protiviti®
Global Business Consulting

BOARD PERSPECTIVES

ISSUE 186

Agentic AI: What It Is and Why Boards Should Care

In 2024, generative artificial intelligence (AI) was all the rage. In 2025, agentic AI has surfaced as the next frontier of AI deployment. What is agentic AI, and why is it important for directors to understand how management intends to use it?

Generative AI is capable of creating new and novel content (such as text and images) in response to prompts or requests. Generative AI models have rapidly evolved in the last three years, progressing from text-only models with limited reasoning abilities to multimodal models with advanced step-by-step reasoning and research capabilities. These advancements are key to unlocking the next era of AI evolution, known as agentic AI.

"Agentic AI" describes an AI system that employs a large language model (LLM), as well as other technologies, as part of its ecosystem to enable autonomous task planning and support reasoning to achieve human-defined objectives. Operating autonomously without continuous human intervention, agentic AI gathers and evaluates data to determine the best course of action and acts to accomplish its stated objective. Unlike generative AI, which produces content for further action, agentic AI systems leverage an LLM to enable autonomous AI agents to process and interpret large volumes of data, create a plan and produce outputs aligned to their assigned goals.

In some sense, agentic AI represents the progression from earlier automation technologies such as robotic process automation (RPA). However, there is significant confusion around the distinguishing features and differences between common automation technologies. The table below contrasts agentic AI with two common automation applications: RPA and intelligent automation (automation supported by AI).

	RPA	Intelligent Automation	Agentic AI
Purpose	Performs a predefined action	Follows a list of steps and uses an LLM in at least one step	Develops a plan and executes it to accomplish an objective
Autonomy	Low	Moderate	High
Adaptability	None	Learns from past experience	Adapts to real-time feedback and changing circumstances
Decisioning	Limited, follows predefined rules and scripts	Advanced, with machine learning algorithms and data analysis	Highly advanced, dynamically adjusting plan to achieve the stated goal
Task Capacity	Repetitive and routine	Complex, requiring reasoning, prediction and problem-solving	Dynamic and goal-oriented; execution steps may differ
Example	Automates responses to customer inquiries with scripted, predefined templates	Interacts with customers through a natural language chatbot, resolves common issues and routes complex problems to human agents	Facilitates a procurement process proactively — researching vendors, negotiating terms, preparing contracts and involving humans for final approval

In summary, agentic AI systems leverage advanced AI techniques, including reinforcement learning, to enable systems to learn from their experiences and improve over time. It is a collection of technologies and models for solving issues with limited human supervision and is often managed by an “orchestrator agent.” Within that framework, AI agents are designed to handle tasks and processes with appropriate autonomy. The orchestrator decides which agent is most suitable for performing specific tasks with intention to optimise workflow and decision-making seamlessly within the overall system.

Why Directors Should Care: The Opportunities

Because AI agents have the versatility to transform the workplace, redefine the nature of work performed by humans, and reimagine customer-facing and back-office processes, the board should understand how management intends to train, deploy, measure and monitor these agents in a responsible manner consistent with the overall corporate strategy. In this respect, AI agents are similar to human employees who are given performance expectations and training and whose performance is measured and monitored. Companies that view AI agents as workforce extensions and integrate them into their operations and monitor their performance, as they would human employees, are likely to be better positioned to optimise this next-gen technology.

As illustrated in the table on the previous page, agentic AI can enhance automation with advanced decision-making capabilities and adapt and learn in unpredictable environments — without continuous human supervision. In particular, AI agents offer value creation opportunities in situations calling for a nonprescriptive AI solution. An example would be lead identification and development, where the agent may need to be creative about how to connect with and reach an individual, but the specific steps it takes to do so are not critically important. For transaction reconciliation, where management cares about the steps taken to perform the task, a more conventional automation or intelligent automation is probably a better fit.

With that important distinction in mind, examples of value-creation opportunities of agentic AI stretch across industries. They include:

- **Boosting productivity and efficiency.** Whether it's manufacturing, logistics or services, AI agents are designed to operate autonomously, make decisions and take actions based on real-time data and learned experiences, adapt to new situations, and continuously improve their performance over time.
- **Scaling dynamically as task and data volumes increase.** As businesses grow, AI agents can scale their operations to handle varying workloads and adapt to new tasks without human intervention. This flexibility aids their deployment in diverse and expansive ways without significant resource increases.

> *Agentic AI can enhance automation with advanced decision-making capabilities and adapt and learn in unpredictable environments — without continuous human supervision.*

- **Improving speed of customer service.** AI agents can handle queries faster and more accurately than humans to deliver personalised experiences at scale. This transformation can improve response times and resolution rates in consumer-facing industries.
- **Enhancing decision-making.** In finance, healthcare and strategy-setting, agentic AI facilitates smarter, more informed decisions by analysing large amounts of data in real time, spotting hidden patterns and highlighting key insights.
- **Innovating products and services.** For example, revolutionising patient care with personalised advice and remote treatment would be game-altering.
- **Adding another dimension of labour.** AI agents can be transformative in performing work along with permanent and temporary employees and contractors, as part of an integrated workforce and thereby reduce costs and minimise human errors.

Opportunities for transforming business processes with agentic AI to increase efficiency and compress time to recognise and react are essentially endless. For example, in cybersecurity, AI agents autonomously monitor and respond to incidents. In IT support, they act autonomously to perform software updates. In human resources, they streamline hiring processes. Supply chain management benefits from AI-driven logistics. Marketing strategies are enhanced by AI-driven analysis of consumer behaviour.

Why Directors Should Care: The Risks

The deployment of agentic AI also comes with certain risks that organisations need to evaluate. In our latest Top Risks survey,[1] one of the top 10 concerns on a global basis was the emergence of new risks from implementing AI.

> *This discussion is relevant to board members because the journey from traditional automation to agentic AI represents a significant and transformative technological evolution that paves the way for a future where technology and human ingenuity work together to achieve remarkable outcomes. But it also comes with challenges.*

[1] *Executive Perspectives on Top Risks for the Near- and Long-Term*, Protiviti and NC State University's ERM Initiative, February 2025: www.protiviti.com/us-en/survey/executive-perspectives-top-risks.

Below are key risks associated with agentic AI deployments that directors should consider. Those that are generally similar to other forms of AI, but may be elevated in importance due to the attributes and uses of agentic AI, include:

- **Accountability.** Who is to blame when an AI agent makes a decision or commits an act that leads to harm or damage — the vendor, the trainer (likely an employee working in concert with the developer), the user or the user's manager? What is the process for dissecting the situation when something goes wrong, why it went wrong and where in the chain things broke down?
- **Potential for bias.** When making hiring and lending decisions or enforcing laws and regulations, biases in the data used to train an AI agent can lead to unfair, discriminatory and even illegal outcomes. This problem can amplify quickly when the agent is rewarded for a biased decision.[2] The risk of bias also surfaces compliance issues — for example, NYC Local Law 144 prohibits employers from using automated employment decision tools without an annual independent audit that assesses such tools for bias.
- **Data privacy and security.** These risks arise from the manner in which agentic AI systems operate, the data by which they are trained and on which they rely, and their interactions with users and other systems. They add complications to the risk of sensitive information being misused or exposed in a data breach.
- **Competition for talent.** The adoption of agentic AI (as well as other emerging technologies) is driving increased demand for professionals with expertise in AI, machine learning, data science and related fields. As organisations recognise the potential of agentic AI, they must seek and retain those skilled individuals who are competent in developing and managing these systems.

In addition to the above risks, there are several risks unique to agentic AI that arise primarily from the autonomous decision-making capabilities of these systems:

- **Loss of control.** As agentic AI systems become more complex and autonomous, they could act in a manner inconsistent with their stated goals. When confronted with situations not outlined as an operating boundary, how will an AI agent behave? While these systems adapt and learn, they operate without direct human intervention. The risk of rogue behaviour can present significant issues in situations where reliability and transparency are critical.

[2] An AI agent learns from its experiences through reinforcement learning when it receives positive feedback for actions beneficial to its stated goal, which encourages repetition of those behaviours. For instance, when handling customer service interactions, each successful interaction resulting in a satisfied customer, without needing human intervention, would result in a "+1" feedback. When rewarded with positive feedback, the AI agent is guided to make similar decisions in the future.

- **A very different kind of collaboration.** As AI agents take on tasks traditionally handled by humans, they will create new job functions while making some existing ones obsolete. This transition will require substantial efforts to upskill and reskill employees to ensure effective interaction with and oversight of AI agents and a smooth shift to a newly transformed work environment. What's unique about agentic AI is it leads to a workplace where human agents and AI agents collaborate, creating a very different culture to manage. As the CEO of Salesforce recently indicated at Davos, "From this point forward … we will be managing not only human workers but also digital workers."[3]
- **Potential training conflicts of interest.** Those whose job functions are performed by an AI agent may be required to assist in training the agent. This process may present a challenging task for companies to manage as it could lead to job insecurity, resistance to change, accountability concerns, potential for bias in the training data and loss of morale.
- **Loss of experiential learning.** As AI agents take over more tasks, there is a risk of a deterioration in human experiential learning and inherent knowledge from engaging in the critical thinking underlying problem-solving and decision-making. How will this affect the resiliency of the workforce over time? However, the transparency an AI system offers can benefit humans by providing traceable, documented steps supporting its decisions.

> *While AI agents are autonomous and function without prescriptive and continuous human intervention, there must be human oversight of their performance and remedial intervention when necessary — just as with the work of humans.*

To mitigate these risks and ensure responsible deployment of agentic AI, robust data governance practices and regular audits and reviews of AI agents and their decisions are essential. It may help to view AI agents as "digital employees" to be supervised in a manner similar to human employees. This means:

- Adopting a customer focus in defining the job and performance expectations of the AI agent
- Setting policies articulating the core values and guardrails for the agent's behaviour
- Deploying metrics and measures to facilitate monitoring of the AI agent's performance against expectations
- Taking remedial action when necessary and ensuring the AI agent can continuously learn and improve

[3] "Today's CEOs are the last to manage all-human workforces, says Marc Benioff," by Anna Cooban, CNN Business, January 23, 2025: www.cnn.com/2025/01/23/business/davos-marc-benioff-salesforce-ai-prediction-intl/index.html.

While AI agents are autonomous and function without prescriptive and continuous human intervention, there must be human oversight of their performance and remedial intervention when necessary — just as with the work of humans. Thus, AI agents become an integral part of the workforce.

The above discussion is relevant to board members because the journey from traditional automation to agentic AI represents a significant and transformative technological evolution. Agentic AI promises more intelligent, efficient and autonomous business processes, paving the way for a future where technology and human ingenuity work together to achieve remarkable outcomes. But it also comes with challenges such as ensuring ethical use, maintaining data privacy and managing employee transitions. Businesses must navigate these challenges to realise the full potential of agentic AI.

Questions Directors Should Ask

When engaged in strategic conversations with management regarding the deployment of agentic AI, the board should consider the following questions:

- How are we preparing for potential future developments in AI technologies and adapting them to our business as our strategy evolves?
- What is our long-term vision for deploying agentic AI in our organisation, and how does its deployment fit within our organisation's overall AI strategy?
  - What specific business problems or opportunities are we addressing with agentic AI?
  - Conversely, are the business problems or opportunities we are trying to address with agentic AI a good fit for that technology?
- How will the deployment of agentic AI affect the customer experience? What do we seek to achieve in the near term with these systems? And how will these systems — and the vendors supporting them — be integrated with our existing processes and technologies to ensure a seamless fit, scalability and ease of use?
- Do we have the necessary talent and expertise within our organisation to design, develop and manage agentic AI systems effectively? What training and development programs are in place to reskill and upskill our workforce to ensure we can fully realise the value proposition underpinning our AI investments?

- What governance structures are in place to oversee the responsible deployment and use of agentic AI within the organisation?
  - Are our AI agents thoroughly trained to identify and address potential issues before deployment, ensuring they meet customer expectations consistent with their stated objectives? What robust data collection and management systems are we using to train and operate them?
  - What measures are in place to protect sensitive data from breaches or misuse attributable to AI agents? Are we satisfied that we remain in compliance with data protection regulations and required safeguards of sensitive data? How do we know?
  - Have we assessed the potential risks associated with agentic AI deployment, including the risk of potential bias and discrimination and other operational, ethical and reputational risks? How do we plan to mitigate these risks, and what safeguards are in place to monitor AI performance and behaviour?
  - Who is accountable for the performance and learning of AI agents and ensuring their continued relevance to changing customer needs and compliance with internal policies and applicable laws and regulations? Given the velocity and pervasive impact of AI agents, how do we ensure that rogue behaviour is identified timely to minimise harm?
  - How will we measure success? That is, what metrics will we use to monitor the success and effectiveness of AI agents in fulfilling their assigned roles as well as in assessing their impact on our business?
  - Are there clear guidelines for when and how human employees should intervene to provide a "safety net" for more complex or sensitive interactions?
- How are we communicating with customers, employees, regulators and other stakeholders about our agentic AI deployments, including their strategic purpose and expected impact? What feedback mechanisms do we have in place?

These questions will help directors obtain a clearer view of the implications of deploying agentic AI. They will support the board's oversight to ensure management is undertaking a strategic, customer-focused and responsible approach to implementing AI agents.

How Protiviti Can Help

AI is rapidly changing the way we do business. Across all industries, from technology to healthcare, financial services and consumer products, organisations are adopting AI, intelligent automation and advanced analytics to improve processes, drive new business opportunities and increase competitive advantage.

As their organisations make this transition, we help boards and senior executives in many ways, which are tailored to their needs, including, but not limited to:

- Running board education sessions
- Leading facilitation and design thinking sessions to help the business prioritise opportunities to use AI
- Defining an overall AI strategy and executing and building against that strategy
- Reviewing and advising on AI governance and policies
- Providing assistance with implementing data-driven solutions that improve customer experiences and increase operational efficiency, speed and reliability
- Helping identify and manage third-party risk and compliance requirements

We have also helped large organisations evaluate their generative AI platforms from end to end (technology, people and processes) to enable deployment with confidence.

About the Author

Christine Livingston
Managing Director, Global Leader, Artificial Intelligence, Protiviti

Christine is responsible for Protiviti's AI/ML capabilities and solutions. With over a decade of experience in AI/ML deployment, she has delivered hundreds of successful solutions, including many first-in-class AI-enabled applications. She has helped several *Fortune* 500 clients develop practical strategies for value-driven enterprise adoption of artificial intelligence, including the creation of capability-based AI-enabled technology road maps. She focuses on incorporating AI/ML capabilities into enterprise solutions, and delivering tangible business value and outcomes through AI.

Contact Christine at christine.livingston@protiviti.com.