers at Tsinghua University (Ramirez, 2025). Staffed by 14 AI doctors and 4 AI nurses, the facility is capable of managing up to 3,000 patients per day without human intervention. These AI agents achieved a 93.06% success rate on the MedQA dataset, demonstrating their ability to handle diagnosis, treatment, and follow-up care autonomously. More than a novelty, this deployment marks the first institutional-scale application of agentic AI, a new class of artificial intelligence systems capable of acting autonomously and purposefully in complex environments.

Since the launch of ChatGPT in late 2022, generative AI has dominated headlines and has been enormously celebrated for its ability to create human-like content based on prompts. However, despite generative AI solutions being quite powerful, generative AI remains fundamentally reactive, only responding to instructions; it cannot independently initiate tasks, plan workflows, or adapt over time. In contrast, agentic AI represents a paradigm shift. Described by McKinsey as "the new frontier of generative AI," agentic systems combine goal-setting, decision-making, and learning capabilities that allow them to operate with a high degree of independence (Lareina Yee et al., 2024).

This transition from passive tools to proactive agents has profound implications. By automating not only content creation but also reasoning, planning, and execution, agentic AI has the potential to reshape how businesses operate. These systems can free up human workers from routine or procedural work, allowing them to focus on creativity, strategic thinking, and innovation. More importantly, agentic AI unlocks new possibilities for organizational design and business models, enabling the delegation of complex, multi-step tasks to autonomous digital collaborators (Roque, 2024).

This ESCP Impact Paper explores the transformative potential of agentic AI. We will begin by examining its foundational capabilities and how they differ from generative AI. We will take a detailed look at the technical architectures required to make agentic AI systems operational, followed by high-impact use cases across industries. The goal of this paper is to equip business leaders, technologists, and policymakers with a clear understanding of how to harness this next wave of intelligent automation.

Agentic AI vs. Generative AI – capabilities and distinctions

While generative AI has proven highly effective at producing content based on patterns in data, it lacks the ability to act independently or reason over time. Generative AI operates within narrow parameters, executing single-turn tasks in response to human input. In brief, generative AI is an assistant, not an agent.

In contrast, agentic AI represents a more advanced and transformative class of artificial intelligence. Rather than merely increasing task efficiency, agentic AI has the potential to transform entire workflows, address system-level challenges, and enable new business models based on continuous human-AI collaboration. **Agentic AI is characterized by three foundational capabilities that distinguish it from earlier forms of AI, namely (1) autonomous decision-making, (2) contextual awareness, and (3) long-term strategic**

planning. Together, these features enable agentic AI systems to function as intelligent collaborators capable of managing complexity and change (Bornet et al., 2025).

Autonomy in artificial intelligence refers to the system's ability to perform tasks, make decisions, and adapt to changes without continuous human intervention. Agentic AI systems are designed to independently perceive their environment, evaluate options, and execute actions aligned with predefined goals. Agentic AI systems can exhibit varying degrees of autonomy. They can be partially or fully autonomous. Partial autonomy is given when the AI system performs specific tasks independently, but may require human oversight for decision-making in complex or unforeseen situations. Full autonomy is reached once the AI system operates entirely independently, managing all aspects of task execution, decision-making, and adaptation without human intervention. The level of autonomy depends on the specific application and the desired balance between automation and human control.

Contextual awareness refers to agentic AI systems' ability to interpret and respond to contextual information, enabling them to adjust their behavior based on the current environment and operational conditions. This includes understanding user intent, environmental variables, and temporal factors that may influence actions and decision-making processes.

Long-term strategic planning captures Agentic AI systems' design to pursue long-term objectives by decomposing them into manageable sub-tasks, planning execution strategies, and monitoring progress over time. This goal-oriented approach enables them to handle complex, multi-step processes that require sustained attention and adaptability. Those planning capabilities allow agentic AI to handle challenges, evaluate alternative courses, and adjust strategies accordingly.

Table 1 summarizes the differences between generative AI and agentic AI along key functions.

Table 1: Comparing Generative AI and Agentic AI		
Function	Generative AI	Agentic AI
Core Function	Content generation (text, images, code)	Autonomous goal pursuit and task execution
Trigger	Responds to human prompts	Initiates action based on goals and context
Decision-making	Prompt-dependent, non-strategic	Self-directed, strategic, and situational
Contextual awareness	Limited to input and short-term memory	Integrates user intent, history, and environment
Planning capability	Single-turn tasks	Multi-step planning and dynamic adjustment
Typical use cases	Drafting content, summarizing data	Managing workflows, virtual assistants, autonomous research, process automation
Human interaction	Required for each step	Optional—depending on autonomy level and application
Business impact	Increases specific task efficiency	Automates complete goals, with subtasks which enables operational transformation and new business models

How agentic AI works

The functionality of Agentic AI can be understood best through a **modular architecture composed of three interrelated components, namely (1) perception, (2) cognition, and (3) action** (see Figure 1). Each module represents a distinct AI agent that can be specialized to perform a set of tasks, forming an integrated system capable of autonomous, context-aware behavior.

In the **perception phase**, the AI agent acts as the sensory system, gathering data from various sources such as cameras, microphones, and digital inputs. It processes this information to extract meaningful features and information, enabling the AI agent to understand its environment.

In the **cognition phase**, the AI agent defines its goals. It plans actions in order to achieve the desired goals, while considering its environment, short-term memory, and relevant knowledge from databases. Based on the developed plans and the environmental context, this component anticipates the outcomes of possible actions and selects the most appropriate one to execute, aligning with the AI agent's objectives.

Finally, in the **action phase**, the AI agent executes the chosen actions, which could involve physical movements in robotics or virtual operations like data processing or communication. The AI agent continuously evaluates and learns from the outcomes of its actions, thereby refining its knowledge and improving future decision-making processes.




Figure 1: Agentic AI Architecture. Authors' own illustration inspired by Sharma (2025)

Technical architectures

Agentic AI is not an out-of-the-box solution. Building an agentic AI system requires a seamless integration of multiple technologies that enable perception, reasoning, and action. These components must work together dynamically and autonomously.

Given the complexity of this setup, several technology providers have introduced frameworks to support the implementation of agentic AI capabilities. To illustrate how these components can interact in practice, this paper uses IBM's watsonx technology stack

(IBM, 2025) as an exemplary case, purely for illustrative purposes and without implying exclusivity or endorsement.

A robust agentic AI architecture typically comprises four functional categories: (1) orchestration and integration, (2) memory and retrieval, (3) cognitive processing via foundation models, and (4) governance and oversight (see Table 2).

Table 2: Overview of the IBM Technology Stack

Platform	Function in Agentic AI
watsonx Orchestrate	Orchestration and Integration: Autonomous task management and workflow execution
watsonx.data	Memory and Retrieval: Insight extraction, document and knowledge retrieval
watsonx.ai	Cognitive engine: planning, reasoning, interaction via foundation models
watsonx.governance	Governance: Oversight, traceability, fairness, and compliance

Orchestration and Integration: In an Agentic AI context, this function acts as the orchestration and execution layer. It executes the agents’ decisions, manages goal-oriented processes based on business logic, and connects with enterprise systems to retrieve or update data. It also enables interaction between AI agents and human users via chat, voice, email, or custom interfaces.

IBM’s *watsonx Orchestrate* provides one implementation of this layer. It is a low-code platform that helps sequence multi-step workflows and trigger tasks through natural language or APIs, allowing AI decisions to be integrated into enterprise operations.

Memory and Retrieval: Within an agentic AI system, this layer functions as the memory and retrieval mechanism. It enables agents to access structured and unstructured information across environments, grounding their decisions in enterprise data. Similar to how humans consult documents or databases to make informed choices, agents use this layer to "look things up" in real time.

In the IBM ecosystem, *watsonx.data* serves this purpose. It is a hybrid, governed data lakehouse platform that supports various formats and includes vector database capabilities for semantic document search and retrieval.

Cognitive Reasoning: In an agentic AI system, this layer serves as the cognitive engine, essentially the "brain" of the agent. It interprets user input, defines goals, and generates the sequence of actions needed to accomplish tasks. By integrating reasoning, planning, and interaction, this component allows agents to move beyond reactive responses and engage in proactive, goal-directed behavior.

IBM’s *watsonx.ai* provides this functionality by supporting large language models, multimodal foundation models, and machine learning systems for tasks such as intent recognition, strategic planning, and contextual adaptation.

Governance and Oversight: In an agentic AI system, governance ensures that autonomous agents operate transparently, ethically, and in compliance with organizational policies and external regulations. This layer provides traceability, fairness auditing, and behavior control mechanisms to ensure agents remain aligned with business objectives, even as they evolve.

watsonx.governance is IBM's implementation of this layer. It offers oversight tools that support responsible AI deployment by enabling monitoring, accountability, and lifecycle management for autonomous systems.

Feasible and impactful use cases for agentic AI in business

Agentic AI represents a new generation of intelligent systems capable of autonomous action, context understanding, and pursuing long-term goals, without the need for continuous human input. Unlike traditional automation or generative models limited to single prompts, agentic AI systems are designed to handle dynamic, multi-step tasks with strategic intent. This unlocks a wide range of business applications that go beyond reactive tools toward proactive, context-aware intelligence.

Among the most feasible and impactful use cases that are already adopted by many companies are (in order of feasibility) **(1) intelligent copilots, (2) customer service automation, and (3) end-to-end process automation.**

Use case 1: intelligent copilots / personal assistants

One of the most familiar and immediately implementable applications of agentic AI is the use of AI copilots or personal assistants in the workplace. These agents support employees in content creation (text, slides, code), summarization, data analysis, and communication. As a result, they significantly reduce cognitive load, help employees make better decisions, and free up time for higher-value work. What distinguishes agentic copilots from conventional productivity tools is their ability to operate across applications and maintain context over time. Such copilots are especially attractive because they build on existing workplace behaviors and digital habits, making adoption smoother and the value quickly tangible. Moreover, these agents can be tailored to individual roles and organizational processes, enabling personalization at scale. For example, a legal assistant agent could autonomously summarize a lengthy contract, draft follow-up questions for clarification, send them to the counterparty, and track responses, all with minimal human intervention. In software development, copilots can manage not just code completion but also bug tracking, testing workflows, and project handoffs.

Citigroup has initiated pilot programs integrating agentic AI into its operations. Their AI suite includes tools like Citi Assist and Citi Stylus, which assist employees in document analysis and real-time chat, aiming to enhance productivity and decision-making across various departments (Bautzer, 2024).

Use case 2: customer service automation

A second highly impactful use case is the transformation of customer service and support. Traditional chatbots and support systems are often limited to predefined responses or rigid scripts, making them unsuitable for handling complex customer issues. Agentic AI changes this by introducing intelligent agents that can understand intent, access relevant systems or data sources, and execute appropriate actions. By reducing friction and improving resolution quality, agentic customer service agents can meaningfully improve customer satisfaction and loyalty. At the same time, they significantly reduce operational costs by minimizing manual escalations and increasing first-contact resolution rates. Over time, these agents can also detect recurring issues and feed insights back into product or service design, closing the loop between operations and innovation. For instance, an AI agent could identify a customer's issue by parsing support tickets and backend logs, cross-checking

account history, issuing a refund, rescheduling a delivery, and sending a confirmation email, without involving a human representative. In multilingual environments, the same agent could dynamically switch languages or adapt tone and response strategies based on customer sentiment.

Ask JADA, developed by MomentPort Inc., is an AI customer service agent tested in a real-world telecom setting. In a controlled trial, it outperformed human representatives at Rogers Communications, answering complex technical questions with 100% accuracy and completing calls in half the time (Doering & Tiu, 2024). Trained only on public data, JADA demonstrated the potential of agentic AI to deliver fast, consistent, and scalable support without human intervention.

Use case 3: end-to-end process automation

A more strategic but equally powerful application is to use agentic AI to automate entire business processes. In contrast to conventional workflow automation, often limited by static rules and brittle logic, agentic systems can adapt to different inputs, unforeseen situations, and evolving business conditions. Consider a scenario in which an AI agent manages the entire procurement cycle: it identifies demand from inventory and usage trends, requests and compares supplier quotes, evaluates risk or ESG factors, secures approvals, and follows through on ordering and tracking. Crucially, if anomalies arise, such as a sudden supplier outage or price spike, the AI agent can autonomously escalate, propose alternatives, or renegotiate terms. This kind of adaptive process automation improves not just efficiency but resilience and decision quality. Unlike rule-based systems that break when exceptions occur, agentic systems can learn from edge cases and improve their reasoning over time. They also reduce dependency on rigid IT integrations by flexibly interfacing with APIs, documents, or even human stakeholders as needed.

For instance, researchers have developed FinRobot, an agentic AI framework for enterprise resource planning in finance. It integrates generative AI with business process modeling and multi-agent orchestration, enabling end-to-end automation of complex tasks such as budget planning, financial reporting, and wire transfer processing (Yang et al., 2025).

Conclusion

Agentic AI represents a fundamental shift in the role of artificial intelligence, from passive tool to autonomous collaborator. Unlike generative AI, which remains confined to content generation, agentic systems can initiate action, plan over time, and adapt to evolving contexts. This capacity unlocks new opportunities for business innovation, process optimization, and organizational redesign. However, agentic AI is not an out-of-the-box solution. Building an agentic AI system requires the integration of multiple technologies, including orchestration platforms, memory and retrieval systems, large language models, and governance frameworks, to function cohesively in real-world environments.

Despite these substantial benefits and the transformative potential of agentic AI, its adoption within companies remains relatively slow due to several notable challenges. First, the technical complexity involved in implementing agentic AI systems is significant. Integrating multiple specialized technologies requires substantial technical expertise and resources, often exceeding the immediate capabilities of many organizations. Second, data protection and privacy concerns represent substantial hurdles, especially under stringent regulations such as the GDPR. Companies must carefully manage the extensive data processing and storage demands inherent in agentic AI to ensure compliance. Third, the

legal landscape surrounding autonomous decision-making technologies remains uncertain, raising liability issues and compliance complexities that businesses hesitate to navigate without clear guidelines. Lastly, end-user uncertainty and resistance further impede adoption. Employees may feel threatened by the autonomy and complexity of agentic AI systems or skeptical regarding their reliability and safety. Addressing these concerns through clear communication, transparency, and robust governance practices will be crucial for accelerating the integration of agentic AI into enterprise environments.

As demonstrated through current use cases in copilots, customer service, and end-to-end automation, agentic AI is not a distant concept but an emerging operational reality. While challenges around implementation and governance remain, the strategic implications for business are clear. Those who understand and embrace agentic systems early stand to gain a lasting competitive advantage in an increasingly automated and complex world.

References

- Bautzer, T. (2024, December 4). "Citigroup rolls out artificial intelligence tools for employees in eight countries". *Reuters*.
<https://www.reuters.com/technology/artificial-intelligence/citigroup-rolls-out-artificial-intelligence-tools-employees-eight-countries-2024-12-04/>
- Bornet, P., Wirtz, J., De Cremer, D., Evergreen, B., Fersht, P., Gohel, R., Khiyara, S., Mullakara, N., & Sund, P. (2025). "Agentic Artificial Intelligence: Harnessing AI Agents to Reinvent Business, Work and Life." Irreplaceable Publishing.
- Doering, T., & Tiu, L. (2024, November 22). "JADA AI Employee Outperforms Humans in Customer Service Study." *MomentPort Inc. - News*.
<https://www.momentport.com/news/ai-outperforms-humans/>
- IBM. (2025). *IBM watsonx*. IBM.
<https://www.ibm.com/watsonx>
- Lareina Yee, Michael Chui, Roger Roberts, & Stephen Xu. (2024, July 24). "Why AI agents are the next frontier of generative AI" . McKinsey.
<https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/why-agents-are-the-next-frontier-of-generative-ai>
- Ramirez, A. (2025, May 3). "China's AI Hospital with 14 Robotic Doctors — The Future of Medicine!" *Engineerine*.
https://engineerine.com/china-launches-ai-hospital/?utm_source=chatgpt.com
- Roque, L. (2024, December 13). "Agentic AI: Building Autonomous Systems from Scratch." *Medium*.
<https://medium.com/data-science/agentic-ai-building-autonomous-systems-from-scratch-8f80b07229ea>
- Sharma, R. (2025, May 12). "Agentic AI Architecture: A Deep Dive". Markovate.
<https://markovate.com/blog/agentic-ai-architecture/>
- Yang, H., Lin, L., She, Y., Liao, X., Wang, J., Zhang, R., Mo, Y., & Wang, C. D. (2025). FinRobot: Generative Business Process AI Agents for Enterprise Resource Planning in Finance. *ArXiv Preprint ArXiv*, 2506.01423.
<http://arxiv.org/abs/2506.01423>