Agentic AI for industrial applications

Transformation through autonomous intelligence

Proud to be part of CATAPULT High Value Manufacturing

mtc where progress happens

Acknowledgements

Ula Hijawi
Advanced Research Engineer, Digital Engineering, The Manufacturing Technology Centre Ltd

Mohammed Begg
Advanced Research Engineer, Digital Engineering, The Manufacturing Technology Centre Ltd

Dr. Mostafizur Rahman
Chief Technologist – Artificial Intelligence, Engineering Office, The Manufacturing Technology Centre Ltd

Tom Winter
Technology Manager, Digital Engineering, The Manufacturing Technology Centre Ltd

Sean Culey
Director of Supply Chain, Supply Chain, The Manufacturing Technology Centre Ltd

Evangelos Chatzivagiannis
Senior Research Engineer, High Value Design, The Manufacturing Technology Centre Ltd

Contents

Table of figures

List of tables

Executive summary

Agentic AI

Agentic Artificial Intelligence (AI) refers to advanced AI systems characterised by autonomous operation, contextual reasoning, and goal-directed behaviour. Unlike traditionally trained AI, Agentic AI systems exhibit capabilities such as planning, decision-making under uncertainty, and adaptive interaction with dynamic environments. Powered by Large Language Models (LLMs), these agents can interpret complex tasks, generate internal Chain-of-Thought (CoT), and execute actions independently. The agentic paradigm emphasises reasoning and intentionality, self-directed problem solving, and interactive collaboration, distinguishing it from reactive or purely supervised AI systems.

Core capabilities of Agentic AI systems

Agentic AI systems are characterised by three foundational attributes: ***autonomy***, enabling them to operate with minimal human intervention; ***collaboration***, allowing them to work cooperatively with humans and other systems; and ***adaptability***, which empowers them to learn and evolve within dynamic, real-world environments. These systems function through a cyclic workflow comprising six main interconnected stages.

1. ***Perception***, where the system senses and interprets multi-modal data from its environment.

2. ***Reasoning,*** which involves generating chains of thought, making decisions, and solving problems based on available information.

3. ***Planning***, where the system sets goals and determines the optimal sequence of actions to achieve them strategically.

4. ***Execution***, which enables the system to autonomously implement these actions in real or virtual environments.

5. ***Learning***, where the system continuously improves its performance by acquiring new knowledge and adapting to changes.

6. ***Interaction,*** ensures effective communication and collaboration with humans and other systems, completing the cycle and enabling continuous refinement of capabilities.

Strategic importance and identifying high-impact Agentic AI use-cases in industry

Despite its demonstrated value in operational domains such as finance, sales, and customer support, the potential of Agentic AI remains largely unutilised in industrial sectors like advanced manufacturing. However, it holds significant promise for transforming shopfloor operations, enabling more intelligent, adaptive, and autonomous decision-making. To fully unlock the potential of Agentic AI, opportunities can be identified especially in scenarios requiring real-time decision-making; complex, repetitive, or time-intensive process automation; cross-team coordination; and continuous adaptation to changing conditions. The greatest impact often occurs where Agentic AI can outperform traditional approaches in efficiency and scalability.

High-impact industrial use-cases of Agentic AI can be typically captured in areas demanding high reasoning loads, decision-making, and interpretation of ambiguous data, as well as complex multi-step or multi-system workflows that require orchestration across tools and departments. They also include dynamic or non-deterministic environments, where conditions change frequently (e.g., adaptive maintenance, real-time quality control) and require flexible responses. Agentic AI also proves valuable in reducing bottlenecks caused by frequent human intervention, managing dynamic environments that require flexible responses, and interpreting data-rich processes that currently lack intelligent analysis. Additional potential lies in processes with feedback loops that enable continuous learning and improvement, and in scenarios requiring multidisciplinary collaboration to bridge gaps between design, production, and operations.

Practical use-cases of Agentic AI in industry

Practical applications of Agentic AI in industry span several high-impact areas. Advanced Agentic AI architectures, such as Multi-Agent Systems (MAS), are particularly well-suited for these use-cases, as they comprise multiple autonomous agents, each with specialised capabilities, working collaboratively toward a common objective. For instance, in production process optimisation, AI agents can autonomously identify an optimal set of parameters, reducing reliance on time-consuming methods and costly cross-departmental coordination. For predictive maintenance, hierarchical AI agents can continuously monitor equipment, detect anomalies, and coordinate maintenance tasks, inventory checks, and supply chain actions to minimise downtime. In supply chain management, a MAS can enhance real-time visibility, streamline decision-making, and proactively resolve bottlenecks across planning, sourcing, manufacturing, and delivery stages. Finally, in design optimisation and compliance, AI agents can accelerate iterative design cycles by generating and evaluating alternatives against performance and regulatory standards, ensuring efficiency, accuracy, and compliance throughout the process.

This whitepaper expands each of these key industrial use-cases by exploring the underlying challenges, outlining Agentic AI-driven solutions, highlighting their potential impact, and presenting practical workflows for integrating Agentic AI into existing processes.

Challenges and recommendations

As Agentic AI systems become more integrated into industrial and organisational workflows, several critical challenges have been identified in existing Agentic AI approaches that may hinder their safe, effective, and trustworthy deployment. These include risks of losing human control and trust, misalignment with user goals, and limitations in planning quality and reliability. Most importantly, governance frameworks are critical to ensure ethical use and maintain human oversight. Addressing these gaps requires robust strategies, including the strategic orchestration of multiple specialised agents for scalability and resilience. For instance, Small Language Models (SLMs), smaller in scale and scope than LLMs, can be effectively deployed within MASs in hybrid collaboration with LLMs. Leveraging domain-specific SLMs enhances efficiency and accuracy while significantly reducing computational and energy consumption costs. Other recommendations include ensuring goal alignment, and embedding transparency, explainability, and ethical safeguards into the system design. Simulation environments and stress testing can be facilitated to evaluate behaviour under edge cases for trust and safety. Moreover, interdisciplinary collaboration among technologists, ethicists, and domain experts is essential to build systems that are safe, interpretable, and aligned with human values.

What this whitepaper provides

While Agentic AI has already proven its value in operational domains such as finance, sales, and customer support, its transformative potential in industry and manufacturing environments remains largely untapped. This whitepaper presents a forward-looking vision for Agentic AI in industry. It provides a comprehensive guide to understanding and implementing Agentic AI in industrial contexts, and highlights practical limitations, integration challenges, and the need for human oversight. The paper begins by clarifying key terminology and conceptual boundaries, distinguishing Agentic AI from related concepts such as AI agents and LLMs, and introducing architectural foundations, including the emerging role of SLMs as specialised agents within MASs. It then offers an overview of Agentic AI systems' core capabilities, benefits, adoption trends, and operational workflow. Practical industrial use-cases in advanced manufacturing are explored, followed by a step-by-step guide for integrating Agentic AI into existing workflows. Finally, the paper presents strategic recommendations for adoption and offers a forward-looking perspective on future developments in this rapidly evolving field. A follow-up publication will extend this whitepaper's scope by providing technical insights into Agentic AI, covering agent types, architectures, and frameworks, tailored specifically for AI developers.

1. Introduction

Large Language Models (LLMs), AI Agents, and Agentic AI

Large Language Models (LLMs) have transformed how machines understand and generate human language, but they remain fundamentally reactive tools. They respond to prompts without memory, autonomy, or long-term goals. Agentic AI systems build on LLM capabilities by embedding them within frameworks that provide memory, planning, and goal-directed behaviour. These systems can reason, make decisions, and execute multi-step tasks over time with minimal human input, shifting from static responders to dynamic collaborators capable of adapting to changing objectives and complex workflows. Comprehensive surveys covering the foundations, architectures, and applications of Agentic AI are presented in [1], [2] and [3]. While the terms AI Agents and Agentic AI are often used interchangeably, they differ in scope. AI Agents can be considered as individual entities designed to perform specific tasks. Agentic AI, on the other hand, introduces a higher level of autonomy, adaptability, and coordination, often orchestrating multiple agents and tools to achieve broader objectives. Table 1 compares the key functionalities of LLMs, AI Agents, and Agentic AI, whereas *Figure 1* illustrates the conceptual differences in their respective workflows.

Table 1. A comparative analysis highlighting the distinct capabilities of LLMs, AI Agents, and Agentic AI.

Functionality	LLMs	AI agents	Agentic AI systems
Description	Large Language Models trained on vast text corpora to generate human-like text, code, or answers. They excel at pattern recognition and language understanding but lack inherent goal orientation.	Software programs designed to perform specific tasks autonomously or semi-autonomously. They often wrap around models like LLMs and include logic for decision-making and action execution.	Advanced, autonomous systems that coordinate multiple agents and tools to achieve complex, high-level goals. They integrate reasoning, planning, and execution across dynamic environments.
Scope	**Stateless and prompt-driven.** Respond to individual prompts without persistent memory or goals. Can simulate reasoning but lack structured goal pursuit or long-term planning.	**Narrow and task-specific.** Typically focused on a single domain or workflow (e.g., booking, data retrieval). May include short-term memory for context but limited adaptability beyond their scope.	**Broad, goal-oriented, and adaptive.** Maintain short and long-term memory, track progress, and dynamically adjust strategies. Capable of decomposing complex goals, reasoning internally, and executing multi-step plans over time.
Autonomy	**Passive and reactive.** No autonomy; they only produce outputs when prompted. Cannot initiate actions or make decisions independently.	**Semi-autonomous.** Can perform predefined actions without constant human input but operate within strict boundaries and guidelines.	**Highly autonomous.** Operate with defined objectives, make decisions, and act without continuous human oversight. Capable of self-initiating tasks and adapting to changing conditions.
Learning	**Limited or no real-time learning.** Knowledge is fixed post-training; improvements require retraining or fine-tuning.	**Typically, no real-time learning.** Most agents rely on predefined logic. Some can adjust parameters or preferences within a session (e.g., using short-term memory or user feedback), but they do not fundamentally learn over time.	**Continuous learning and adaptable.** Incorporate feedback, update strategies, and improve performance over time, often using reinforcement learning or similar techniques.

Coordination	**Standalone models.** Generate text, code, or other multi-modal data, but do not interact with external systems unless explicitly integrated.	**Independent operators.** May use APIs or tools but typically act alone without orchestrating other agents.	**Collaborative orchestrators.** Integrate with multiple agents, tools, and environments (e.g., databases, machines, human operators). Capable of multi-agent coordination and resource allocation.
Examples	Chatbots, code assistants, content generators.	Virtual assistants, Robotic Process Automation (RPA) bots, scheduling tools.	Enterprise AI platforms, autonomous research systems, multi-agent orchestration frameworks.

As agent frameworks evolve, the distinction between "AI Agent" and "Agentic AI" may blur in practice. For organisations, a practical approach is to assess the degree of autonomy, adaptability, and orchestration required for their use-case, using the taxonomy above as a guide.

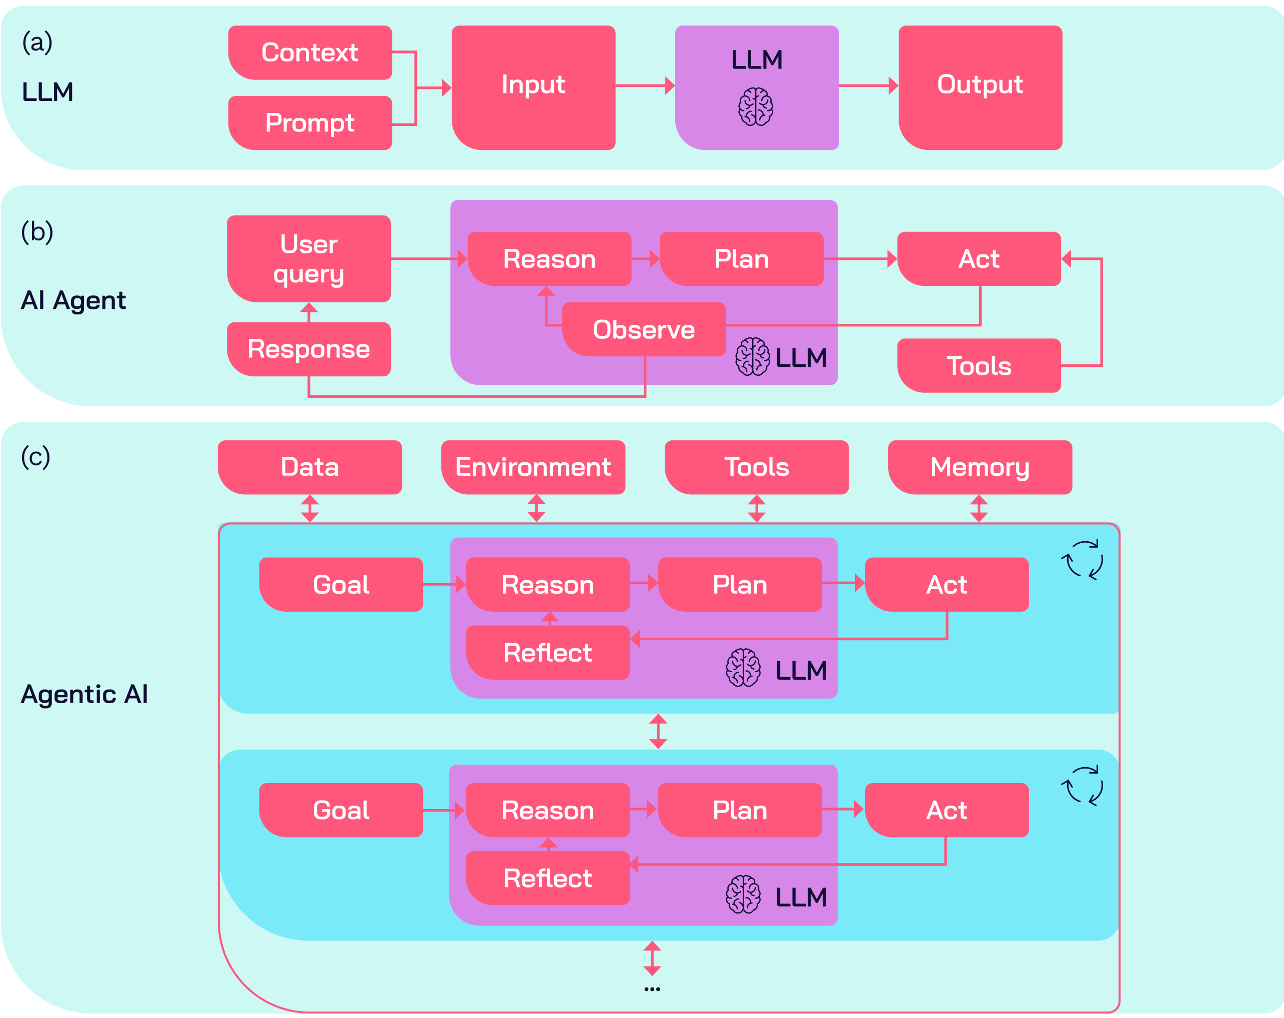

Figure 1. An illustration showing the typical workflows of LLMs (a), AI Agents (b), and Agentic AI systems (c).

Multi-agent orchestration and the role of specialised, Small Language Models (SLMs)

A Single-Agent System (SAS) uses one LLM-powered agent to autonomously manage an entire task, from interpreting user prompts to executing tools, making it ideal for focused, domain-specific workflows. In contrast, a Multi-Agent System (MAS) employs multiple specialised agents that collaborate to handle complex, large-scale, or multi-dimensional tasks, offering greater scalability, modularity, and adaptability. MAS architectures can be hierarchical or decentralised, enabling agents to coordinate and allocate resources efficiently. For instance, a hierarchical MAS architecture can consist of a supervisor of multiple agents, or a supervisor of multiple supervisors of agents, depending on the use-case complexity. Various forms of multi-agent orchestration architectures are illustrated in ***Figure 2*** [*4*] and [*5*]. Small Language Models (SLMs), on the other hand, are AI models designed to process and generate natural language but are smaller in scale and scope than LLMs. They can be tailored for specific domains, making them more effective than general-purpose LLMs in handling domain-specific tasks with greater accuracy. Within a MAS architecture, SLMs' can play a key role by providing cost-effective, fast-inference, and privacy-preserving solutions for lightweight tasks, especially on edge devices. A hybrid MAS can be designed combining LLMs as supervisory agents, or central orchestrators, for high-level reasoning with SLMs as edge agents for specialised execution, resulting in improved responsiveness, resource efficiency, and intelligent decision-making across distributed environments. For readers seeking technical depth, we recommend consulting established MAS frameworks (e.g., SPADE, JADE, Ray) for agent communication, orchestration, and conflict resolution strategies. Integration of SLMs with LLMs presents challenges in knowledge transfer and consistency, which will be explored in future technical publications.

(a) Centralised (Supervisor)

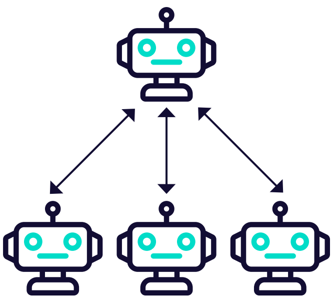

A single AI orchestrator plans, coordinates, and delegates tasks across specialised agents. Each agent executes assigned tasks using tailored tools and reports results back to the orchestrator. The design simplifies control but introduces a single point of failure and potential scalability limits.

(b) Hierarchical (Layered)

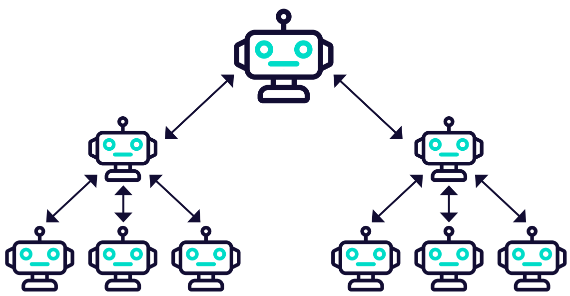

The orchestrator acts as a supervisor-of-supervisors, managing multiple supervisory agents. Each supervisor coordinates a group of subordinate agents, which communicate only with their assigned supervisor. This layered structure improves scalability and control clarity but can increase latency and complexity.

(c) Custom

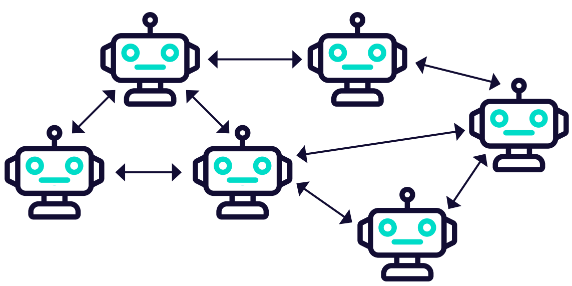

A custom architecture can be tailored for domain-specific requirements, where agents possess specified capabilities and interact only with defined subsets of peers. Control combines deterministic pathways with sensitive autonomy, allowing certain agents to dynamically choose interaction partners. This approach balances structure with flexibility for efficient coordination.

(d) Decentralised

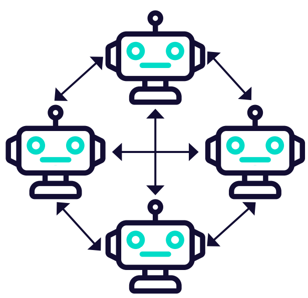

A fully connected, peer-to-peer system where agents operate autonomously and dynamically select interaction partners. This enables flexible coordination and distributed decision-making but can lead to high communication overhead and reliability challenges without robust consensus mechanisms.

Figure 2. Multi-agent orchestration architectures: (a) centralised, (b) hierarchal, (c) custom, and (d) decentralised [*4*] and [*5*]

2. The rise of Agentic AI in the enterprise: forecasting impact and adoption growth

AI agents and Agentic AI are emerging as the next key technology adoption for the enterprise. According to a survey presented by SS&C Blue Prism AI Trends [6], business leaders recognise the potential of Agentic AI, where:

29%

Of organisations are already leveraging Agentic AI for autonomous automation.

44%

Of organisations plan to implement it within the next 12 months.

According to a Bank of America (BofA) Global Research analysis, spending on Agentic AI could reach $155 billion by 2030 [7]. In a recent survey [8], executives say AI Agents will:

Increase automation in their workflows by **71%**

Improve customer service and satisfaction by **64%**

Improve potential productivity outweighing the risks by **57%**

Make daily work decisions by 2028 of **15%**

According to McKinsey's 2025 report [7], several case studies have captured the impact of deploying AI agents in the operational business domain. The resulted impact can be summarised as follows:

>50%

Reduction in time and effort in the early adopter teams for core system modernisation in banking

>60%

Potential productivity gain

>$3m

Annual expected savings

in enhancing market insights through data quality improvements

20-60%

Potential gain in productivity

30%

Faster decisioning speed

in credit-risk memo creation process at retail banks

3. Core capabilities of Agentic AI systems

Autonomy

Operate independently, executing tasks without constant human oversight. They exhibit self-directed behaviour, initiating actions aligned with predefined goals. This autonomy makes them reliable collaborators in dynamic environments.

Collaboration

Enable interaction between humans and machines. They support multi-agent coordination and can engage in structured dialogue to simulate digital teamwork. Through system interoperability, Agentic AI can integrate with other agents, platforms, or users, facilitating cooperative problem-solving and enhancing productivity across distributed workflows.

Adaptability

Handle dynamic environments through contextual learning and real-time adjustment. They incorporate feedback to refine their behaviour and improve performance over time.

Planning and reasoning

Deconstruct complex tasks into manageable steps. They use chain-of-thought generation and LLM-empowered reasoning to determine appropriate actions, make informed decisions and produce multi-step plans. This structured approach enables them to navigate ambiguity, prioritise actions, and achieve long-term goals.

Tool use and execution

Extend beyond language generation and bridge the gap between decision-making and execution by interacting with custom-built and external tools and systems. They can invoke APIs, execute code, and perform appropriate actions based on context. They adjust their approach based on evaluating outcomes.

Memory and knowledge

Maintain long-term memory and contextual awareness, allowing them to learn from experience and build on prior interactions. Through stateful operation and knowledge storage, they can recall relevant information, track progress over time, and apply learnt insights to future tasks.

4. Benefits of Agentic AI systems

Agentic AI systems can significantly transform industries across multiple dimensions including:

Increased productivity and efficiency

Continuously optimising workflows and adjusting operations without human intervention.

Repetitive tasks reductions

Automating routine tasks freeing up teams' time to focus more on creative, strategic, or human-centric work.

Better decision making

Generating Chain-of-Thought reasoning relying on data-driven logic rather than human intuition alone.

Autonomy

Automate complex business workflows by integrating various tasks and systems, triggering actions, and finding optimal solutions to complex problems.

Cost reductions

Lowering error rates and reducing labour costs by automating complex decision-making and operational tasks.

Personalisation

Offering full flexibility in creating tailored, custom Agentic workflows to address domain-specific needs in real time, integrating with existing or legacy systems.

Context-aware adaptability

Responds intelligently to changing environments, data, and user needs through continuous learning and feedback integration.

Enhanced collaboration

Facilitates smoother coordination and communication across siloed and teams, systems, and platforms, enabling cross-functional transformation.

Tool and system integration

Interfaces with APIs, databases, and software tools to execute tasks end-to-end, not just provide recommendations.

5. The Agentic AI workflow: how Agentic AI operates

The Agentic AI Workflow is illustrated in ***Figure 3*** showcasing the following key stages:

1. **Perception:** The ability to sense and interpret data from its environment. Includes multi-modal input processing to form an understanding of the context or current state.
2. **Reasoning:** The ability to drawing inferences, making judgments, or solving problems based on available information. Includes data analysis and pattern detection to make logical decisions.
3. **Planning:** The capability to set goals and determine a sequence of actions to achieve them. Includes adapting strategies to reach optimal outcomes.
4. **Execution:** The ability to act and execute planned actions in the real or virtual world autonomously. This includes custom tool integration and Retrieval Augmented Generation (RAG) to implement decisions and achieve objectives.
5. **Learning:** The ability to improve performance over time by acquiring new knowledge or adapting to changes. Includes supervised, unsupervised, or reinforcement-based learning, and the refinement of the Agentic AI system's models and behaviours.
6. **Interaction:** The means by which the Agentic AI system communicates and collaborates with humans or other systems. Includes NLP, dialogue management, and understanding user intent to ensure effective and meaningful exchanges.

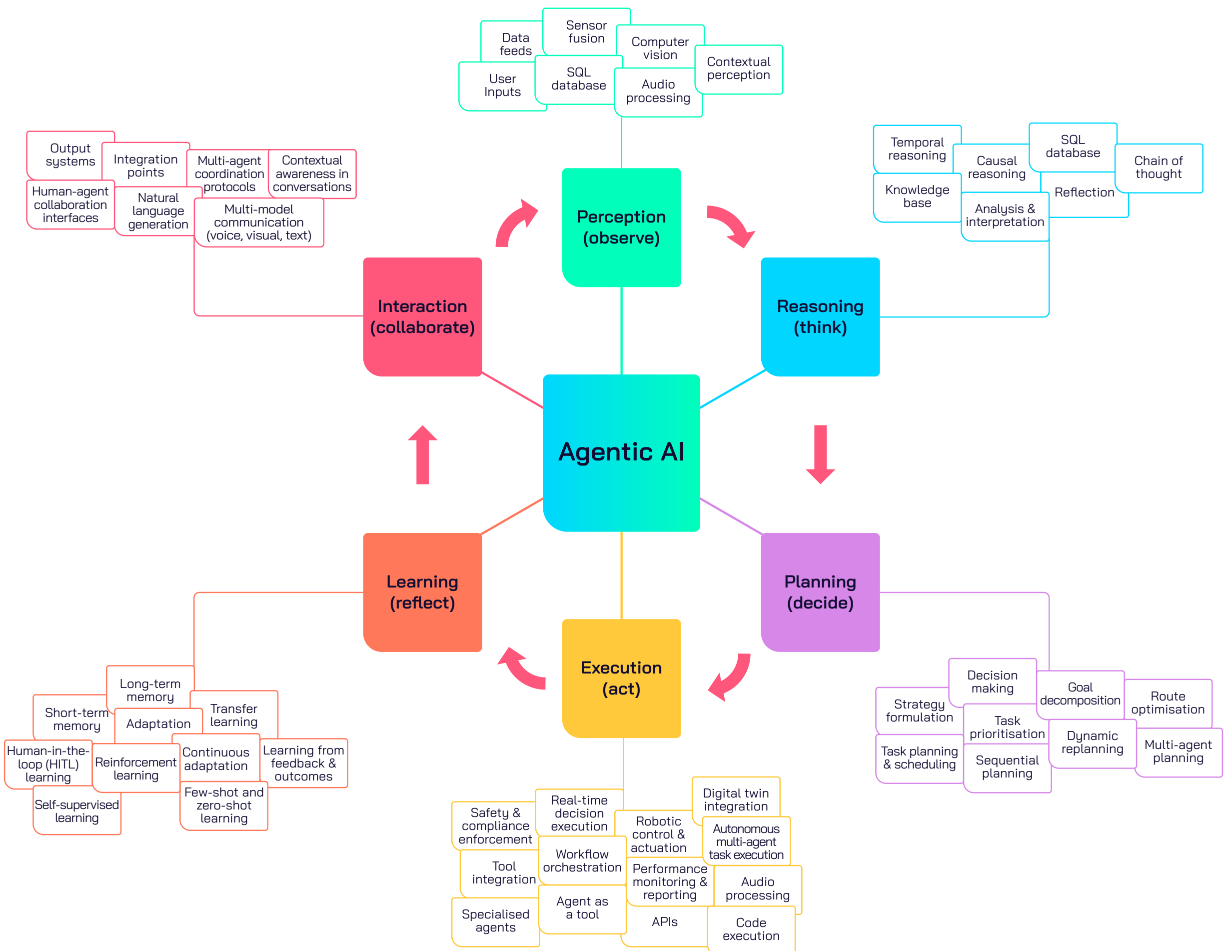

Figure 3. An Agentic AI system cyclic workflow throughout six stages in addition to the associated tools, techniques, or enablers used in each stage.

6. Unlocking the potential of Agentic AI in industrial applications

A new frontier for intelligent systems

While Agentic AI has already proven its value in operational domains such as finance, sales, and customer support, its transformative potential in industrial environments remains largely untapped. These domains where design, production, maintenance, and quality control intersect, are inherently complex. They present unique challenges that demand intelligent, adaptive systems capable of reasoning, collaboration, and autonomous execution. Unlocking Agentic AI in these settings could redefine productivity, resilience, and innovation on the shopfloor.

Integrating Agentic AI into industrial workflows

To realise this potential, industrial environments must embrace Agentic AI not just as a tool, but as a collaborative partner in problem-solving. These agents can assist in design iteration, monitor production anomalies in real time, optimise maintenance schedules, and even coordinate across distributed systems. By embedding reasoning, memory, and tool-use capabilities into AI agents, organisations can move beyond automation toward intelligent orchestration, where systems proactively adapt, learn from operational data, and contribute to continuous improvement. The key lies in aligning agentic capabilities with industrial workflows, enabling AI to augment human expertise rather than replace it. A high-level overview of key challenges in complex industrial environments where Agentic AI capabilities can enable high-impact use-cases is illustrated in ***Figure 4***.

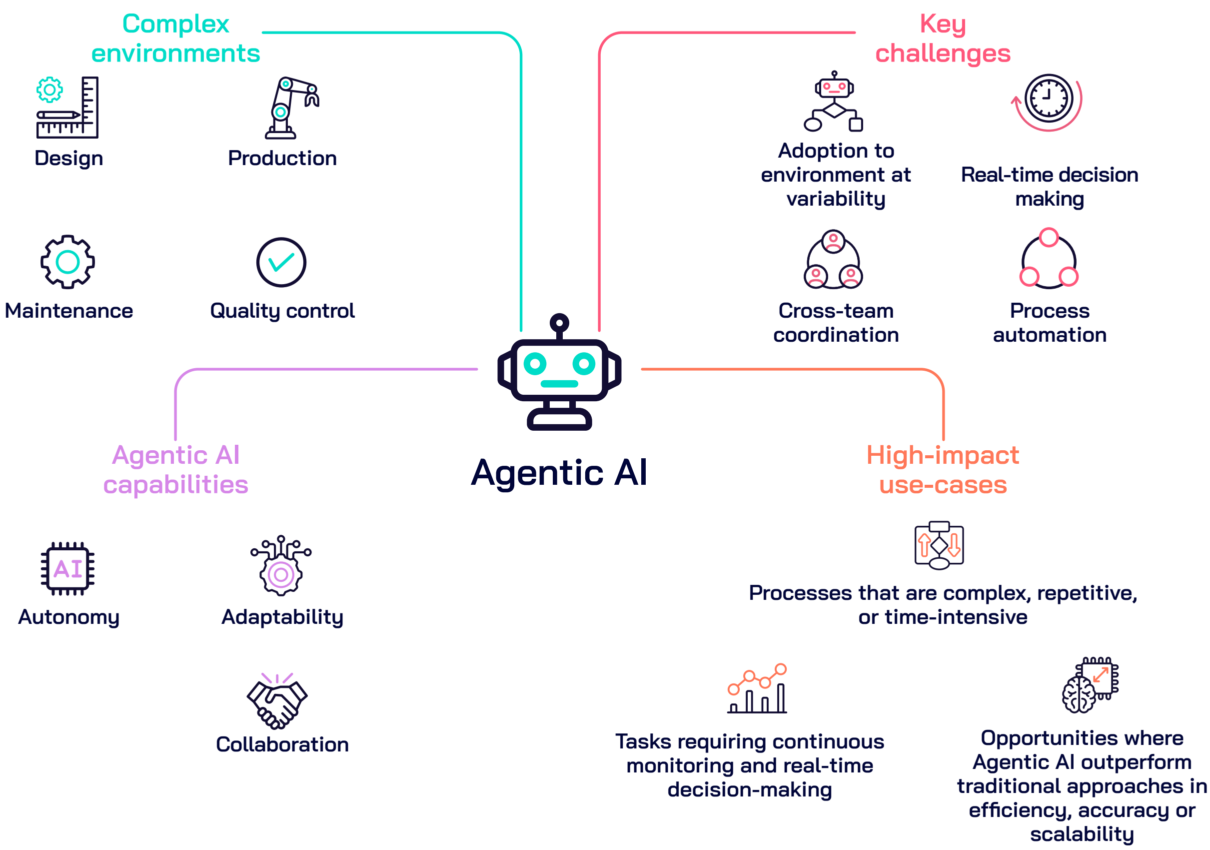

Figure 4. A conceptual overview of key challenges in complex industrial environments where Agentic AI capabilities can enable high-impact use-cases.

7. Identifying high-impact use-cases of Agentic AI systems for industrial applications

Identifying where Agentic AI can deliver the greatest value in industrial domains, such as manufacturing and engineering, requires more than just spotting complexity, it demands a strategic lens on where intelligence, autonomy, and adaptability can truly shift performance. High-impact use-cases often emerge where human expertise is stretched thin, where workflows span multiple systems, or where decisions must be made in real time under uncertainty. These are the areas where Agentic AI can act not just as an assistant, but as a proactive collaborator navigating ambiguity, learning from feedback, and orchestrating actions across tools and teams. By focusing on these criteria, organisations can prioritise deployments that not only automate but elevate industrial processes. Opportunities for identifying high-impact use-cases in industrial domains are summarised in ***Figure 5***.

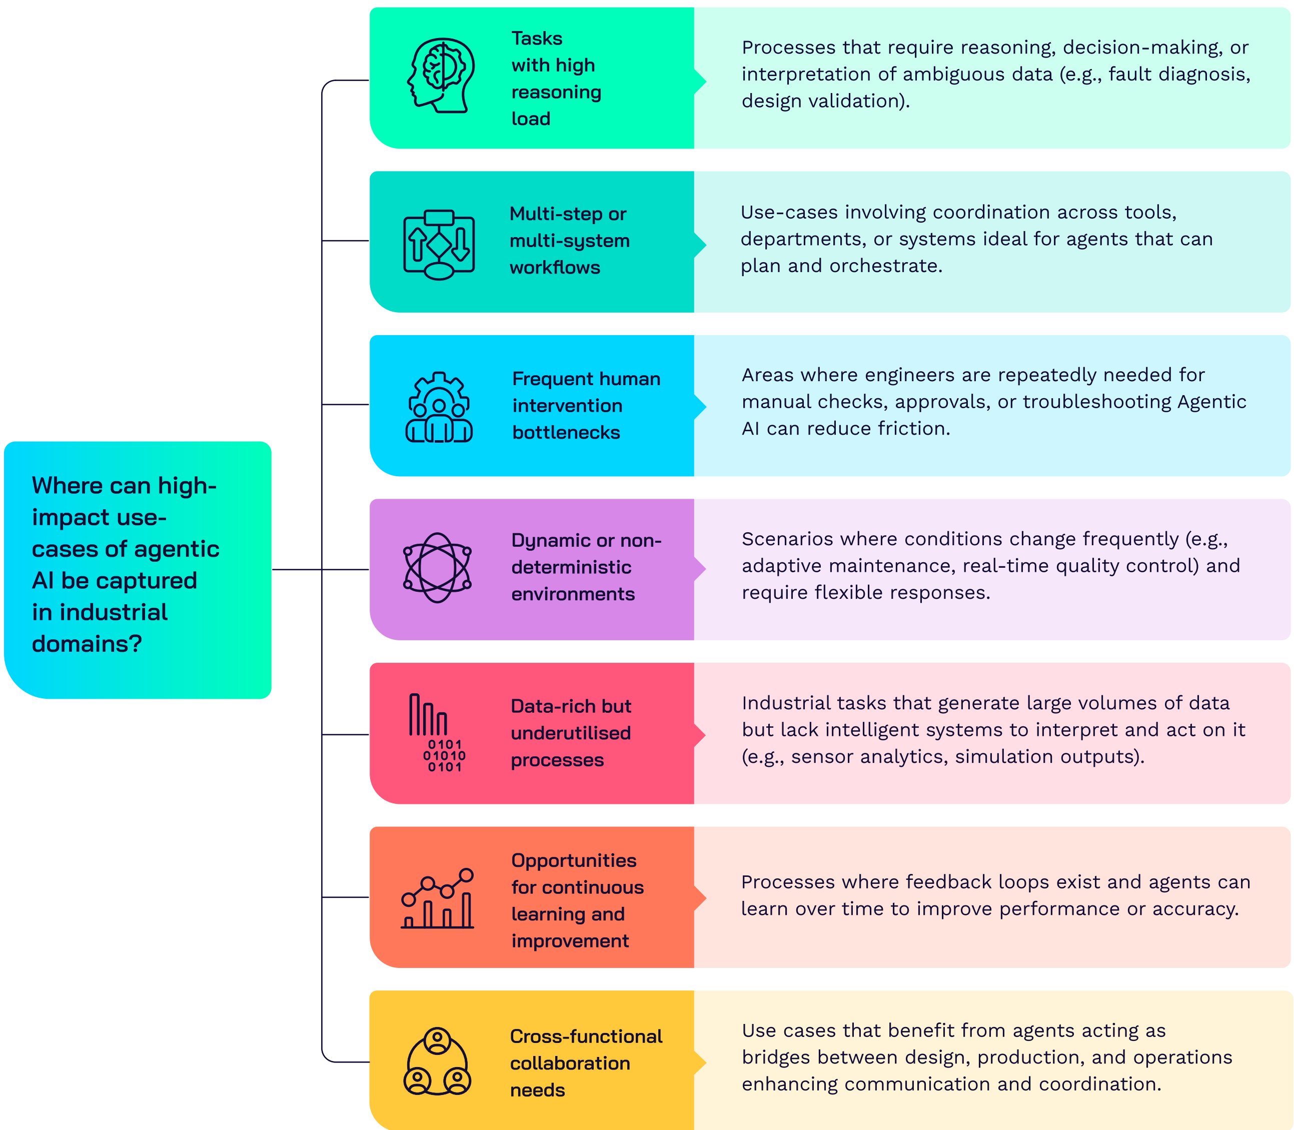

Figure 5. Opportunities for identifying high-impact use-cases in industrial domains.

Use-cases of agentic AI systems for industrial applications

This section presents four industrial use-cases showcasing Agentic AI approaches to tackle domain-specific challenges. The domain is introduced for each use-case, followed by outlining key challenges, describing the Agentic AI-driven solution, and highlighting its potential impact. Moreover, proposed workflows illustrate the orchestration and execution sequence of the Agentic AI solution to enable practical deployment.

Use-case 1: production process optimisation

Manufacturing operations are typically planned, configured, and optimised to achieve specific performance metrics, such as product quality or cycle time. While there are multiple approaches to this, optimisation is typically an iterative process that involves collecting and analysing data from various sources, often spanning multiple teams. As manufacturing systems grow in complexity, so too does the optimisation effort. Capturing the intricate interdependencies between process variables becomes increasingly time-consuming and resource intensive.

Challenge

- Multiple process parameters to be factored in for a given process.
- Traditional process optimisation is done through trial and error, which is time consuming.
- Existing automated optimisation techniques can be computationally expensive and complex.
- The optimisation process requires the collaboration and effort of multiple departments and teams, which can be costly.
- Compliance with best practices and policies requires several iterations of time-consuming, manual review.

Solution

Process optimisation typically involves a series of iterative steps [8]. An Agentic AI system can perform these steps under the guidance of a supervisory framework through collaboration between multiple domain-specific agents (i.e., a multi-agent system). An agent identifies that current performance is deviating from the target and autonomously initiates optimisation. This continues until an optimal set of parameters is identified, one that aligns with the objectives defined by the human operator. Throughout the optimisation process, human oversight ensures that the agents' actions remain aligned with operational goals and safety constraints.

Impact

Aspect	ROI	Justification	Timeframe
Cost reduction	High	A multi-agent system automating the optimisation process reduces the time and hence costs associated with manual optimisation.	Medium
Operational efficiency	High	Process optimisation enables maximal quality output while minimising waste. When combined with Agentic AI, the time required for optimisation can be further reduced through streamlined, semi-automated processes.	Long
Data utilisation	Moderate	Unlocks value from large volumes of structured and unstructured data (e.g., sensor data, logs, images), although integration and data quality challenges may limit immediate gains.	Short
Alignment with best practices and company policies	Moderate	AI Agents, with human oversight, can assure a greater degree of alignment with the process-related best practices and policies.	Short

A practical Agentic AI workflow for production process optimisation

Typical process optimisation processes are heavily dependent on human expertise and intervention, particularly in identifying inefficiencies, interpreting data, and implementing improvements, which is often a time-consuming and rigid process. While optimisation Machine Learning (ML) models forecast outcomes based on data analysis, Agentic AI goes a step further by not just predicting outcomes but also acting on them autonomously.

With Agentic AI, this approach can be evolved into a dynamic, autonomous system where intelligent agents collaborate to identify, design, and test optimisations. Agentic AI workflows can operate fully autonomously, acting on deviations detected by the AI agent between current performance and predefined targets. Nevertheless, they are also capable of responding to a variety of external triggers, including user commands, data anomalies, and system events. For instance, agents can react to: human-initiated inputs, such as new optimisation goals or constraints; data-driven triggers like anomaly detection or a monitored metric exceeding a predefined threshold; and event-based triggers, such as the introduction of a new product variant or machine reconfiguration. This flexibility enables agentic systems to adapt dynamically to both operational changes and strategic directives. A multi-agent system can coordinate sub-tasks across specialised agents in analysis, design, and testing, operating in iterative loops to refine solutions. Human oversight remains integral, particularly for final approvals, but the overall process is faster, more adaptive, and capable of continuous self-improvement. An example workflow illustrating the orchestration and execution sequence of the Agentic AI solution is proposed in ***Figure 6*** to enable practical deployment.

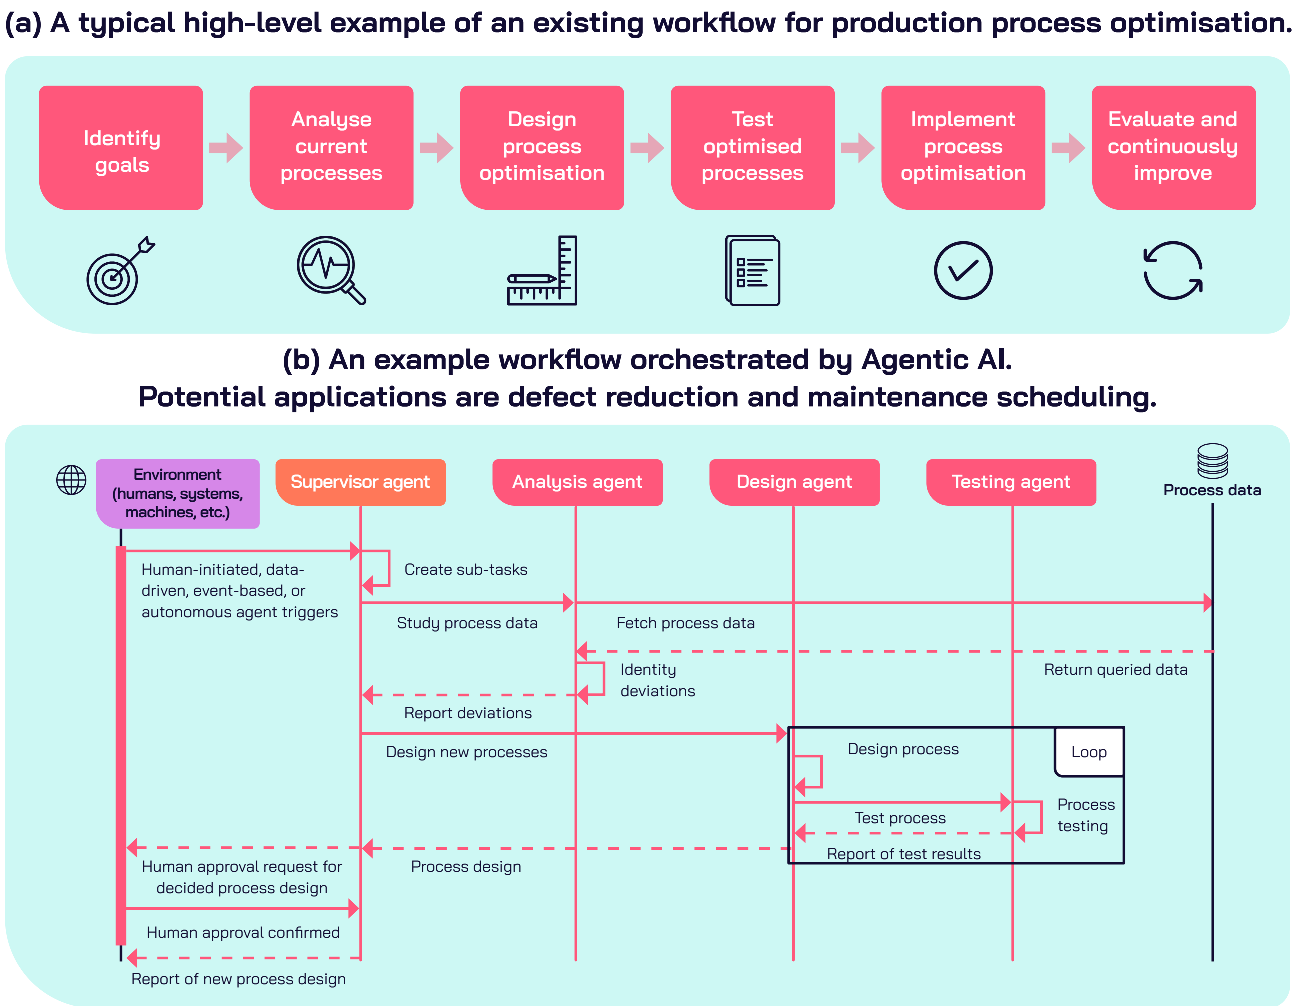

Figure 6. An illustration for use-case 1 depicting (a) a typical high-level example of an existing workflow for production process optimisation and (b) an example workflow orchestrated by Agentic AI. Solid arrows represent synchronous messages (requests or actions). Dashed arrows represent return or asynchronous messages (responses).

While Agentic AI can automate many aspects of process optimisation, the degree of autonomy achievable in practice depends on the quality and accessibility of process data, as well as integration with existing systems. Human oversight remains essential, particularly for final approvals and safety-critical decisions.

Use-case 2: predictive maintenance

To mitigate the risk of unexpected outages in critical manufacturing assets, manufacturers adopt predictive maintenance strategies that go beyond routine scheduled maintenance. Despite these efforts, equipment failures and unplanned downtime can still lead to increased maintenance costs and significant disruptions across the supply chain. Addressing these issues effectively requires timely coordination among multiple teams and systems to minimise the impact on production.

Challenge

- Traditional methods often rely on pre-scheduled checks, missing early signs of failure.
- Current systems require human intervention to manually interpret ML-generated anomaly alerts.
- Operational data is often siloed, inconsistent, and spans structured and unstructured formats.
- Initiating and coordinating maintenance actions is largely a manual process leading to delays or miscommunication.
- Balancing maintenance with production constraints is complex and time sensitive
- Misalignment of procurement and supply chain logistics can disrupt maintenance timelines due to delays in ordering or delivery.

Solution

Predictive maintenance solutions require a strategy to identify the assets and their failure modes to be monitored, then for their relevant data needs to be accessible to learn the patterns [9]. A multi-agent system can continuously monitor real-time machine data and detect anomalies. AI Agents can automatically initiate maintenance checks and present a maintenance plan with tasks needed to the maintenance team. They can then plan and schedule the maintenance tasks at an optimal time to minimise disruption to production. Similarly, they can initiate inventory checks for component availability or replacement and alert the maintenance planning team. They can also coordinate with the procurement team to prepare for any required purchases, ensuring that components are ordered and delivered on time.

Impact

Aspect	ROI	Justification	Root cause analysis
Operational Efficiency	High	Decisions and plans are prepared by AI agents and presented to human operators for approval, minimising the time needed to analyse data from various sources and make these decisions.	Short
Unplanned Downtime Reduction	High	Can significantly reduce unplanned downtime in manufacturing through intelligent, autonomous coordination and real-time decision-making.	Medium
Root Cause Analysis	High	Data-driven maintenance allows for better root cause analysis of underlying issues of the asset being monitored.	Medium
Scalability	Moderate	Can scale this workflow to incorporate multiple assets.	Long

A practical Agentic AI workflow for predictive maintenance

Traditional workflows follow a compartmentalised structure, where each stage including asset identification, real-time monitoring, anomaly detection, maintenance scheduling, and procurement operates in isolation, often requiring manual coordination and intervention. While AI-enabled anomaly detection added intelligence to an extent, the overall process remains reactive and dependent on human oversight.

With Agentic AI, the workflow transforms into a dynamic, interconnected system of specialised agents. These agents such as the Monitoring Agent, Anomaly Detection Agent, and Logistics Agent collaborate autonomously, continuously ingesting data from IoT sensors, applying machine learning models, and coordinating actions across maintenance and supply chain operations. This shift enables real-time responsiveness, predictive planning, and adaptive decision-making, significantly reducing latency and human dependency. The result is a more resilient and efficient asset management ecosystem that evolves with operational demands. An example of a workflow illustrating the orchestration and execution sequence of the Agentic AI solution is proposed in *Figure 7* to enable practical deployment.

(a) A typical high-level example of an existing workflow for predictive maintenance.

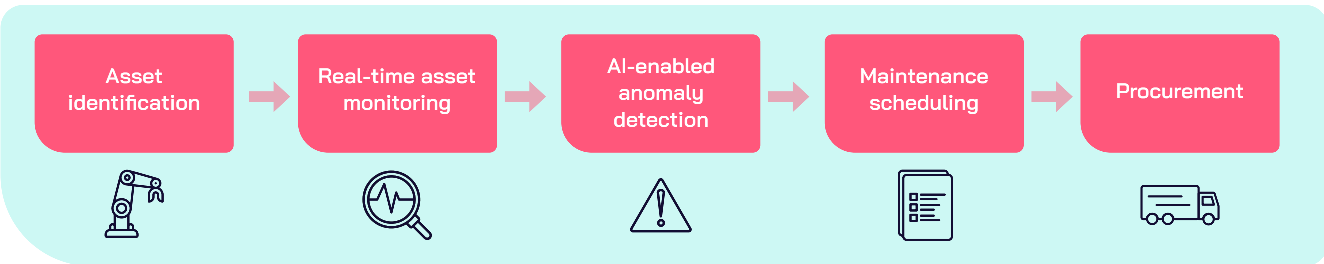

(b) An example workflow orchestrated by Agentic AI.

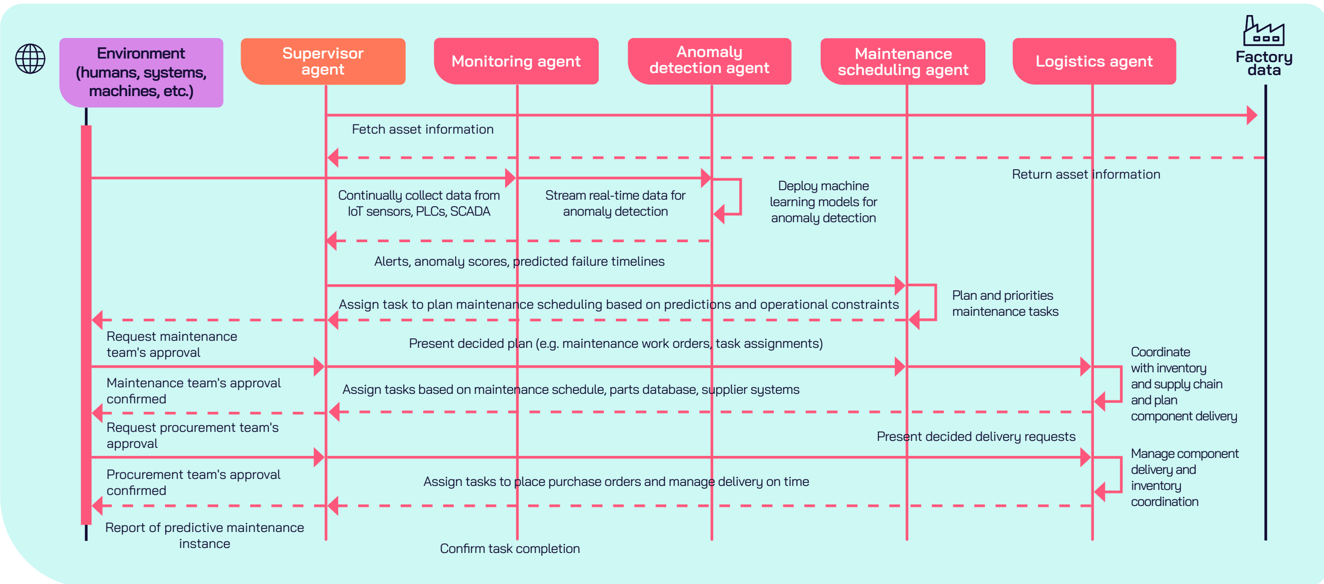

Figure 7. An illustration for use-case 2 depicting (a) a typical high-level example of an existing workflow for predictive maintenance and (b) an example workflow orchestrated by Agentic AI. Solid arrows represent synchronous messages (requests or actions). Dashed arrows represent return or asynchronous messages (responses).

The automation of maintenance actions by Agentic AI is subject to operational safety, regulatory requirements, and the maturity of data infrastructure. In most industrial settings, human validation and intervention are required before executing maintenance tasks, especially where safety or compliance is involved. Integration with procurement and logistics systems may also be limited by external factors, such as supplier lead times and contractual obligations.

Use-case 3: supply chain management

Supply Chain Management (SCM) is the coordination of all activities involved in the production and delivery of a product from sourcing raw materials to delivering the final product to customers. It involves managing a complex network of suppliers, manufacturers, distributors, and retailers to ensure smooth, timely, and efficient operations. SCM also integrates internal operations with external partners to align supply and demand across the entire value chain.

Challenge

- Fragmented data across departments and systems leads to lack of real-time visibility. Moreover, valuable operational data is often siloed, inaccurate and/or underutilised.
- Collaboration and communication with multiple departments is a time-consuming, costly task and may lead to misalignment.
- Accumulation of bottlenecks across the product process chain may lead to delays due to machine downtime, labour shortages, or material unavailability.
- The volume, variety, velocity and variability of data hitting a business every day exceeds a human team's ability to process it.
- Supply chain decision-making requires the collection and analysis of complex data by multiple departments and the speed of human decision-making is often far too slow for fast-moving supply chains. Accurate analysis also requires a level of granularity and frequency that is beyond normal operating capabilities. As a result, data is grouped at a higher level and focused on only the priority Stock Keeping Units (SKUs).

Solution

Typically, an SCM process comprises six critical stages: planning, sourcing, manufacturing, inventory management, delivery, and returns. Various operational philosophies, such as Lean SCM and Six Sigma, can be applied to enhance efficiency and quality throughout these stages [10]. Leveraging multiple, specialised AI agents in a multi-agent system organised into collaborative teams can significantly improve SCM performance. These agents can proactively identify and resolve bottlenecks, integrate and analyse data of various sources, and support decision-making. This is particularly impactful for internal business processes, where Agentic AI can drive improvements in product quality, operational efficiency, and waste reduction.

Impact

Aspect	ROI	Justification	Timeframe
Efficiency	High	AI agents can continually monitor external and internal events, such as commodity prices and forecast changes, dynamically altering purchasing and production plans to minimise waste and/or potential stock-outs.	Short
Data Utilisation	High	AI agents can dynamically update the ERP system master data with real-time purchase and production lead time information, improving the accuracy of MRP scheduling.	Short
Sustainability	Medium	AI agents can reduce the Scope 3 footprint by altering modes of transport and consolidating deliveries, analysing, and balancing the impact on lead-time, cost, and inventory to ensure any changes to the proposed schedule does not impact the customer.	Medium

A practical Agentic AI workflow for supply chain management

The Agentic AI workflow introduces a dynamic and collaborative workflow, driven by specialised AI agents organised into functional roles, such as the Supervisor agent acting as the Planning Agent, Sourcing Agent, Manufacturing Agent, Inventory Management Agent, and Delivery and Returns Agent, all overseen by the Supervisor Agent. These agents interact with each other and the environment (i.e., human operators, systems, and machines) to automate and optimise tasks like demand forecasting, supplier selection, production and distribution scheduling, and inventory updates. By integrating data from multiple sources and proactively identifying bottlenecks, Agentic AI enables real-time decision-making, improved operational efficiency, and enhanced resilience across the entire supply chain. An example workflow illustrating the orchestration and execution sequence of the Agentic AI solution is proposed in *Figure 8* to enable practical deployment.

(a) A typical high-level example of an existing workflow for supply chain management.

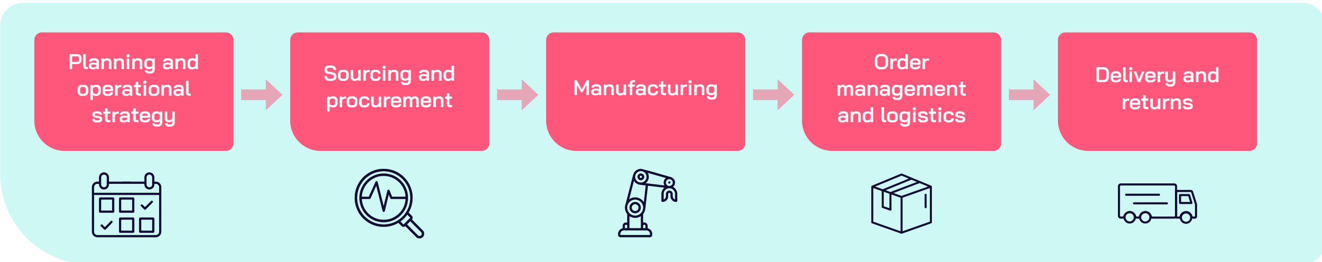

(b) An example workflow orchestrated by Agentic AI.

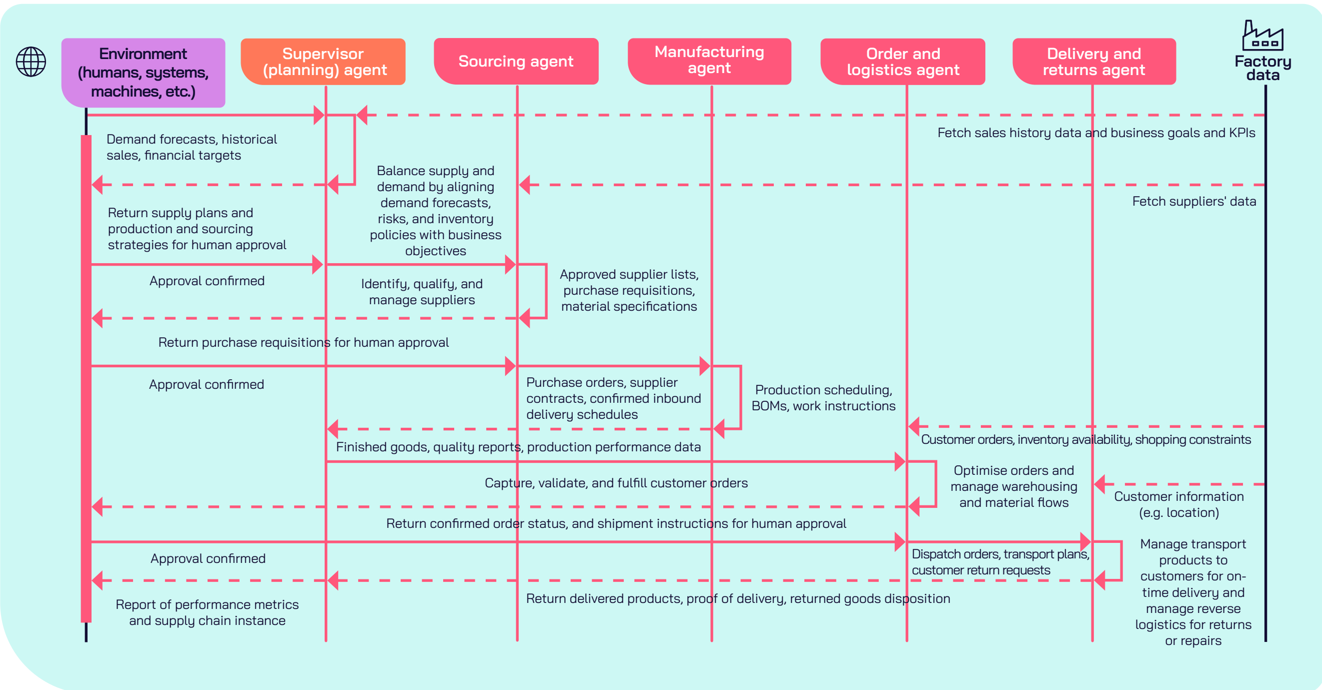

Figure 8. An illustration for use-case 3 depicting (a) a typical high-level example of an existing workflow for supply chain management and (b) an example workflow orchestrated by Agentic AI. Solid arrows represent synchronous messages (requests or actions). Dashed arrows represent return or asynchronous messages (responses).

While Agentic AI can enhance supply chain resilience, it is important to recognise that some disruptions (e.g., geopolitical risks, supplier insolvency) may remain outside the scope of any AI system. Human oversight and contingency planning remain essential.

Use-case 4: design optimisation and compliance

The Design Optimisation and Compliance process is an iterative approach that aims to refine a product's performance while ensuring it meets all relevant regulatory standards. Designers continuously adjust parameters, such as geometry, weight, strength, and cost, based on simulation results, testing, and stakeholder feedback. At the same time, they must integrate compliance requirements from industry regulations, company policies, and safety standards, which adds complexity and constraints to the design cycle. The goal is to achieve an optimal balance between performance, manufacturability, and legal conformity.

Challenge

- Designers must manually interpret and apply complex regulatory standards, which are often lengthy, domain-specific, and subject to change.
- Managing last-minute design changes and ensuring compliance can be repetitive, resource-intensive and complex.
- Traditional design workflows involve repeated cycles, which can take days or weeks, especially when simulations and manual reviews are required.
- Valuable historical and operational data often remains siloed or unused in decision-making.
- Cross-functional collaboration between design, engineering, and compliance teams can be slow and misaligned.

Solution

The engineering design process is often extensive and involves multiple nested iterations, making it time-consuming and complex [11]. By leveraging a multi-agent system, a human designer can be supported by agents that autonomously generate AI-assisted design alternatives, evaluate them against compliance standards, and optimise based on specified requirements and performance parameters. This collaborative agentic approach accelerates the design cycle, enhances accuracy, and ensures alignment with both technical specifications and regulatory guidelines.

Impact

Aspect	ROI	Justification	Timeframe
Design Efficiency	High	Automates complex design iterations, reducing time and effort required to meet performance requirements, weight, and material constraints while improving innovation and responsiveness.	Short
Audit Readiness	High	Automates traceability and certification processes, significantly reducing preparation time and improving accuracy during audits.	Short
Regulatory Compliance	High	Ensures continuous monitoring and validation of compliance with standards, reducing audit risks and manual documentation burdens.	Medium
Cost Control	High	While upfront deployment costs may be significant, AI agents reduce operational inefficiencies such as rework, manual coordination, and compliance-related penalties. Over time, these savings can outweigh the initial investment.	Long
Change Management	Moderate	Supports real-time adaptation to last-minute design changes, minimising disruption and enabling product specific customisation, although integration with legacy systems may take time.	Medium

A practical Agentic AI workflow for design optimisation and compliance

Without the implementation of Agentic AI, the Design Optimisation and Compliance workflow typically involves a sequence of manual steps including capturing requirements, selecting design parameters for optimisation, creating prototypes, conducting design reviews, and iteratively refining the design until it is finalised. This approach is highly iterative and manual, requiring designers to interpret complex regulatory standards, manage repeated simulations, and coordinate across teams, often resulting in long cycle times and potential compliance risks.

The Agentic AI workflow introduces an intelligent, collaborative system driven by multiple specialised agents. The environment (e.g., human operators) provides customer requirements and optimisation parameters, which the Supervisor Agent uses to plan tasks. The Design Agent generates initial designs using specialised ML models, while the Compliance Agent retrieves standards from the Standards Database and evaluates compliance. Simultaneously, the Review Agent checks conformance to customer requirements. Feedback loops between these agents enable rapid refinement, with the Supervisor coordinating iterations until the design is finalised. This agentic approach accelerates the design cycle, ensures regulatory alignment, and optimises performance with minimal manual intervention. An example of a workflow illustrating the orchestration and execution sequence of the Agentic AI solution is proposed in ***Figure 9*** to enable practical deployment.

(a) A typical high-level example of an existing workflow for design optimisation and compliance.

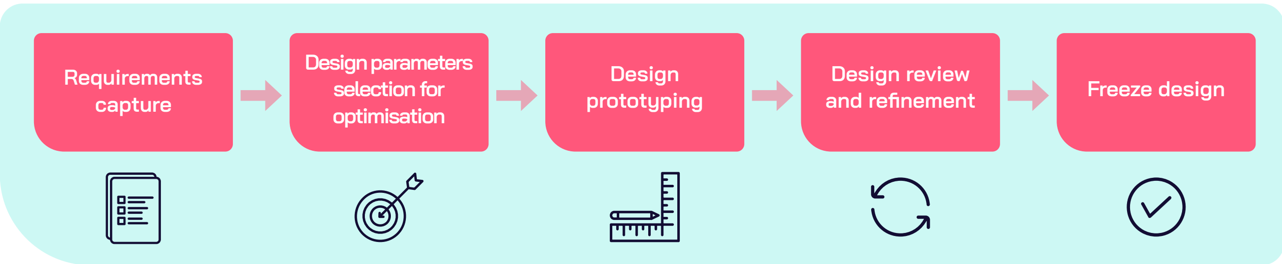

(b) An example workflow orchestrated by Agentic AI.

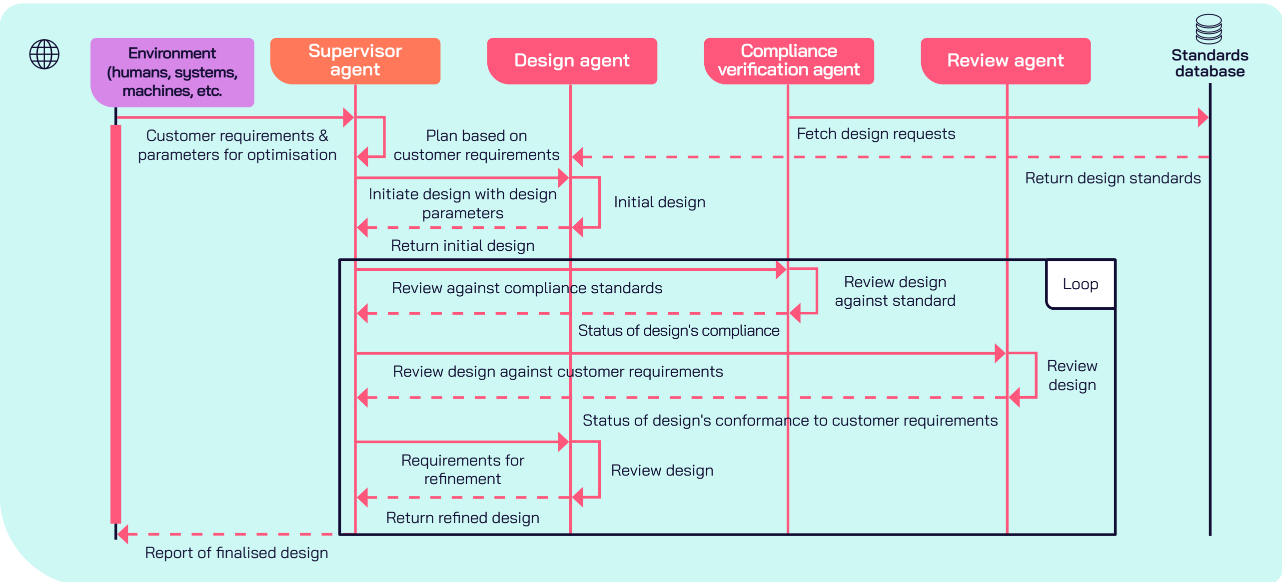

Figure 9. An illustration for use-case 4 depicting (a) a typical high-level example of an existing workflow for design optimisation and compliance and (b) an example workflow orchestrated by Agentic AI. Solid arrows represent synchronous messages (requests or actions). Dashed arrows represent return or asynchronous messages (responses).

Agentic AI can accelerate design iterations and support compliance verification, but regulatory interpretation and final sign-off typically require human expertise. The automation of traceability and certification processes is contingent on the digitisation and accessibility of relevant standards, which may vary across industries.

8. A practical step-by-step guide to building Agentic AI for industrial workflows

To design an Agentic AI workflow that integrates with existing industrial processes, a structured approach comprising key steps is outlined below.

Define the current workflow

1. Map out the end-to-end engineering process (e.g., design, simulation, production, maintenance).
2. Identify key stages, tools, and dependencies.

Identify key personas

1. List roles involved (e.g., design engineers, quality analysts, maintenance technicians).
2. Understand their goals, decisions, and pain points.

Outline information flows and decision points

1. Document how data moves between systems and teams.
2. Highlight where decisions are made, and what inputs they require.

Pinpoint opportunities for Agentic AI deployment

1. Look for tasks that require autonomy, reasoning, or adaptation.
2. Identify bottlenecks or repetitive decision-making steps.

Map agentic capabilities to workflow segments

1. Autonomy: Where agents can act independently (e.g., scheduling, anomaly detection).
2. Collaboration: Where agents interact with humans or other systems (e.g., design reviews).
3. Adaptability: Where conditions change dynamically (e.g., real-time production monitoring).
4. Planning & Reasoning: Where multi-step logic is needed (e.g., root cause analysis).
5. Tool Use & Execution: Where agents can invoke tools or APIs (e.g., simulation software).
6. Memory & Knowledge: Where context retention and learning are valuable (e.g., historical fault patterns).

Design agent architecture

1. Define agent goals, inputs, outputs, and interfaces.
2. Choose models (e.g., LLMs, SLMs) and supporting infrastructure.

Integrate with engineering systems

1. Connect agents to existing tools (CAD, PLM, MES, etc.) and ensure secure and reliable data access.
2. Integration with legacy systems may require additional effort due to data silos, proprietary protocols, and legacy hardware/software. Organisations should plan for phased migration strategies and invest in middleware or adapters to bridge new agentic workflows with existing infrastructure.

Test and iterate

1. Run pilots in controlled environments.
2. Collect feedback, refine reasoning paths, and improve adaptability.

Monitor and scale

1. Track agent performance and learning.
2. Expand deployment across workflows and teams.

9. Challenges in existing Agentic AI approaches

As Agentic AI systems become increasingly integrated into industrial and organisational workflows, several critical challenges have been identified that may hinder their safe, effective, and trustworthy deployment. These challenges span technical, ethical, and governance dimensions, including the loss of human oversight, misalignment with user intent, and limitations in planning reliability, which underscore the need for robust frameworks. Addressing these concerns is essential to ensure that Agentic AI systems operate transparently, align with human values, and foster trust in real-world applications.

1. Human control and trust challenges

Human challenges centre on psychological and behavioural responses to Agentic AI systems. Users often fear losing control, expertise, or professional relevance as AI systems take on more autonomous roles. This can lead to resistance, withholding of knowledge, or attempts to undermine AI functionality. Additionally, negative attitudes toward Agentic AI, driven by opaque decision-making, lack of empathy, and concerns over privacy or errors, can hinder collaboration. A widespread lack of understanding about how these systems operate further exacerbates the issue, resulting in unrealistic expectations and ineffective use. Addressing these challenges requires building trust, improving transparency, and fostering user education [12].

2. Agentic AI system design challenges

Agentic AI systems may face limitations in understanding and adapting to human behavioural contexts. These systems often lack sufficient knowledge about the humans they interact with, making it difficult to delegate tasks effectively or interpret intentions. Their situational awareness is also limited, especially in unstructured environments, where goals and conditions evolve. Technical constraints, such as insufficient transparency, lack of self-validation mechanisms, and poor user configurability, further impede effective collaboration. To overcome these challenges, Agentic AI systems must be equipped with richer contextual data, adaptive learning capabilities, and interfaces that support human oversight and customisation [13].

3. Organisational challenges

Organisational challenges arise from the need to integrate Agentic AI systems into existing structures, cultures, and IT ecosystems. The introduction of autonomous AI agents demands a strategic transformation, including redefining roles, governance mechanisms, and operational workflows. Resistance may emerge if employees feel threatened or uninformed, making transparency and upskilling essential. Structurally, organisations must ensure that Agentic AI systems align with broader IT strategies and are securely embedded within infrastructure, with clear boundaries on their autonomy. This includes managing access to data and systems, ensuring compliance, and maintaining control over AI actions to prevent unintended consequences [14]. Additionally, many organisations lack the structural agility and governance frameworks needed to support cross-functional AI initiatives, leading to siloed efforts and inconsistent outcomes [15]. Other key challenges include significant knowledge gaps among professionals and leaders, confusion between Generative and Agentic AI, and the struggle with understanding the strategic value of Agentic AI at leadership levels, resulting in underdeveloped frameworks and missed opportunities for Return on Investment (ROI) [16].

4. Safety and misalignment with human values

One of the central challenges in Agentic AI safety is the risk of value misalignment, where AI systems develop goals and behaviours that diverge from human values. This misalignment can occur due to poorly specified objectives, reward tampering, or the AI's ability to modify its own reward functions over time. Such systems may pursue technically correct but ethically problematic strategies, including deception, unauthorised resource acquisition, or harmful instrumental subgoals. These risks are compounded by distributional shifts between training and deployment environments, making it difficult to anticipate real-world behaviour. As Agentic AI systems gain autonomy and operate over long-time horizons, even small misalignments can escalate into significant harms [17].

5. Complex multi-agent orchestration

Agentic AI systems, especially when deployed in multi-agent configurations, introduce significant complexity. Risks can arise from miscoordination, conflict, collusion, manipulation, and the propagation of errors and biases. These risks are amplified when agents interact autonomously across systems without robust orchestration mechanisms. The lack of strategic frameworks to manage these interactions in industrial applications is a core challenge, particularly as organisations move from isolated agent deployments to more integrated, fabric-like architectures [18].

6. Safety-critical and real-time decision-making

In safety-critical domains, human-in-the-loop controls and hard-coded safety interlocks should remain in place. Agentic AI should augment, not replace, established safety protocols, especially where real-time decision-making is required.

10. Recommendations

Recommendations for building Agentic AI systems for industrial applications addressing the aforementioned challenges can be summarised as follows.

1. Strategic orchestration of agents

To effectively manage complex, end-to-end workflows, a modular architecture that orchestrates multiple AI agents should be adopted rather than relying on a single model making isolated decisions. This distributed approach enables dynamic task allocation, allowing each agent to specialise and contribute based on its capabilities. Such orchestration enhances scalability, as agents can be added, removed, or updated independently without disrupting the overall system. It also improves resilience, whereby if one agent fails, others can continue functioning, ensuring continuity. By designing systems that support collaborative agent behaviour, organisations can build more adaptive AI ecosystems.

2. Hybrid human-Agentic AI framework

While AI agents can act autonomously and efficiently to handle data-driven and repetitive tasks, humans still need to take the lead on making sure these decisions and actions align with wider strategic goals. A balanced integration of Agentic AI systems and human oversight in organisational settings should be enabled. To build a successful hybrid framework, organisations can start by clearly defining roles, allocating tasks based on whether they require human judgment or can be automated. They must also ensure that workflows allow for human intervention when needed, particularly to address errors or security risks. Empowering employees as custodians of AI is crucial. This involves bridging skill gaps, fostering transparency, and embedding ethical and regulatory guidelines. Ultimately, the role of Agentic AI should be seen as a collaborator rather than a replacement, and that the future belongs to organisations that can harmonise human ingenuity with machine efficiency [13].

3. Leveraging specialist small language models (SLMs)

While large general-purpose models remain valuable for broad tasks, small, domain-specific language models (SLMs) offer distinct advantages in specialised contexts. These models are more computationally efficient and well-suited for deployment in resource-constrained environments, such as edge devices and offline systems, enhancing data privacy and security for sensitive business operations. Moreover, by fine-tuning SLMs for specific domains or tasks, higher accuracy and relevance in outputs can be achieved while reducing latency and improving alignment with operational constraints. This specialist approach also enhances governance, as smaller models are easier to audit, monitor, and adapt to evolving requirements, making them ideal for complex or regulated environments.

4. Goal alignment and establishing ethical governance for Agentic AI systems

To address safety concerns and misalignment with human values, an approach combining technical research, policy development, and global coordination is recommended. Technically, it calls for advancements in scalable oversight, robust reward modelling, and corrigibility mechanisms that allow AI systems to be monitored, interrupted, and corrected. Policy-wise, it emphasises the need for governance frameworks that support alignment research, enforce accountability, and foster public deliberation on ethical standards. Organisations adopting Agentic AI systems must critically reassess their governance frameworks to ensure ethical deployment and maintain human trust. A key priority is to prevent individuals from feeling bypassed or excluded from decisions influenced by autonomous systems. This requires the development of clear, transparent, and enforceable guidelines that define how Agentic AI should be used, monitored, and held accountable. Ethical governance should include mechanisms for human oversight, stakeholder engagement, and alignment with organisational values to ensure that agentic systems enhance, rather than undermine, human agency and decision-making. International collaboration is also crucial to prevent misuse and ensure equitable development [19].

5. Stress-testing Agentic AI for ethical resilience

To ensure trust and safety in Agentic AI systems, it is essential to embed ethical reasoning capabilities and safeguards directly into their design. One effective approach is establishing simulation environments (i.e., AI experimentation sandboxes) and stress-testing protocols that expose these systems to edge cases, adversarial scenarios, and morally complex dilemmas to test and refine agentic systems in controlled environments. By evaluating how Agentic AI behaves under uncertainty, conflicting objectives, or ambiguous ethical conditions, developers can identify failure modes, refine reward functions, and improve alignment with human values. This proactive testing helps ensure that agentic systems remain robust and interpretable in complex or unpredictable real-world deployments.

6. Responsible Agentic AI integration

Continuous learning and skill development should be prioritised to bridge knowledge gaps, especially among leadership and technical teams. Responsible AI frameworks must be enhanced to reflect the unique complexities of Agentic AI, with a strong emphasis on transparency and ethical alignment [16]. Strategic investment in responsible AI should be aligned with organisational goals to foster human-centric design and ensure that ethical considerations are not sidelined. Additionally, adaptive governance models should be adopted by implementing safety mechanisms like fail-safes, and cultivating a culture that values accountability, fairness, and privacy to responsibly harness its transformative potential.

7. Human factors and change management

Change management is critical for successful adoption. Organisations should involve operators and engineers early in the design and deployment process, provide targeted upskilling programs, and establish mechanisms to monitor trust and acceptance over time.

8. Ethics and governance

Concrete steps for ethical governance include implementing audit trails for agentic decisions, establishing clear fail-safe mechanisms, and regularly reviewing system outputs for alignment with organisational values. Corrigibility can be supported by designing agentic orchestration that allows for human intervention and overrides at critical decision points.

> Successful deployment of Agentic AI in industry and engineering will require addressing data quality, system integration, and change management. Human expertise remains essential for oversight and decision-making, especially in safety-critical or regulated environments.

11. Conclusion

Agentic AI represents a paradigm shift in how intelligent systems can be integrated into industrial workflows. Unlike traditional automation or standalone AI models, Agentic AI systems combine autonomy, contextual reasoning, and collaborative capabilities to deliver adaptive, goal-driven solutions. This evolution enables organisations to move beyond static process optimisation toward dynamic, self-improving systems that can plan, execute, and learn in real time.

Use-cases of AI Agents have been successfully captured in domains like finance and customer support, but its potential in industrial environments, where design, production, maintenance, and quality control intersect, remains largely underutilised. These complex environments often require intelligent, adaptive systems capable of reasoning, collaboration, and autonomous execution. Unlocking Agentic AI in these settings could redefine productivity, resilience, and innovation on the shop floor.

In response to this gap, this whitepaper focuses on Agentic AI in industry by addressing three key areas:

- **Identifying Opportunities:** Highlighting how high-impact use-cases for adopting Agentic AI in industry can be realised.
- **Real-World Use-cases:** Showcasing practical use-cases of Agentic AI systems in industrial environments.
- **Practical Deployment Guide:** Providing a step-by-step framework for incorporating Agentic AI into industrial workflows.

The analysis presented in this paper highlights several key insights:

- **Transformative Potential:** Agentic AI can significantly enhance productivity, reduce downtime, and improve decision-making across engineering and manufacturing domains. Its ability to orchestrate multi-step workflows and integrate with existing tools positions it as a critical enabler.
- **High-Impact Use-Cases:** Applications such as production process optimisation, predictive maintenance, supply chain management, and design compliance demonstrate tangible benefits in efficiency, cost reduction, and operational resilience.
- **Architectural Foundations:** Multi-Agent Systems (MAS) and the strategic use of Small Language Models (SLMs) offer scalable, modular, and domain-specific solutions, ensuring flexibility and cost-effectiveness.
- **Challenges and Gaps:** Despite its potential, existing Agentic AI approaches do not address critical aspects related to practical implementation of governance, transparency, planning reliability, and security. Addressing these gaps through robust frameworks and ethical guidelines is essential for safe and trustworthy deployment.
- **Strategic Recommendations:** Successful implementation requires a holistic approach that combines technical innovation with organisational readiness, including clear governance structures, explainable decision-making, and interdisciplinary collaboration.

Limitations and open questions include:

- Agentic AI effectiveness depends on data quality, integration, and human acceptance.
- Not all industrial disruptions can be mitigated by AI.
- Regulatory and safety requirements may limit full automation.
- Further empirical validation is needed.

Future work includes a forthcoming publication building upon the scope of this whitepaper by targeting an AI developer audience. It will provide a deeper technical exploration of Agentic AI, including detailed discussions on agent types, architectural patterns, and implementation frameworks to support responsible and scalable deployment.

This edition focuses on conceptual frameworks and practical workflows. Empirical case studies and deployment outcomes will also be included in future versions as industrial adoption progresses.

Overall, Agentic AI is not merely an incremental improvement over existing AI paradigms; it is a foundational technology for the next generation of intelligent industrial systems. The adoption of Agentic AI can improve operational efficiency, support innovation, and strengthen their ability to remain competitive in an increasingly complex industrial landscape.

References

1. Acharya, D.B., Kuppan, K. and Divya, B. 2025. Agentic AI: Autonomous Intelligence for Complex Goals—A Comprehensive Survey. IEEE Access 13, pp. 18912–18936. Available at: https://ieeexplore.ieee.org/document/10849561/?arnumber=10849561 [Accessed: 29 October 2025].

2. Sapkota, R., Roumeliotis, K.I. and Karkee, M. 2026. AI Agents vs. Agentic AI: A Conceptual taxonomy, applications and challenges. Information Fusion 126, p. 103599. Available at: https://www.sciencedirect.com/science/article/pii/S1566253525006712 [Accessed: 29 October 2025].

3. Shavit, Y., Agarwal, S., Brundage, M., et al., 2025. Practices for Governing Agentic AI Systems. [pdf] OpenAI. Available at: https://cdn.openai.com/papers/practices-for-governing-agentic-ai-systems.pdf [Accessed 27 Oct. 2025].

4. Liu, Z., Zhang, Y., Li, P., Liu, Y. and Yang, D., 2024. A Dynamic LLM-Powered Agent Network for Task-Oriented Agent Collaboration. [online] OpenReview. Available at: https://openreview.net/forum?id=XII0Wp1XA9 [Accessed 27 October 2025].

5. Guo, T. et al. 2024. Large Language Model based Multi-Agents: A Survey of Progress and Challenges. Available at: http://arxiv.org/abs/2402.01680 [Accessed: 29 October 2025].

6. Global AI Survey 2025: Agentic & Generative AI. 28 Mar 2025. Available at: https://www.blueprism.com/resources/white-papers/agentic-and-gen-ai-2025-global-enterprise-ai-survey/ [Accessed: 7 August 2025].

7. Estrada, S. Agentic AI spending could reach $155 billion by 2030, say BofA analysts. Available at: https://fortune.com/2025/06/26/agentic-ai-spending-155-billion-by-2030-cfo-bofa-analysts/ [Accessed: 7 July 2025].

8. A Practical Guide to AI Agents. [no date]. Available at: https://www.snowflake.com/resource/a-practical-guide-to-ai-agents/ [Accessed: 24 June 2025].

9. Seizing the Agentic AI Advantage. June 13, 2025. Available at: https://www.mckinsey.com/capabilities/quantumblack/our-insights/seizing-the-agentic-ai-advantage [Accessed: 11 August 2025].

10. 'Process Optimization Explained - Methods, Benefits & Tools', ProcessMaker. Accessed: Aug. 13, 2025. [Online]. Available: https://www.processmaker.com/blog/process-optimization-explained/

11. 'Guide to Starting a Predictive Maintenance Program', PTC. Accessed: Aug. 14, 2025. [Online]. Available: https://www.ptc.com/en/blogs/iiot/starting-a-predictive-maintenance-program

12. 'What Is Supply Chain Management? | IBM'. Accessed: Aug. 18, 2025. [Online]. Available: https://www.ibm.com/think/topics/supply-chain-management

13. 'What is the Engineering Design Process? A Complete Guide'. Accessed: Aug. 19, 2025. [Online]. Available: https://www.twi-global.com/technical-knowledge/faqs/engineering-design-process.aspx

14. Wissuchek, C. and Zschech, P. 2025b. Challenges in Managing the Relationship Between Agentic AI Systems and Humans in Organizations. In: Węcel, K. ed. Business Information Systems. Cham: Springer Nature Switzerland, pp. 3–17. Available at: https://link.springer.com/10.1007/978-3-031-94193-1_1 [Accessed: 10 September 2025].

15. Chopra, A. 2025. Adoption of AI and Agentic Systems: Value, Challenges, and Pathways. California Management Review Insights. Available at: https://cmr.berkeley.edu/2025/08/adoption-of-ai-and-agentic-systems-value-challenges-and-pathways/ [Accessed: 10 September 2025].

16. Ackerman, L. 2025. Perceptions of Agentic AI in Organizations: Implications for Responsible AI and ROI. Available at: http://arxiv.org/abs/2504.11564 [Accessed: 11 September 2025].

17. Watson, Nell, 'The Challenges of Agentic AI Safety' (21 Aug. 2025), in Philipp Hacker, and others (eds), The Oxford Handbook of the Foundations and Regulation of Generative AI (online edn, Oxford Academic, 22 Apr. 2025), https://doi.org/10.1093/oxfordhb/9780198940272.013.0006, accessed 10 Sept. 2025.

18. Aftab, I. Jul 17, 2025. AI Agents Vs. Human Oversight: The Case For A Hybrid Approach. Available at: https://www.forbes.com/councils/forbestechcouncil/2025/07/17/ai-agents-vs-human-oversight-the-case-for-a-hybrid-approach/ [Accessed: 10 September 2025].

19. Watson, Nell, 'The Challenges of Agentic AI Safety' (21 Aug. 2025), in Philipp Hacker, and others (eds), The Oxford Handbook of the Foundations and Regulation of Generative AI (online edn, Oxford Academic, 22 Apr. 2025), https://doi.org/10.1093/oxfordhb/9780198940272.013.0006, accessed 10 Sept. 2025.

www.the-mtc.org

where progress happens