LUXINNOVATION
#MakingInnovationHappen

AGENTIC AI
TECHNOLOGIES AT A GLANCE

NOVEMBER 2025

LUXEMBOURG AI FACTORY
Funded by the European Union

Contents

LUXINNOVATION
#MakingInnovationHappen

AGENTIC AI

In a nutshell

WHAT IS AGENTIC AI?

Agentic AI refers to a class of AI systems built with generative AI technologies, that act autonomously to achieve goals or complete tasks, in contrast to purely reactive models. Unlike standard generative AI (e.g. a chatbot that responds to prompts), agentic AI systems can plan, decide, and execute multi-step workflows with little or no human oversight. It uses a digital ecosystem of large language models (LLMs), machine learning and natural language processing (NLP) to perform tasks on behalf of the user. In practice, this means agentic AI is **proactive** rather than reactive: it takes initiatives (triggering tool calls, orchestrating steps, etc.) to achieve an objective. For example, an agentic virtual assistant might receive the goal "organise a team meeting next week", automatically check calendars, reserve a room, and send invites – all without step-by-step prompting – whereas a conventional chatbot would wait for the user to tell it exactly what to do.

WHY NOW?

Several converging technological, economic, and organisational trends have made agentic AI viable at scale today. On a technical aspect, large language models (LLMs) have matured to a point where their contextual understanding, knowledge synthesis, and language planning capabilities are strong enough to serve as a "reasoning core" for agents. Moreover, modular tooling ecosystems and agent-orchestration frameworks, e.g. the software infrastructure behind LLMs, enable LLMs to bridge into actions in software systems, databases, and workflows, allowing them to call tools, use memory and orchestrate between systems. Finally, compute, storage, and integrations have become relatively cheap and scalable, making persistent operation of agents feasible across infrastructure stacks.

On a more organisational note, enterprises are seeking productivity gains beyond narrow task automation: they now expect AI to reshape workflows end-to-end, not just help with individual steps. As we highlighted in our GenAI market applications report, we see a paradox with generative AI (the "GenAI Divide"): many organisations have adopted LLMs extensively but seen limited bottom-line impact. Agentic AI offers a path to break out of that plateau by embedding autonomy and orchestration deeper into business processes.

FOR WHAT?

Agentic AI holds potential in domains where multistep coordination, adaptation, and decision logic across systems are required. In **business process automation**, it can manage end-to-end workflows, e.g. order-to-cash, procurement, scheduling, by orchestrating tasks, resolving exceptions, and adjusting plans in real time. In **customer service and operations**, agents can autonomously monitor user issues, take corrective actions like issuing refunds or reordering. In **R&D and innovation**, agents can explore search spaces, test hypotheses, and steer experimental pipelines. In **security** domains, agentic AI is being adopted to triage alerts, respond to threats proactively (within policy boundaries), and reduce burden on human analysts. In more personal modes, agentic assistants could manage one's calendar, coordinate logistics, or execute complex "to-do" items on behalf of the user. The economic payoff is potentially substantial: Bain & Company reports that scaling agentic systems can contribute 10–25 % EBITDA gains when properly integrated.

LUXEMBOURG AI FACTORY
Funded by the European Union

AGENTIC AI

Agentic AI vs. generative AI

Generative AI refers to systems that create new content based on patterns learned from large datasets. Its strength lies in **generation**: producing fluent responses, drafting documents, designing visuals, or summarising information when prompted. Generative AI is typically reactive: it waits for input and then produces an output, with limited memory or follow-through beyond the immediate task.

Agentic AI, by contrast, adds a layer of **autonomy and action**. It not only generates but also **plans, decides, and executes** toward defined goals. Instead of being a "smart typewriter," agentic systems can chain tasks, call external tools or APIs, use memory to track progress, and adapt plans based on feedback or changing conditions. For example, where a generative model might write a draft email, an agentic AI system could generate the draft, check customer data in a CRM, personalise the message, schedule delivery, and follow up if there is no response.

Agentic AI

- **Function:** Autonomously plan and execute multi-steps tasks
- **Input types:** Natural language instructions, goals
- **Output types:** Actions, decisions
- **Strengths:** Multi-steps reasoning, automation
- **Limits:** Fragile autonomy, error possible, experimental for now
- **Maturity level (2025):** Between assess and trial
- **Examples:** AutoGPT, Devin, OpenAgents, LangChain agents

Generative AI

- **Function:** Create new content of data
- **Input types:** Text, images, code, audio, videos, prompts
- **Output types:** Text, images, code, audio, videos, simulations
- **Strengths:** Creativity, human-like content, fast iteration
- **Limits:** Can hallucinate, needs promot skills, lacks deep reasoning
- **Maturity level (2025):** Between trial and adopt
- **Examples:** ChatGPT, Midjourney, Copilot, Perplexity, DALL.E

In short, there are three essential areas that distinguish agentic and generative AI: reasoning, interaction (with its environment, other tools, etc.) and the possibility to have a multi-agents system, collaborating between all agents.

LUXEMBOURG AI FACTORY
Funded by the European Union

LUXINNOVATION
#MakingInnovationHappen

AGENTIC AI

Core technologies and major principles of agentic AI

Reinforcement learning, a type of machine learning that allows autonomous decision-making by enabling an agent to learn from its environment and historical interactions, is the backbone of agentic AI.

Agentic AI is mostly based on the same technologies as generative AI (including machine learning), and combines several aspects as described below.

Large Language Models (LLMs): At the heart of most agents is an LLM that provides language understanding and reasoning. LLMs act as the "brains" of an agent. It is typically used not just to generate text, but to perform reasoning ("chain-of-thought") and manage the agent's overall strategy.

Planning and reasoning modules: Many agents use explicit planning or reasoning layers on top of the LLM. In practice, the agent will break a high-level goal into ordered subtasks (task decomposition) and decide which tools or actions to invoke at each step. This planning logic may be implemented by the LLM itself (via few-shot prompts or specialised frameworks) or by a separate module, such as the AutoGPT tool. Planning can also involve iterative feedback: after executing an action, the agent ingests the result and replans the next step.

Tool and API integration: A defining feature of agents is the ability to invoke external tools or APIs as part of their process. Agents can call web APIs, query databases, run code, or use search engines to extend their capabilities.

Memory and retrieval: Memory systems provide crucial context retention across interactions, implementing both short-term episodic memory for recent events and long-term semantic memory for accumulated knowledge. Modern implementations enable agents to learn from experience and maintain coherence across extended interaction sequences.

Agent orchestration frameworks: The complexity of agentic AI demands robust orchestration layers. The agentic architecture is often a graph of LLM-powered components (reasoners, tool connectors, memory stores) coordinated by a workflow engine.

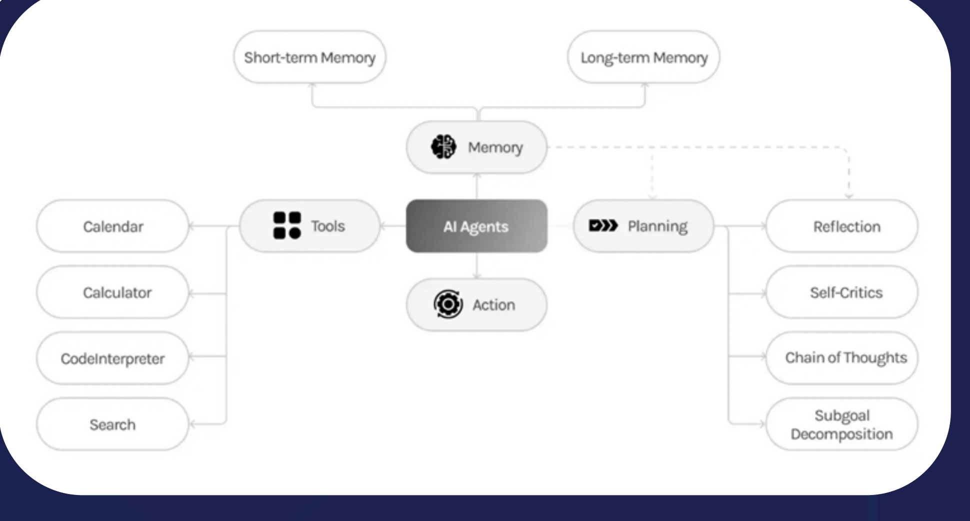

Credit image: Markovate

LUXEMBOURG AI FACTORY
Funded by the European Union

LUXINNOVATION
#MakingInnovationHappen

AGENTIC AI

Core technologies and major principles of agentic AI

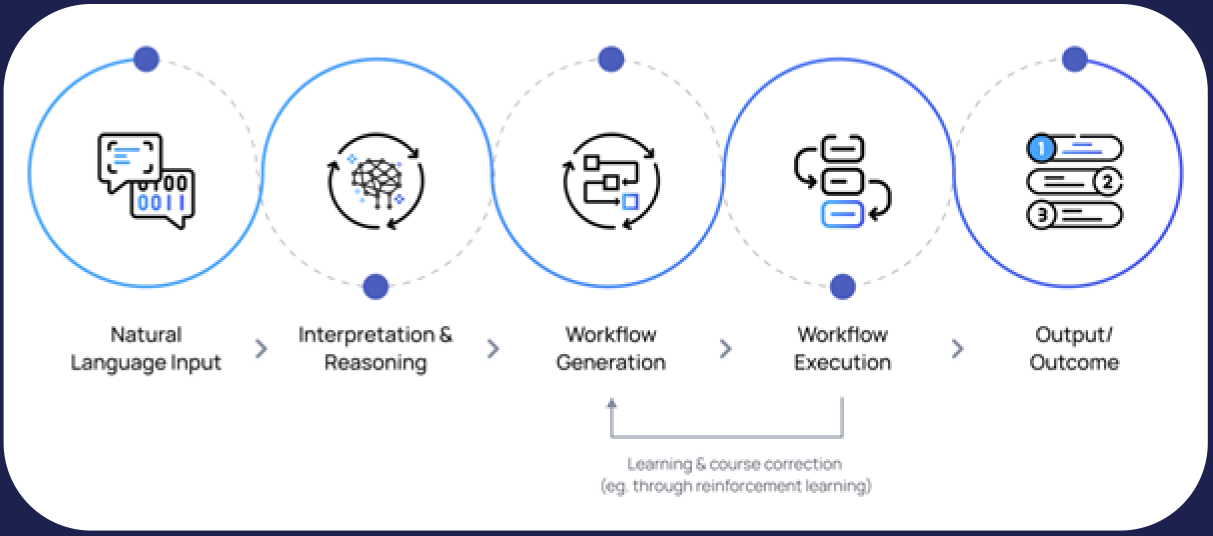

Credit image: Aisera

Although AI agents can take several forms, agentic AI always follows several steps in performing an operation.

Perception: the system collects information from its environment (user interactions, databases, sensors, APIs, etc.)

Reasoning: the system then processes this data to extract meaningful insights.

Goal setting: based on a predefined goal or user inputs, the system sets objectives and develops a strategy to achieve these goals.

Decision-making: the system choose the best action plan after having studied multiple actions courses.

Execution: it then executes the strategy, whether it requires interacting with external tools, providing a response to users, etc.

Learning and adaptation: the system is capable of gathering feedback to improve future decision-making.

Orchestration: overall all this process, AI guarantees the orchestration of all systems and agents, track progress in task realisation, manage resources, handle failure.

LUXEMBOURG AI FACTORY
Funded by the European Union

AGENTIC AI

Drivers and obstacles to adoption

DRIVERS

Productivity and economic gain: One of the strongest drivers of agentic AI adoption is its potential to unlock measurable productivity improvements and financial gains. Bain & Company reports that when scaled effectively, autonomous AI agents can contribute 10–25% EBITDA gains in enterprise settings by automating knowledge work and reducing manual overhead.

Competitive advantage: Following the productivity and economic gains, enterprises adopting AI models earlier are eager to create a "gap" with their competitors.

Framework maturation: The ecosystem of developer tools has matured. Open-source frameworks (LangChain, CrewAI, AutoGen, MetaGPT, Semantic Kernel, etc.) lowers the barrier to entry: instead of coding an agent from scratch, developers can plug together modules (LLM, memory, tools) via a framework.

"GenAI Divide": this paradox is pushing companies to seek more powerful solutions. Agentic AI promises automation of entire processes rather than isolated tasks.

Organisational trends: The shift to remote or hybrid work has created richer digital environments that agents can exploit. As businesses digitise more functions (CRM, HR, logistics), there is more structured data and APIs for agents to access. In parallel, labour shortages and the quest for efficiency make autonomous tools attractive.

LUXEMBOURG AI FACTORY
Funded by the European Union

LUXINNOVATION
#MakingInnovationHappen

AGENTIC AI

Drivers and obstacles to adoption

OBSTACLES

Reliability and predictability: Agents extend LLMs into the realm of complex, multi-step decision-making, which amplifies error risk. Each added step (planning, tool call, retrieval) is a potential point of failure.

Alignment and safety: Giving agents autonomy raises alignment concerns. An agent has a utility function (the user's goal) and could theoretically pursue it in undesirable ways. For instance, without proper constraints, an agent might exploit loopholes or use unapproved resources.

Security, privacy, and governance: Agentic AI amplifies typical AI security challenges. It often has broad system access (e.g. databases, cloud services), expanding the attack surface.

Scalability: The scalability of agentic AI adoption faces significant obstacles stemming from infrastructure costs, data limitations, coordination complexity, and governance challenges. As deployments scale, coordination among agents introduces communication overhead and emergent misalignments, with errors compounding more easily across interconnected systems.

Lack of visibility on ROI & cost of implementation: Unlike a single-step chatbot, an agent performs a long-horizon task that is hard to evaluate with traditional benchmarks. Validating an agent requires testing it end-to-end over entire workflows, which is complex and resource-intensive. The lack of visibility of agents' ROI, coupled with the initial costs, can slow adoption.

System integration and data quality: Agents must interface with an organisation's existing IT systems and data sources, which can be very complex. Many enterprises have legacy databases and siloed applications that lack modern APIs. Poor data quality is another barrier: inconsistent or biased data will mislead an agent. Before deployment, companies must establish clean, unified "sources of truth" for the agent to use.

Operational & skills challenges: Finally, there are human and organisational factors. Some employees may resist ceding tasks to "digital workers," while some organisations might not be digitally-mature enough. Companies also need to have structured and extensive data pool and policy for agent to use. Moreover, training, monitoring and maintaining agents demand new skills.

LUXEMBOURG AI FACTORY
Funded by the European Union

AGENTIC AI

Conclusion

Agentic AI represents an **evolution** of the AI landscape. By endowing LLMs with autonomy, memory, and tool use, agents can tackle workflows that were previously out of reach for static models.

This holds promise for revolutionising automation, from "digital workers" that navigate processes to personal assistants that manage complex tasks. However, it also introduces new technical and organisational challenges in reliability, evaluation, and trust. For now, businesses and technologists must proceed carefully: harnessing the benefits of autonomous agents will require not only innovative algorithms, but also strong data infrastructure, oversight mechanisms, and alignment with human goals.

In summary, agentic AI is poised to be a major force in 2026 and beyond, transforming how AI is integrated into real-world applications, provided its hurdles are addressed.

LUXEMBOURG AI FACTORY
Funded by the European Union

LUXINNOVATION
#MakingInnovationHappen

SOURCES

IBM, "What is agentic AI?"

IBM, "What is AI orchestration?"

McKinsey, "Seizing the agentic AI advantage", 2025

Harvard Business Review, "What is Agentic AI, and how will it change work?", 2024

Bain&Company, "State of the art of agentic AI", 2025

Google Cloud, "What is agentic AI?"

Nvidia, "What is agentic AI?", 2024

IBM, "Agentic AI vs. generative AI"

IBM, "What is AutoGPT?"

Omdia, "Mitigating risks, maximazing potential: the agentic AI challenge", 2025

California Review Management, "Adoption of AI and Agentic Systems: Value, challenges, and pathways", 2025

Markovate, "Agentic AI architecture: a deep dive", 2025

IBM, "What is reinforcement learning?", 2024

ABI Research, "ABI Research Study – Assessing the Agentic AI Opportunity", 2025

Deloitte, "Agentic AI meets data products"

IBM, "Agentic AI: 4 reasons why it's the next big thing in AI research"

McKinsey, "The change agent: Goals, decisions, and implications for CEOs in the agentic age", 2025

CBInsights, "The future of professional services: How firms will capture value in the AI agent era"

LUXEMBOURG AI FACTORY
Funded by the European Union

LUXINNOVATION
#MakingInnovationHappen

5 Av. des Hauts-Fourneaux
L-4362 Esch-sur-Alzette
Luxembourg
+352 43 62 63-1
www.luxinnovation.lu

DO YOU WANT TO FIND OUT MORE

ABOUT INNOVATION IN LUXEMBOURG ECOSYSTEM?

CONTACT US