FUJITSU

AI agent innovates

Pushing the boundaries of Generative Tech

Contents

AI agent innovates: Pushing the boundaries of Generative Tech

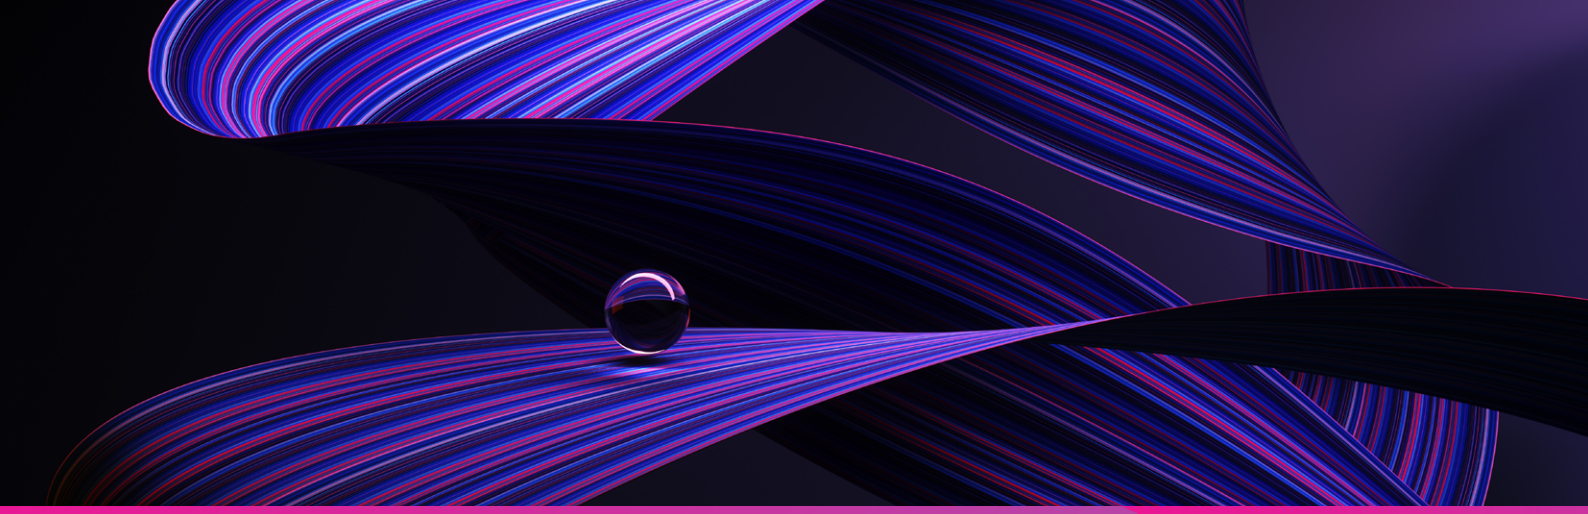

Introduction

The Limits of Generative AI and the Attention of AI Agents

At the end of 2022, generative AI technologies such as ChatGPT and Gemini emerged, showcasing impressive capabilities in text generation, translation, summarization, question answering, and more, all based on large-scale language models (LLMs). These advancements enable the generation of natural, fluent text, the ability to answer complex questions, and the creation of creative content, all of which set generative AI apart from traditional AI. Human society has been captivated by generative AI's natural language processing abilities, versatility, and customizability.

However, generative AI also faces several challenges. Ethical concerns include "hallucinations," where AI generates information that is factually incorrect, as well as the potential for biased or harmful content. Additionally, LLMs consume significant computational resources and energy, and their understanding of logical reasoning and causal relationships remains limited. As a result, they are not yet fully equipped to solve critical problems in fields such as drug discovery, materials science, and physics.

In the industrial sector, generative AI is nearing the end of its learning phase. Moving forward, there is a demand for improvements in quality, cost efficiency, and the ability to handle more complex tasks to facilitate large-scale adoption. Generative AI technology is evolving rapidly, with innovations like the introduction of retrieval-augmented generation (RAG) contributing to quality improvements. Efforts to reduce the costs associated with training and deploying underlying models, as well as to enhance the architecture of LLMs, are also underway.

As generative AI continues to evolve, AI agents[*1] that perform tasks autonomously are rapidly gaining attention, shifting away from passive assistants. Autonomous AI agents integrate generative AI with traditional AI technologies to deliver more advanced and flexible capabilities. This enables AI agents to tackle complex tasks and meet diverse user needs.

This paper explains the mechanisms and innovations behind AI agents, provides use cases and examples from leading companies, and offers insights into AI agents for top management.

*1 In this article, the term 'AI agent system' refers to a system that autonomously makes decisions and carries out tasks or workflows on behalf of users or other systems, leveraging a range of technologies such as large-scale language models (LLMs), inference models, machine learning (ML), reinforcement learning, and knowledge representation. Generally, a collection of AI agents with these capabilities is referred to as 'agentic AI.'

1. The Evolution of Generative AI Models and Future Reasoning Models

Generative AI models (LLMs) are indeed well-versed in language patterns, but their capabilities are largely limited to pattern recognition. Numerous studies have shown that they still have significant limitations in understanding, establishing true causal relationships, and handling complex, multi-step reasoning.[*2] In response to these limitations, leading vendors are developing models with advanced reasoning capabilities.

Notable Efforts by Leading Vendors

For example, DeepMind's Gemini 1.5 Flash, released in May 2024, is a lightweight version but claims to have a highly effective multimodal reasoning ability over vast amounts of information.[*3] Additionally, OpenAI's o1 model, released in September 2024, has entered the domain of complex reasoning, marking a significant advancement in fields such as physics and coding. The o1 model integrates reasoning capabilities directly into the model, making it a reasoning model.[*4] Moreover, in January 2025, DeepSeeK-R1, which made a global impact, was released as another high-capacity reasoning model.[*5]

Approaches to Achieving Reasoning Functionality

Reasoning, of course, comes in many forms.[*6] While language models (LLMs) excel at certain types of reasoning, particularly pattern-based reasoning, there are several promising approaches being explored by leading vendors to enhance reasoning capabilities in language models. Based on these efforts, we believe the following three approaches hold promise:

(1) Enhancing Transformers with a Reasoning Layer

One approach is to add a dedicated reasoning layer to the transformer architecture, enabling it to perform more advanced logical and planned reasoning.

(2) Leveraging Expert Models, Such as MoE, for Reasoning

The MoE (Mixture of Experts) architecture allows specific parts of the model to specialize in certain tasks, such as reasoning. This approach involves "calling upon different experts for each task to perform optimal reasoning."

*2 Fangzhi Xu et al. (September 15, 2024) "Are Large Language Models Really Good Logical Reasoners? A Comprehensive Evaluation and Beyond". Iman Mirzadeh et al. (October 7, 2024) "GSM-Symbolic: Understanding the Limitations of Mathematical Reasoning in Large Language Models"

*3 Demis Hassabis (May 14, 2024) "Gemini breaks new ground with a faster model, longer context, AI agents and more"

*4 OpenAI (September 12, 2024) "Introducing OpenAI o1-preview"
James O'Donnell archive page (September 17, 2024) "Why OpenAI's new model is such a big deal"

*5 DeepSeek-R1 Release (January 20, 2025) "DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning"

*6 Parser (November 13, 2024) "Understanding LLMs' Reasoning Limits Today: Insights to Shape Your Strategy"

With the addition of higher reasoning capabilities to LLMs, questions may arise about whether this comes at the expense of the model's linguistic abilities. However, reasoning and language capabilities are closely related, and both have improved together during the evolution of the models. In fact, enhancing reasoning ability often positively influences language understanding and generation as well. In the future, the distinction between "language models" and "reasoning models" may blur, leading to more integrated AI models with broader capabilities.

(3) Incorporating Reasoning Modules into Agent Systems as Applications

The third approach involves leveraging traditional LLMs while designing systems, such as AI agents, that incorporate reasoning-specific modules. This approach goes beyond generating outputs from inputs alone, enabling models to react flexibly to environments and tasks and have decision-making processes for reasoning.

2. Innovation in Application Systems and The Role of AI Agents

Now, let's explore what AI agents are, what their characteristics are, and how they approach problem-solving and creating new value.

What is an AI Agent?

There are various definitions of AI agents, but simply put, an agent system refers to a digital system that can independently interact in a dynamic environment. These software-based agent systems existed before the rise of LLMs (large language models), but within the context of generative AI (LLM-based), AI agent systems can understand given goals and contexts, plan their actions, use online tools and data to complete tasks, and (depending on the system design) collaborate with other agents or humans, learning and improving performance through memory. In other words, AI agent systems include not only LLMs but also components such as database access, API integrations, rule-based engines (traditional AI), and interfaces with external systems.

Indeed, the enhanced reasoning capabilities in the models mentioned earlier (1) and (2) allow for high-performance reasoning models that can handle tasks like language generation, reasoning, decision-making, and problem-solving. By adding simple interaction features, it is possible to create a simple agent that can perform specific tasks. The AI assistants we use daily have evolved thanks to such reasoning models (e.g., OpenAI-o1).

However, when it comes to executing complex, multi-stage workflows in the real world, the reasoning model approach alone is insufficient. Therefore, the approach of AI agents (3) is more effective.[*7] What society needs is an AI agent capable of solving important real-world problems, such as drug discovery, materials science, coding, and physics, with high reliability.

*7 Lareina Yee et al. (July 2024) "Why agents are the next frontier of generative AI"
Deloitte AI Institute (November 2024) "Prompting for action: How AI agents are reshaping the future of work"

Figure 1 is a conceptual diagram comparing the workflow of generative AI (LLM) and AI agents. Generative AI, in principle, involves humans being actively involved in the workflow, directly influencing the AI system's decision-making and tasks.[*8] AI agents, on the other hand, can autonomously complete given tasks, with humans functioning as supervisors, intervening only in case of abnormalities or unexpected situations. However, evolved models like GPT-4 or reasoning models like OpenAI-o1 can autonomously carry out multi-step tasks from a single prompt, making it difficult to clearly distinguish them from AI agents, and this should be understood as part of an evolutionary process.

Figure 1 Conceptual diagram comparing generative AI and AI agent workflows

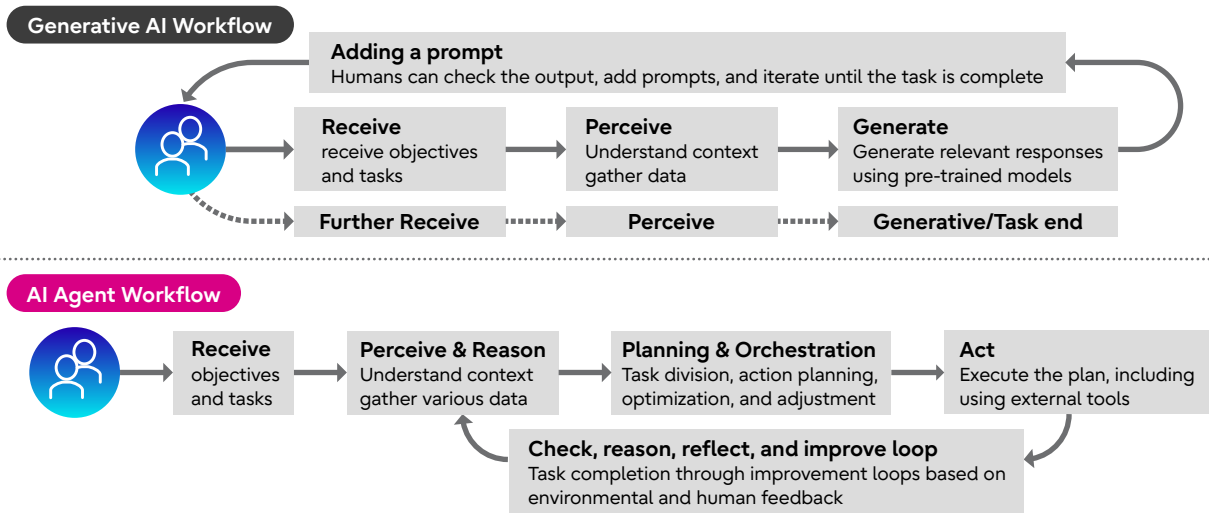

Source: WEF (December 2, 2024) "How Agentic AI will transform financial services with autonomy, efficiency and inclusion". revised and compiled by the author

AI Agent Functionality and Structure

As mentioned earlier, AI agents utilize their capabilities, such as reasoning, planning, memory, action, learning, and introspection, to address the limitations of language or reasoning models. They can support highly complex use cases across industries and business functions. This is especially useful for workflows that involve time-consuming tasks or specialized analyses, such as drug discovery or materials science.

The Basic Process of AI Agents

The basic process of an AI agent starts with the user's instruction (goal setting). This is followed by:

① Understanding the intent and context
② Creating an action plan
③ Optimization and decision-making
④ Executing the plan
⑤ Feedback for learning and improvement
⑥ Ensuring the harmony of the agent's overall functioning (see Figure 2).

*8 WEF (December 2, 2024) "How Agentic AI will transform financial services with autonomy, efficiency and inclusion"

Figure 2 Conceptual diagram of AI agent system configuration

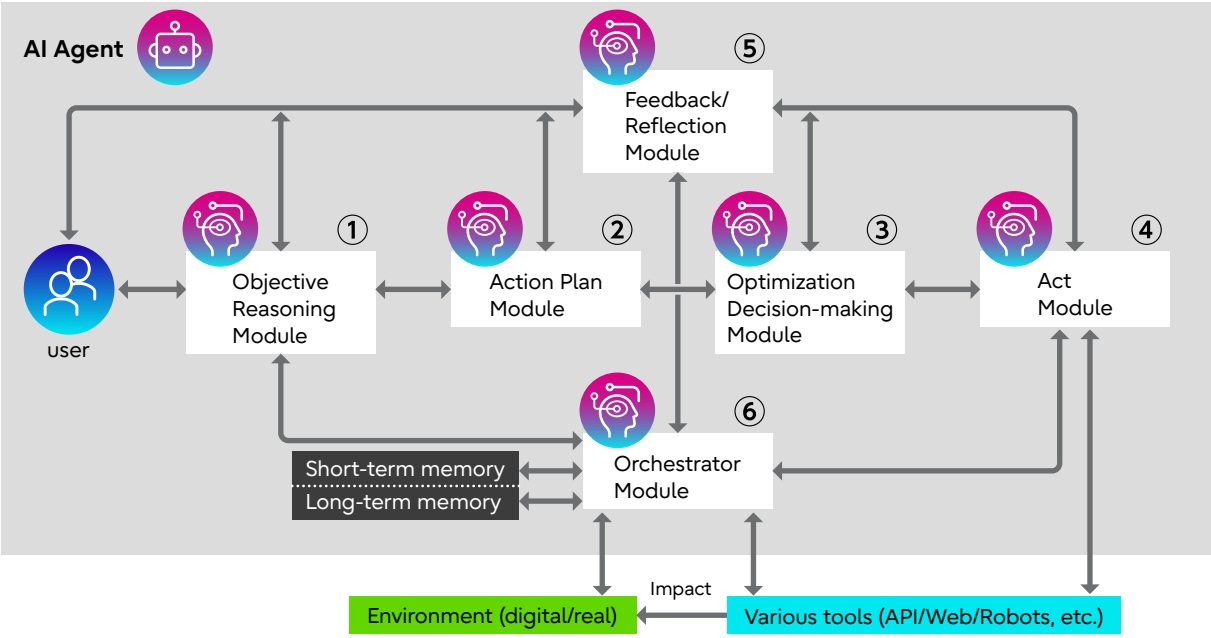

Source: Created by the author with reference to Yuheng Cheng et al. (January 7, 2024) "Exploring Large Language Model based Intelligent Agents: Definitions, Methods, and Prospects" and others.

AI agents are designed with modules or components corresponding to each of these processes.

Key Functional Modules of AI Agents

(1) Objective Reasoning Module The goal inference is typically handled by a language model or reasoning model (LLM). Depending on the design, natural language understanding (NLU) or other reasoning modules may also be involved. This module identifies the user's intent, defines the goals, and determines the actions the agent should take based on the context. If necessary, it may ask for clarification from the user.

(2) Action Plan Module Once the AI agent understands the user's intent, a separate module is usually required to break down the task and generate an action plan.*9 While LLMs excel at natural language processing, they are not specialized in task decomposition or planning. This is why task breakdowns often involve planning algorithms or reinforcement learning to handle more complex tasks and workflows. Additionally, domain-specific knowledge may be incorporated for more accurate planning.

(3) Optimization and Decision-Making Module To determine the optimal action for a planned task, AI agents need an optimization or decision-making module (such as reinforcement learning algorithms). This process involves assessing the current environment, generating options, evaluating and scoring them, and selecting the best course of action.

(4) Act Module The execution module is responsible for carrying out the planned actions and affecting the environment. It interfaces with the physical or virtual environment, translating actions into commands (such as robot instructions or software operation procedures). This module also ensures precise timing and sequence management, handles errors, collects feedback, and maintains safety and reliability.

*9 Yuheng Cheng et al. (January 7, 2024) "Exploring Large Language Model based Intelligent Agents: Definitions, Methods, and Prospects"

(5) Feedback/Reflection Module The feedback module is crucial for continuous improvement. It collects feedback from users and performance data (such as response time or success rate), analyzes it, and identifies patterns or issues. Based on this analysis, improvement measures are suggested and implemented, ensuring the agent's performance is continually enhanced.

(6) Orchestrator Module The orchestration module is essential for coordinating the movement of each component. It ensures that the processes are well-managed and that modules exchange information and cooperate effectively. The module also facilitates feedback delivery, promoting the agent's overall learning and adaptation.

Addressing and Expanding LLM Limitations

In addition to the basic modules, AI agents have advantages that help mitigate and expand upon the limitations of LLMs:

(1) Memory and Context Retention Unlike LLMs, which are stateless and lack memory of previous interactions, AI agents incorporate memory mechanisms. This allows the agent to maintain consistency in long-term engagements, leverage past context, and enhance future interactions, improving the user experience.

(2) Asynchronous and Parallel Processing Unlike LLMs, which process inputs synchronously and sequentially, AI agents can manage multiple tasks simultaneously in an asynchronous manner. This enables more effective real-time interactions and enhances efficiency and responsiveness in use cases requiring the handling of multiple queries or tasks.[*10]

(3) Fact-Checking and Real-Time Information Access By integrating real-time data validation and accessing external tools or information sources (such as databases, web APIs, or the internet), AI agents can address issues such as LLM hallucinations (incorrect or misleading outputs). This is particularly valuable in applications where accurate and up-to-date information is crucial, such as technical and scientific fields.

Multi-Agent Systems (MAS)

The concept of a single AI agent is useful for handling tasks within small teams or individual workflows. However, when it comes to complex workflows across larger organizations or ecosystems, single agents have limitations. Large organizations or domains require diverse knowledge and workflows, which means specialized agents are needed for each task. This is where Multi-Agent Systems (MAS) come into play.[*11] MAS excels in handling tasks across multiple domains through interactions and information sharing between agents.

As shown in Figure 3, the orchestration of multiple AI agents plays a crucial role in MAS. The orchestration system intelligently routes query, maintains context across interactions, and manages multiple agents to work together effectively.[*12]

*10 Janaki ram MSV (June 11, 2024) "AI Agents: Key Concepts and How They Overcome LLM Limitations"

*11 See note 8

*12 Another slightly different definition is that an orchestrator is a lead agent responsible for breaking down tasks, planning and directing other agents in the execution of subtasks, tracking overall progress, and taking corrective action when necessary. Kerem Aydın (November 17, 2024) "Which AI Agent framework should i use?"

Figure 3 Concept of a multi-agent system and orchestrator

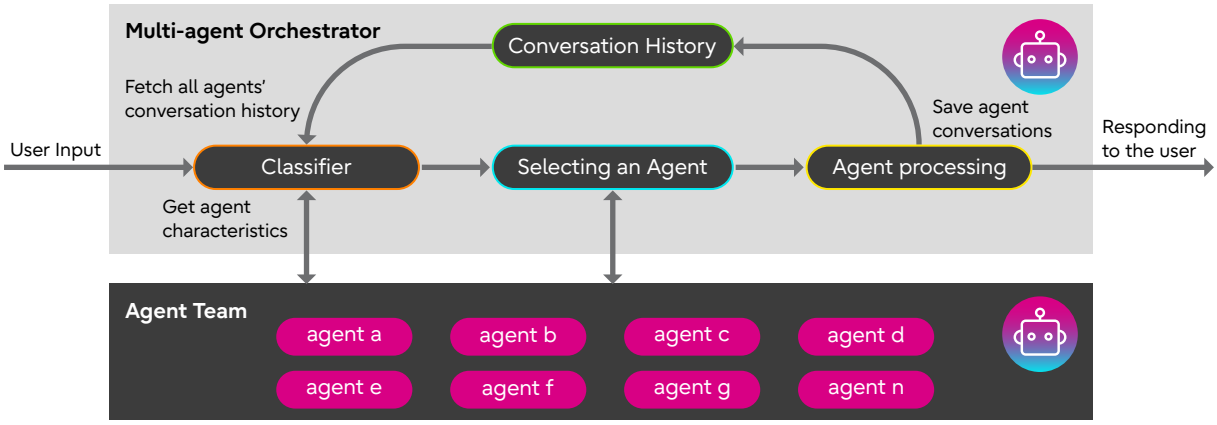

Source: AWSLABS (2024) "multi-agent-orchestrator", etc., modified by the author.

Orchestration Flow in Multi-Agent Systems

1) The process begins with user input, which is analyzed by an intelligent classifier.
2) The classifier uses the agent's characteristics and conversation history to select the most suitable agent for the task.
3) Once the agent is selected, the user input is processed.
4) The orchestration system then updates the conversation history and delivers the response to the user.

In this way, MAS doesn't just perform reasoning and actions through standalone AI agents but can orchestrate a team of agents to handle more complex workflows. Some leading vendors have already started developing pre-built components for quick deployment and scalable solutions, which are expected to be widely adopted across industries.*13

3. Use Cases for AI Agents and Examples of Their Adoption by Leading Companies

In recent years, AI technology has advanced rapidly. The industrial society has entered the implementation phase after passing through the learning stage of generative AI, though many challenges remain. On the other hand, AI agents/multi-agent systems are expected to have direct effects on improving business operations, leading to rapidly growing interest in these technologies across industries. As AI agent technology advances, proof of concepts and demonstrations are being conducted, and use case development is progressing.

AI Agent/Multi-Agent Use Case Examples

AI agent/multi-agent use cases are still in the development stage, but potential use cases are being developed for customer service, workflow management, healthcare services, financial services, and risk management. Table 1 summarizes seven representative use cases proposed by various organizations. Use Cases 1 and 2 are for the financial industry, while Use Cases 3 to 7 can be applied across all industries.

*13 For example: Microsoft's "Semantic Kernel documentation"

Table 1 Seven Examples of AI Agent/Multi-Agent Use Cases

Overview	Key benefits
Use Cases 1. Financial Advisory and Wealth Management	
• Traditional financial advice categorizes clients by age and income, but can miss individual needs. • There is a growing demand for personalized advice. • Multi-agent AI analyzes a variety of data to create a personalized financial plan that can be adjusted based on the situation.	• Hyper-personalization • Continuous fine-tuning of financial plans and strategies • Improved customer satisfaction • Improved scalability
Use Cases 2. Loan Underwriting	
• Financial institutions prepare credit risk reports to assess the credit risk of borrowers. • This process is time-consuming because it involves analyzing a lot of information and collaborating with relevant parties to perform specialized analysis. • The agent system allows multiple specialized agents to share tasks and efficiently handle credit risk scenarios.	• Reduce review cycle time by 20% to 60%, improving efficiency • Integrate data from multiple systems to produce high-quality content • Quickly review details of generated text and numbers, making it easy to verify output
Use Cases 3. Dynamic pricing and personalized promotions	
• Traditional pricing strategies use static models that don't account for market changes or customer behavior. • Multi-agent AI analyzes real-time data such as competitor prices and customer purchase history to dynamically adjust prices. • Personalize promotions based on individual customer preferences and purchasing habits to improve conversion rates and customer satisfaction.	• Instant price adjustments when things change • Personalized offers • Maximize margins, minimize discounts
Use Cases 4. Personalized customer support	
• Traditional support systems rely on scripts and can't handle complex inquiries, which can lead to customer dissatisfaction. • Multi-agent AI considers customer history and preferences to handle these inquiries and generate natural responses. • This reduces the need for escalation and improves customer satisfaction	• Improved consistency and scalability • Improved customer experience, including improved quality and time savings by learning from interactions
Use Cases 5. Code documentation and modernization	
• Legacy software in large enterprises can pose security risks and slow business innovation. • Modernizing these systems requires extensive code understanding and manual documentation of business logic, which is complex, time-consuming, and costly. • AI agents streamline the analysis, documentation, and translation of legacy code, while quality assurance agents generate test cases to improve accuracy.	• Streamline the software migration process and improve productivity • Repeatable process can be reused for other software migrations, reducing costs • Provide accurate documentation and test cases according to organizational standards to ensure quality
Use Cases 6. Talent Acquisition and Recruitment	
• Traditional hiring processes are inefficient, with a lot of manual resume review and repetitive assessments. • AI agents automate the hiring process by using natural language processing to analyze resumes, score candidates based on skills and experience, and conduct initial interviews. • These systems work in tandem with HR to efficiently screen the right candidates and make hiring decisions in a compliant manner.	• Improved efficiency • Improved candidate matching • Reduced bias • Dynamic scalability
Use Cases 7. Cyber Security	
• Cybersecurity professionals face a global talent shortage of 4 million people • Malicious attackers increasingly use generative AI to infiltrate security systems • New agent-based cybersecurity systems improve professional efficiency by automating attack detection and reporting	• Reduces the workload of human experts by up to 90%. • Automatically detects vulnerabilities in new code and provides immediate feedback to development teams • Improves system security and enables early detection and response to attacks

Source: Created by the author with reference to Deloitte "Prompting for action How AI agents are reshaping the future of work"; Deloitte "Autonomous generative AI agents: Under development", and McKinsey "Why agents are the next frontier of generative AI"

Figure 4 Schematic of the multi-agent use case workflow: Credit Risk Memo Creation

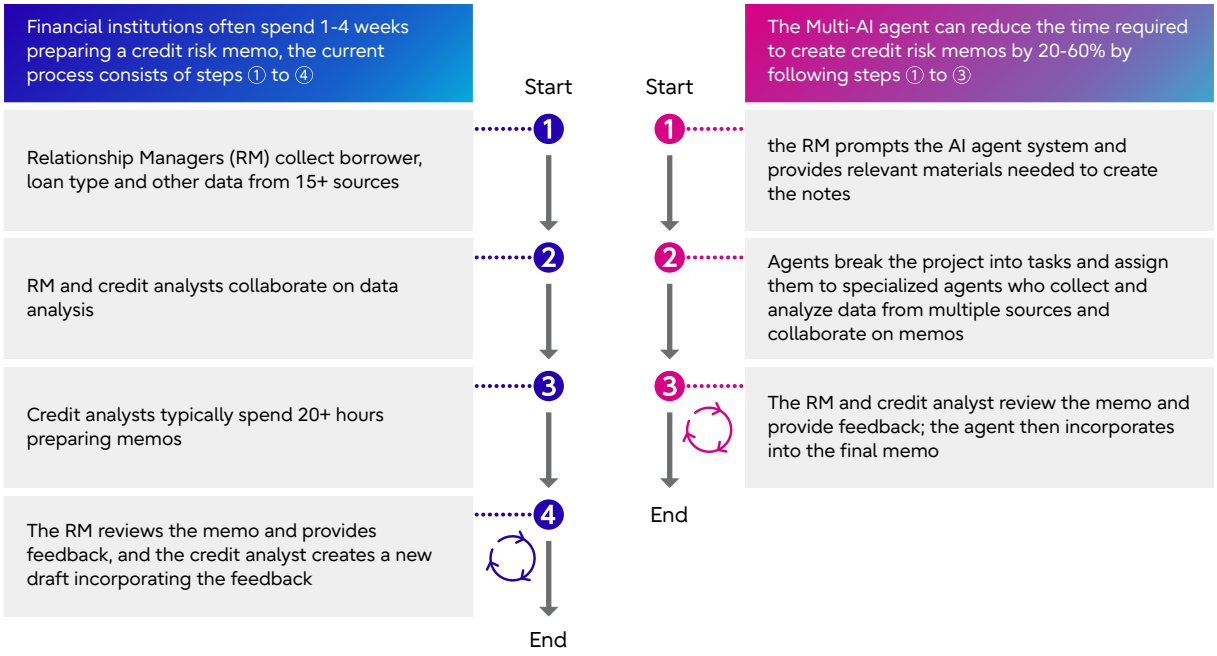

Source: Created by the author with reference to McKinsey (July 2024) "Why agents are the next frontier of generative AI"

Figure 4 shows a multi-agent use case for a financial institution to generate credit risk memos for loan underwriting.[*14] As a solution based on agent-based (multi-agent) systems, a system using agents with specialized roles can process various credit risk scenarios. When a human user provides task plans in natural language, the agent breaks them down into executable sub-tasks. For example, one agent takes on the role of a relationship manager and handles communication between the borrower and the financial institution. Other agents handle document collection, financial analysis, check and correction of data, and after verification by the RM and credit analysts, the final credit memo is completed. This process significantly reduces review cycle time by 20–60%, enabling rapid analysis of data across multiple systems and quick validation of output results.

Examples of AI Agent/Multi-Agent Development and Use

As mentioned earlier, industries (end users) are gaining a deeper understanding and knowledge of AI agents/multi-agent systems. The AI agent (including multi-agent systems) market is expected to expand significantly from 5.7 billion USD in 2024 to 52.1 billion USD by 2030, with a compound annual growth rate (CAGR) of 45% during the forecast period.[*15] Large tech companies and startups are accelerating the development of technology and applications, while industries are accelerating the process from proof of concept to pilots and full adoption. Let's take a look at a few examples.

*14 Lareina Yee et al. "Why agents are the next frontier of generative AI"
*15 BCG "AI Agents"

(1) Auquan: Specializing in AI Agent Development for the Financial Sector

Auquan, a startup in the AI agent field, has developed an approach to automate the creation of financial analyst reports (including credit risk memos) using multi-agent systems, as described earlier.[*16] Unlike OpenAI's "Deep Research," which handles general reports, Auquan's system is specifically designed for financial applications. The agent team reviews data from internal documents such as presentations and PDFs, as well as from the broader internet, and answers them in the reports. Depending on the underlying model used, the cost of reports generated by Auquan's system is estimated to be 80–90% cheaper than human-created reports. This results in significant improvements in productivity.

Currently, Auquan's clients include MetLife, UBS, and Capital Group, and they are expanding the scope of their system to cover the entire risk, compliance, and sustainability workflow, including due diligence, monitoring, KYC/KYB, and LP reports.[*17]

(2) Fujitsu: Working on Industry-Wide AI Agent Development

Fujitsu, an IT service vendor, is at the forefront of AI agent technology and application development, particularly in industry-specific generative AI. Let's look at three key results.

1) Security-Specific Multi-AI Agent System

Fujitsu has developed a multi-AI agent security technology that links agents specialized in attack and defense to protect IT systems of companies and public organizations from new threats.[*18] This system consists of three parts: orchestration functionality that transparently links AI agents, security AI agents that perform attack, defense, and testing, and technologies to strengthen security for generative AI (see Figure 5).

Figure 5 Multi-AI Agent Security Technology Overview

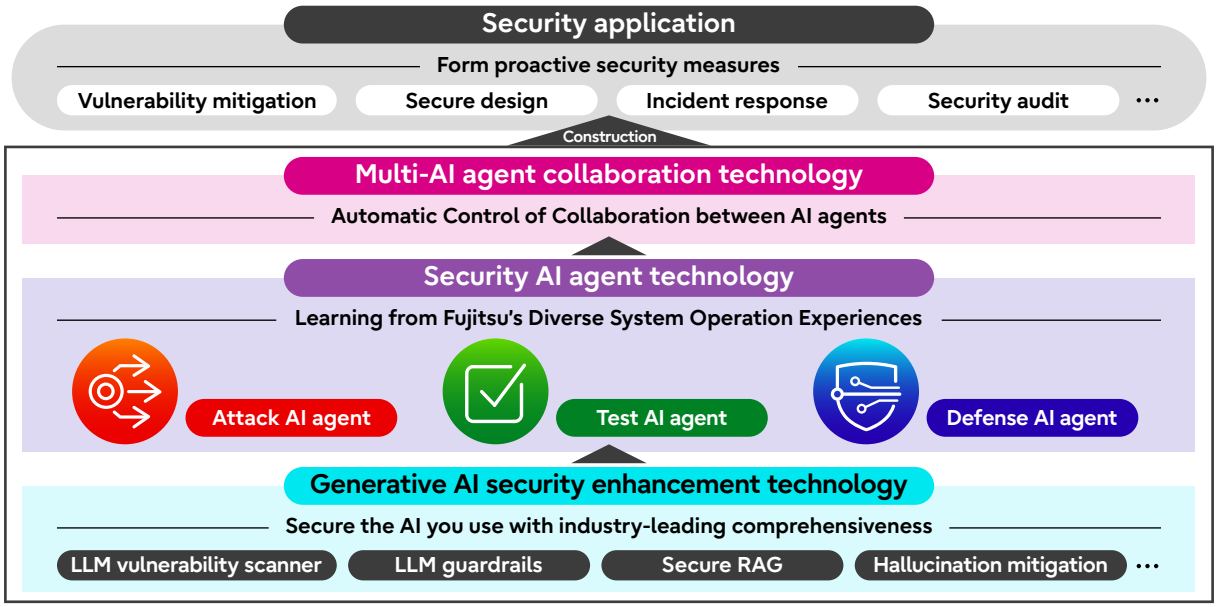

Source: Fujitsu press release (December 12, 2024) "Fujitsu develops world's first multi-AI agent security technology to protect against vulnerabilities and new threats"

*16 EVIDENT (February 6, 2025) "Interview with Chandini Jain"
*17 Auquan (January 14, 2025) "Auquan Enters 2025 with Breakthrough AI Innovations and Financial Services Impact"
*18 Fujitsu press release (December 12, 2024) "Fujitsu develops world's first multi-AI agent security technology to protect against vulnerabilities and new threats"

Through this, AI agents perform attack and defense simulations on a "cyber twin" and strengthen measures against new vulnerabilities. This allows companies to take effective defensive actions against new attacks. Fujitsu's AI agents are particularly notable for addressing complex problems through collaborative learning between agents with different skills.

This technology allows IT administrators, even those without security expertise, to build proactive security measures and operate IT systems securely while safely utilizing generative AI.

2) AI Agent for Video Analysis

Most AI agent use cases are related to digital environments, but Fujitsu has developed AI agent technology for the real world.[*19] This system helps create safe and efficient work environments in warehouses and factories by reading real-time video footage and performing video analysis. Based on situational awareness, the AI agent autonomously executes tasks such as providing alerts or suggestions to site managers and workers (see Figure 6). This is also a use case for a single AI agent.

Figure 6 Image of an AI agent for real-world video analysis

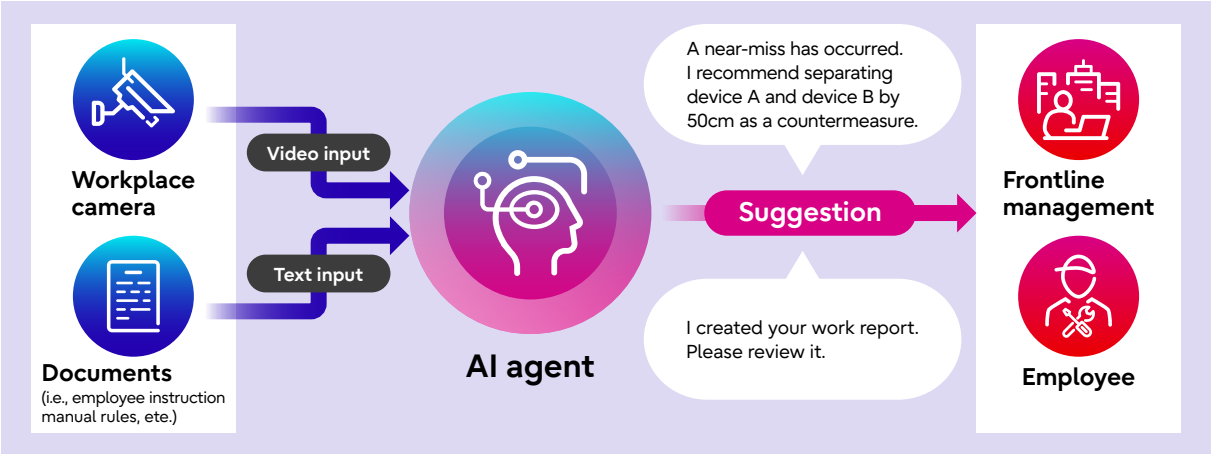

Source: Fujitsu press release (December 12, 2024) "Fujitsu develops video analytics AI agent to support safe, secure, and efficient frontline workplaces"

Moving forward, Fujitsu plans to expand the use of AI agents to include specialized agents for production management, with applications expected in many industries such as healthcare and retail management.

3) Meeting AI Agent

This AI agent, developed by Fujitsu, can autonomously participate in meetings related to profits, losses, and business negotiations and has the ability to share relevant information and propose strategies.[*20] For example, when the meeting AI agent participates in a meeting and hears a statement such as 'Sales in the Asia region are half of what they were last year,' it performs data analysis. As a result, it presents sales by region in a bar graph, showing that sales in the Asia region are 54% of last year's figures. This helps smooth the progression of the meeting and supports the generation of productive conclusions.

*19 Fujitsu press release (December 12, 2024) "Fujitsu develops video analytics AI agent to support safe, secure, and efficient frontline workplaces"

*20 Fujitsu Press Release (October 23, 2024) "Fujitsu to offer AI agents that can both collaborate and engage in high-level tasks autonomously"

(3) Adoption Examples from Leading Companies

Agent-based AI has been ranked as the number one strategic technology trend for 2025 by Gartner.[*21] Deloitte forecasts that by 2025, 25% of companies using generative AI will start AI agent pilots or proofs of concept, with this number expected to rise to 50% by 2027.[*22]

While it will take time for AI agents to be widely adopted across companies, some advanced organizations are already integrating them into existing tasks and workflows for certain use cases. Table 2 summarizes examples from eight companies that have piloted or implemented AI agents.

Table 2 Examples of end users who are early adopters of AI agents (8 companies)

Enterprise	Use Case Overview
Johnson & Johnson (J&J)[*23] Drug Discovery with AI Agents	• J&J uses AI agents to optimize processes such as solvent switching (replacing one solvent with another to crystallize molecules to make a drug). • Without AI, scientists would manually test conditions over and over again, but now the agents use machine learning and digital twins to speed up this process, making it more cost-effective and reliable. • J&J continues to monitor the agents to prevent errors such as hallucinations.
KG Steel[*24] Autonomous control agent for steel manufacturers	• KG Steel faced two challenges: high liquefied natural gas (LNG) energy costs and inconsistent product quality due to skill gaps caused by an aging workforce. • Solution: Using predictive control optimization models, system integration was performed and partial automation of furnace operations was achieved by inputting agent output directly into the furnace control system. • Benefits: LNG consumption was reduced by approximately 2% and product quality variability was reduced.
Siemens Electronics Works Amberg (EWA)[*25] Autonomous quality control AI agent	• EWA's goal is to achieve a First Pass Rate (FPY) of over 95% and a Defect Per Million Connections (DPMC) of less than 10, but this goal has been elusive because a PCB has up to 3,800 quality features that are very difficult to monitor. • EWA has developed a patented autonomous AI quality control agent. The agent helps configure solder paste printers, reduces process times for complex tasks, continuously improves process parameters, and ultimately enables automatic adjustment.
Cosentino[*26] Customer Service Agent	• Cosentino, a Spanish manufacturer, uses AI agents as a "digital workforce" for its customer service operations, which are trained to ensure accuracy and designed to follow strict protocols. • Cosentino's digital workforce replaces three to four employees who were previously responsible for processing orders, freeing up human staff to focus on higher-priority tasks. The company continually monitors and retrains these agents to maintain their performance. • The company has also successfully deployed AI agents in credit management.
Moody's[*27] Financial Analysis Agent	• Moody's uses AI agents to analyze SEC filings and industry data to enhance its research capabilities. The agents specialize in specific tasks and use a multi-agent system in which other agents validate the results. • Moody's has also developed 35 agents, including agents responsible for smaller tasks such as project management, that work with supervising agents to create a "multi-agent system."

*21 Gartner (October 21, 2024) "Gartner Top 10 Strategic Technology Trends for 2025"
*22 Ariane Bucaille et al. (November 19 2024) "TMT Predictions 2025: Bridging the gaps"
*23 J&J (October 10,2024) "6 ways Johnson & Johnson is using AI to help advance healthcare".
Belle Lin (January 6, 2025) "How Are Companies Using AI Agents? Here's a Look at Five Early Users of the Bots"
*24 WEF (January 2025) "Frontier Technologies in Industrial Operations: The Rise of Artificial Intelligence Agents"
*25 See note 23
*26 Ryan Stevens (November 5, 2024) "Cosentino leverages AI to optimize global operations and drive efficiency". Belle Lin (January 6, 2025)
*27 Moody's (2024) "GenAI's transformative potential in the financial sector: the evolution of agents".
Ari Lehavi et al. (2024) "The rise of the digital colleague: evaluating companies with Moody's AI agents"

Enterprise	Use Case Overview
Capital One[*28] Chat Concierge Agent	• Capital One has launched a chat concierge that uses multiple AI agents to provide comprehensive answers to customers' questions about buying a car. It helps with all aspects of the buying process, from comparing cars to booking test drives, and reduces the cognitive load on customers. • The customized model, based on Llama AI, answers customers' questions instantly and allows them to complete multiple tasks at once, such as getting a trade-in quote or setting up an appointment with a salesperson.
Ebay[*29] Code Development and Marketing Agent	• eBay developed its own AI agent framework to streamline tasks like writing code and creating marketing campaigns. The agents help buyers find products and sellers list products. • The framework integrates multiple language models to handle tasks such as translation and suggesting code snippets. • eBay also introduced OpenAI's AI agent operator, taking another step toward the future of agent ecommerce.
Deutsche Telekom[*30] Agent for employees	• Deutsche Telekom has deployed askT, an AI agent that answers employee questions about policies, benefits, and products, to be used by approximately 10,000 employees each week. • The company is testing whether askT can automate tasks such as entering vacation requests into its HR system, reducing administrative work and freeing employees to focus on higher-value tasks. • Deutsche Telekom has also developed its own multi-agent architecture and is developing various applications, such as a billing agent.

Source: Compiled by the author from related company press releases and related reports.

From these use cases, several insights can be drawn:

1) Expanding from Simple to Complex Task Automation

In general, the adoption begins with automating simple tasks, gradually progressing to more complex ones. In particular, examples in manufacturing and healthcare start with single tasks like parameter setting and quality control. On the other hand, more complex tasks requiring multiple steps are being applied in industries like finance (analyst tasks) and e-commerce, as well as for employee agents. A single agent is typically used for individual tasks, while multi-agent systems are applied to tasks that require multiple steps.

2) Virtual AI Agents Leading the Way, While Embodied AI Agents Are Under Development

AI agents are classified into two types: virtual AI agents, which operate entirely in digital environments and autonomously achieve predefined goals in digital applications, and embodied AI agents, which are integrated into physical systems like robots and interact with the physical world.[*31] Most of the development examples and end-user use cases discussed earlier fall under "virtual AI agents." This is because advancements in AI models (such as large language models and improved inference capabilities) are progressing rapidly in digital environments, making virtual AI agent technology practical.

*28 Jocelyn Mintz (January 31, 2025) "Capital One's new AI agent will help you buy your next car"
*29 eBay (January 23, 2025) "Announcing a New Collaboration Between eBay and OpenAI".
Belle Lin (January 6, 2025)
*30 InfoQ (November 8, 2024) "Launching AI Agents across Europe at Breakneck Speed with an Agent Computing Platform".
Belle Lin (January 6, 2025)
*31 BCG "BCG-WEF Project: AI-Powered Industrial Operations"

On the other hand, embodied AI agents require integration into physical systems like robots, which must recognize and interact with the environment through dynamic and complex movements. These agents are still in the research and development phase.[*32] These agents recognize the world through sensors (such as cameras, radar, lidar, microphones, etc.) and take action through actuators like advanced grippers.

Figure 7 Evolution of embodied AI robot capabilities

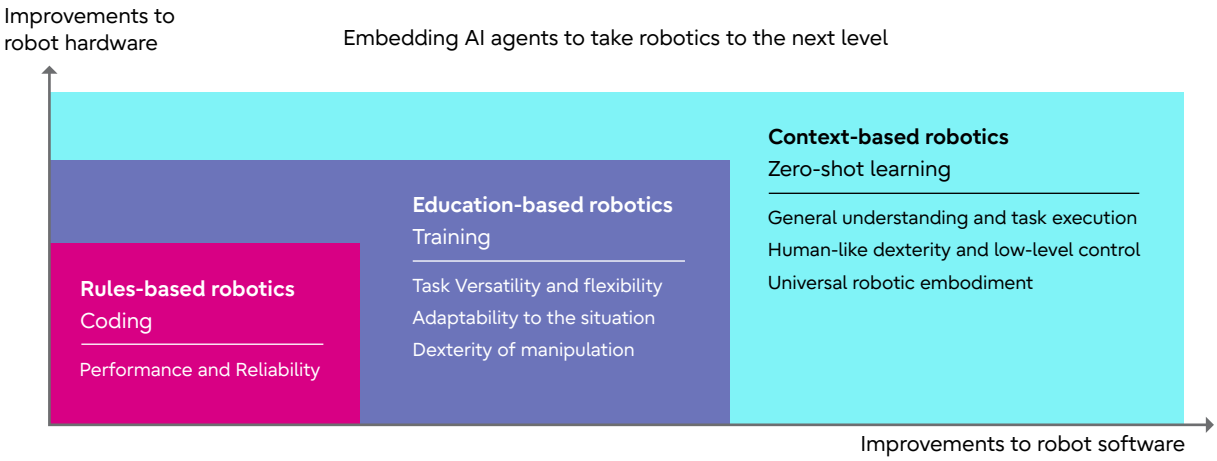

Source: Created by the author with reference to WEF "Frontier Technologies in Industrial Operations: The Rise of Artificial Intelligence Agents"

BMW is conducting experimental trials using humanoid robots (a type of "embodied AI agent") in the assembly preparation at their Spartanburg plant.[*33] The practical application of these agents is accelerating due to advancements in robotics foundation models (RFMs), reinforcement learning technologies, and progress in physical robotics (hardware) (see Figure 7). Depending on the definition of "embodied AI agents," we consider that fully realized "embodied AI agents" will be achieved when robots reach context-aware capabilities at the robot level.

3) Balancing AI Agent Potential and Reliability

Companies like Johnson & Johnson, Casentino, Moody's, and eBay are working to unlock the potential of AI agents while also monitoring and mitigating risks such as hallucinations and security threats. In fact, research firm Gartner predicts that by 2028, at least 15% of business decision-making will be autonomously performed by AI agents (up from 0% in 2024).[*34] However, Gartner also warns that by 2028, 25% of corporate breaches will be caused by the malicious use of AI agents by external or internal attackers.[*35] Therefore, balancing the benefits and reliability of AI agent adoption is crucial.

*32 Jianmin Jin (January 2025) "Designing the Next Generation of Intelligent Manufacturing with Generative AI"
*33 See note 22
*34 Gartner (October 21, 2024) "Gartner Identifies the Top 10 Strategic Technology Trends for 2025"
*35 Gartner (October 22, 2024) "Gartner Unveils Top Predictions for IT Organizations and Users in 2025 and Beyond"

4. The Future of Companies Co-Creating with Human and AI Agents

Over the past two years, the technological evolution has rapidly progressed from traditional analytical AI to generative AI, and now to AI agents. While the entire industry recognizes the high potential of AI agents beyond the limitations of generative technology, there are still challenges to overcome, such as significant upfront investments in core technologies, the reliability of AI agent systems, and the need for skilled personnel and new expertise. As seen in Chapter 3, major tech companies and startups are accelerating AI agent technology, while end-users are experimenting with learning, proof of concepts, and pilot programs to explore the impact on industries and businesses, and to innovate the methods and functions of AI agent systems.

As machine (AI) natural language understanding advances, interactions between humans and machines will become more intuitive and fluid. In this future business environment where humans and machines collaborate, AI agents will take the lead in autonomously executing business processes. Humans will redefine their roles from operators of daily tasks to orchestrators of AI-driven systems. Specifically, humans will oversee, intervene in, and adjust AI systems when necessary (see Figure 8). With the improvement of reasoning capabilities, breakthroughs in hallucination mitigation technologies, and advancements in techniques to reduce risks of violating ethics and privacy, the degree of human intervention will gradually become minimal. Additionally, humans will focus more on tasks requiring creativity, emotional intelligence, and physical dexterity, in collaboration with machines.

AI agent technology, one of the leading innovations of the intelligent era, is advancing rapidly. For industries to fully harness the transformative power and potential of AI agents, it is essential to set value-driven goals, enhance skills and capabilities, and establish a governance framework that includes change management and risk management. By putting these foundations in place, companies will realize a future where humans and AI agents work in harmony to co-create the workplace.

Figure 8 The future vision of enterprises co-created by humans and AI agents

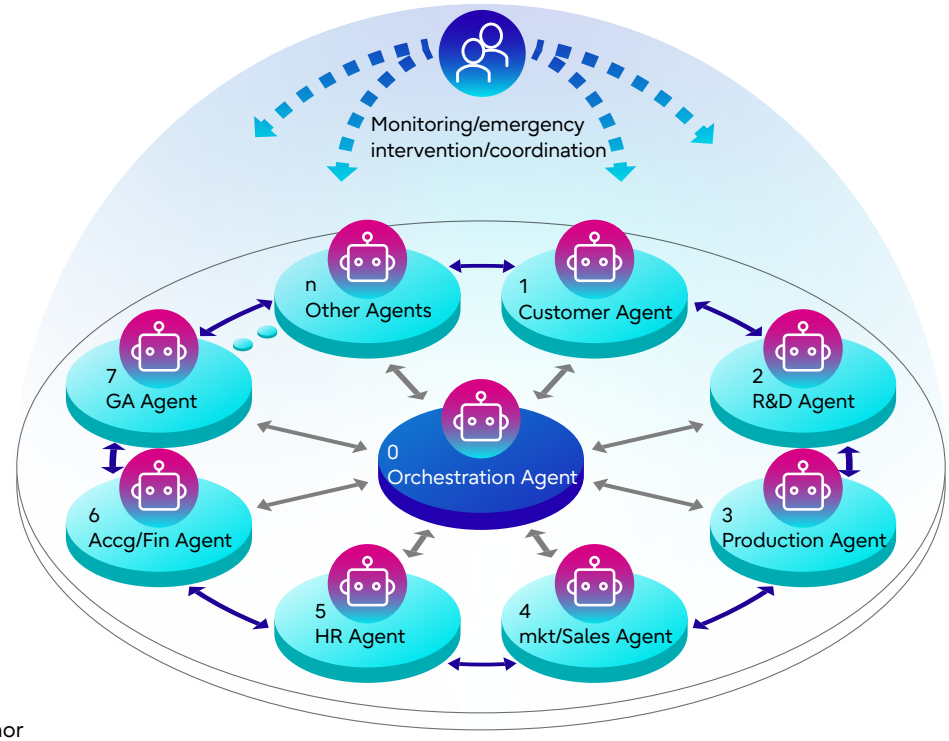

Source: Author

Key References

1. BCG "AI Agents"
2. BCG "BCG-WEF Project: AI-Powered Industrial Operations"
3. Deloitte AI Institute (November 2024) "Prompting for action: How AI agents are reshaping the future of work"
4. Deloitte "Autonomous generative AI agents: Under development"
5. Fangzhi Xu et al. (September 15, 2024)
   "Are Large Language Models Really Good Logical Reasoners? A Comprehensive Evaluation and Beyond"
6. Iman Mirzadeh et al. (October 7, 2024)
   "GSM-Symbolic: Understanding the Limitations of Mathematical Reasoning in Large Language Models"
7. Janaki ram MSV (June 11, 2024) "AI Agents: Key Concepts and How They Overcome LLM Limitations"
8. Jianmin Jin (January 2025) "Designing the Next Generation of Intelligent Manufacturing with Generative AI"
9. Lareina Yee et al. (July 2024) "Why agents are the next frontier of generative AI"
10. WEF (December 2, 2024) "How Agentic AI will transform financial services with autonomy, efficiency and inclusion"
11. WEF (January 2025) "Frontier Technologies in Industrial Operations: The Rise of Artificial Intelligence Agents"
12. Yuheng Cheng et al. (January 7, 2024)
    "Exploring Large Language Model based Intelligent Agents: Definitions, Methods, and Prospects"

About the author

Dr. Jianmin Jin

2020 Fujitsu Ltd., Chief Digital Economist
1998 Fujitsu Research Institute, Senior Fellow

Dr. Jin's research mainly focuses on global economic, digital innovation/ digital transformation, and Dr. Jin has published books such as "Towards the create of a Japanese version of Silicon Valley", etc.
Recent writings: the following Fujitsu Insight Paper, etc.

- Designing the Next Generation of Intelligent Manufacturing with Generative AI (2025)
- Innovative Banking with Generative AI (2024)
- Leveraging the LLM: Strategy from Model Selection to Optimization Insight for top management (2024)
- Generative AI: Use Cases as the Pathway to Value Creation (2024)
- Transforming Supply Chains to Be More Productive, Resilient, and Sustainable (2023)

The author would like to thank Jungo Okai, Ryoma Ohashi, Koichi Shirahata, Hiroshi Nishikawa, Hiromichi Kobashi, Kenichi Kobayashi, Naomi Hadatsuki and Takashi Shinden for their review and invaluable advice during the development of this paper. The author would also like to thank Hiroko Meguro, Yukiko Sato, and Mitsuo Tsukihara for their unwavering daily support.

FUJITSU

March, 2025 v1.0