© 2026 IJERR | Volume 9, Issue 7 July 2026 | ISSN: 2321-9912 | ijerr.com

Sustainable Agentic AI for Smart Manufacturing: A Layered Framework for Autonomous, Energy-Aware Industrial Decision-Making

Dr. K. Devika Rani Dhivya[1] (HOD), G. V. Poornima Radhalakshmi[1], E. John Joshua[1]
[1]Department of Computer Science, Sri Krishna Arts and Science College, Coimbatore, Tamil Nadu, India

Abstract—Manufacturing companies face increasing demands to enhance productivity while also reducing energy usage, material waste, and carbon emissions. Traditional automation and even Industry 4.0 connected systems primarily operate in a reactive manner: they gather data and deliver it to human operators, yet the understanding, prioritization, and long-term decision-making still depend on individuals. This study explores agentic artificial intelligence — systems made up of goal-oriented, tool-utilizing, and self-organizing software agents — as a method for bridging that divide in a sustainability-conscious way. A five-layer reference architecture is suggested, integrating a physical shop-floor layer, an edge/IIoT layer, a shared knowledge and digital-twin layer, a multi-agent decision layer, and a sustainability-governance layer that consistently evaluates all potential actions based on energy, emissions, and waste criteria prior to execution. A closed operational cycle, called Monitor–Analyze–Plan–Execute–Sustain (MAPE-S), is outlined along with a multi-objective reward framework that enables agents to weigh throughput against environmental expenses. The framework undergoes qualitative evaluation compared to traditional automation and Industry 4.0 benchmarks across seven operational dimensions, with five typical industrial use cases examined for their anticipated sustainability effects. The conversation also highlights the practical obstacles — trust and transparency of autonomous agents, compatibility with existing operational technology, the energy consumption of the agents, and workforce preparedness — that need to be addressed before these systems can be scaled effectively. This work aims to provide researchers and practitioners with an organized foundation for creating agentic AI systems where sustainability is regarded as a primary design goal instead of a retroactive evaluation.

Index Terms*—Agentic AI, intelligent manufacturing, sustainability, multi-agent systems, Industry 5.0, digital twin, energy-conscious scheduling, eco-friendly manufacturing.*

I. INTRODUCTION

The manufacturing industry represents a significant portion of worldwide energy usage and industrial greenhouse gas emissions, while also being the sector most vulnerable to fluctuating demand, supply interruptions, and increasing energy costs. In the past ten years, Industry 4.0 efforts linked machines, sensors, and business systems via the Industrial Internet of Things (IIoT), providing plant managers exceptional insight into their operations. Visibility, nonetheless, differs from autonomy. The majority of existing systems continue to send dashboards and alerts to human operators who must analyze multiple conflicting signals — machine condition, order backlog, energy costs, material supply, and emission limits — before determining their next steps. As production systems become increasingly distributed and more focused on sustainability, this human-in-the-loop bottleneck shifts from being a safeguard to a limiting factor.

Agentic AI has arisen as a potential remedy for this limitation. In contrast to a traditional AI model that delivers one prediction or classification, an agentic system comprises software agents capable of perceiving their surroundings, reasoning about multi-step objectives, utilizing external tools, engaging in negotiations with other agents, and interacting with the physical world with minimal human oversight. When these agents are integrated with large language models for advanced reasoning and more lightweight rule-based or smaller-model agents for quick edge execution, the resultant hybrid systems can strategize, orchestrate, and adjust in ways static automation logic cannot achieve.

Sustainability adds a second requirement that is frequently treated as an afterthought in the agentic-AI literature: every autonomous decision an agent makes — accelerating a machine, rerouting a batch, switching to a backup supplier — carries an energy, emission, or waste consequence. If sustainability metrics are only measured after the fact, autonomous agents can inadvertently optimise for throughput or cost while degrading environmental performance. This paper argues that sustainability criteria should instead be embedded directly into the agents' perception, planning, and reward mechanisms, so that every autonomous action is evaluated jointly against operational and environmental objectives.

The contributions of this paper are threefold. First, we propose a five-layer reference architecture that separates physical execution, edge sensing, shared knowledge, agentic decision-making, and sustainability governance into distinct but continuously communicating layers. Second, we describe a closed operational cycle — Monitor, Analyze, Plan, Execute, and Sustain (MAPE-S) — that extends the classical autonomic-computing MAPE-K loop with an explicit sustainability feedback path. Third, we compare the proposed approach qualitatively against traditional automation and Industry 4.0 practice, examine five representative manufacturing use cases, and consolidate the open challenges that remain before such systems can be trusted with real production authority.

The remainder of the paper is organised as follows. Section II reviews recent literature on agentic AI in manufacturing and on sustainability-oriented Industry 5.0 research. Section III presents the layered architecture. Section IV details the MAPE-S operational cycle and its multi-objective decision logic. Section V discusses representative use cases. Section VI provides a comparative analysis against existing paradigms. Section VII outlines open challenges, Section VIII sketches future research directions, and Section IX concludes the paper.

II. RELATED WORK

Interest in agentic AI for industrial settings has grown quickly alongside the broader rise of large-language-model-based agents. Farahani et al. proposed a hybrid architecture that pairs LLM-based agents, which provide strategic orchestration and adaptive reasoning, with rule-based and small-language-model agents that execute domain-specific tasks at the edge; their case study on prescriptive maintenance reported improved robustness and explainability compared with either approach used alone [1]. In a related direction, an autonomous multi-agent system built around a five-level agentic-intelligence maturity model was applied to predictive machinery-health monitoring in ceramic-tile manufacturing, combining federated learning with distributed agents monitoring presses, kilns, and glazing lines; the reported outcomes included a marked reduction in false positives and unplanned downtime together with a favourable economic payback period [2].

Lee and Su formally distinguished agentic AI from earlier generations of AI agents in a manufacturing context and introduced a framework built around multiple LLM-based agents coordinated through a unified Data–Model–Knowledge lake, demonstrating feasibility through a retrieval-augmented-generation question-answering case study while also cataloguing the technical barriers that remain [3], [4]. A broader conceptual review situates agentic AI as the most recent stage in an evolution from reactive agents through cooperative and cognitive agents, and links this evolution to a cloud–edge–device deployment structure that clarifies where each agent type is best hosted [5], [6]. Complementing this, a systematic review on collaborative and sustainable engineering design proposed an agentic-AI-enabled framework spanning five concurrent dimensions — sensing, smart, sustainability, social, and safety — illustrating that sustainability is increasingly being folded into agent design rather than treated as an external constraint [7].

A parallel body of work has examined sustainability within the Industry 5.0 paradigm without necessarily invoking agentic AI. Reviews of AI-driven optimisation techniques report that metaheuristic and reinforcement-learning methods can jointly balance productivity, energy efficiency, and environmental impact in smart-factory settings [8].

Several recent studies on Industry 5.0 describe **human-centricity, resilience, and sustainability** as the key principles shaping the future of manufacturing. Researchers have shown that technologies such as digital twins, collaborative robots, and IoT-based monitoring systems help industries reduce unnecessary energy usage and improve the utilization of available resources. Other studies also explain that the level of automation directly influences energy consumption and overall production efficiency. In addition, affordable sensing technologies, including IoT-based monitoring devices, make it easier to observe manufacturing activities that were previously difficult to track. Researchers have also demonstrated that AI optimization techniques can improve productivity while reducing energy usage and supporting environmentally responsible manufacturing.

The existing research clearly indicates two important observations. First, agentic AI has reached a stage where it can be applied to several manufacturing tasks, particularly in areas such as maintenance planning and production coordination. Second, Industry 5.0 research has already introduced effective approaches for improving sustainability through better energy management and emission control. However, most existing studies still consider sustainability as only one performance metric instead of making it a central decision-making factor. Likewise, many sustainability-focused approaches do not yet make full use of autonomous multi-agent systems. This gap provides the motivation for the framework proposed in this paper.

TABLE I
Summary of Representative Related Work

Ref.	Focus Area	Approach	Sustainability Emphasis
[1]	Prescriptive maintenance	Hybrid LLM + rule-based multi-agent system	Indirect (efficiency-driven)
[2]	Predictive machinery health monitoring	Federated multi-agent, 5-level maturity model	Moderate (downtime reduction)
[3],[4]	General agentic AI for manufacturing	LLM agents + unified Data-Model-Knowledge lake	Not a primary objective
[5],[6]	Agent-technology evolution & deployment	Cloud-edge-device mapping framework	Discussed conceptually
[7]	Sustainable engineering design	5-dimension (S5) agentic collaboration framework	Core design dimension
[8],[15]	Multi-objective process optimisation	Metaheuristic / reinforcement learning	Explicit optimisation target
[9]-[12]	Industry 5.0 sustainability reviews	Systematic literature review	Central theme
[13]	AI and energy consumption in Industry 4.0/5.0	Evidence-based analytical framework	Central theme
Proposed	Sustainable agentic AI for manufacturing	Layered architecture + MAPE-S agent cycle	Embedded as a governing layer

III. PROPOSED FRAMEWORK: SUSTAINABLE AGENTIC AI ARCHITECTURE

The proposed framework separates industrial decision-making into five layers so that sustainability considerations can be enforced consistently regardless of which agent or subsystem is acting. Fig. 1 illustrates the layered structure and the direction of information exchange between layers. Data flows upward from the physical shop floor to the sustainability layer, while decisions and control commands flow back down; the digital-twin and communication layers sit in the middle so that any agent can query the current or predicted state of the plant before proposing an action.

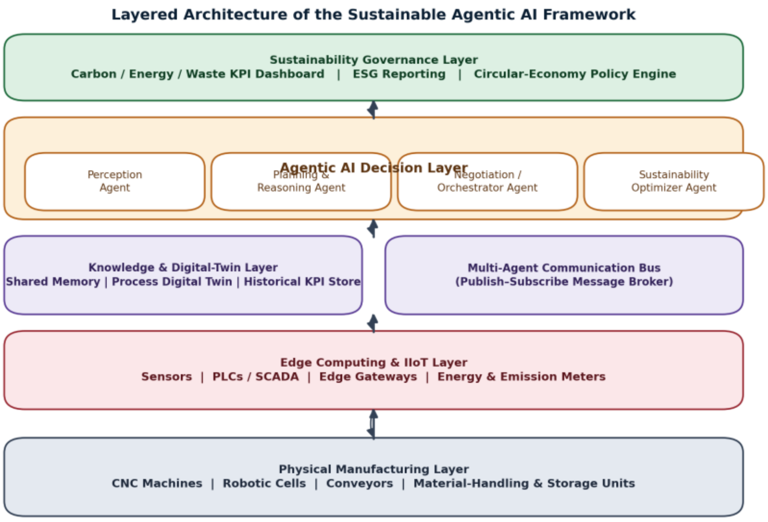

Fig. 1. Layered architecture of the proposed sustainable agentic AI framework.

A. Physical Manufacturing Layer

The foundation layer consists of the physical assets that carry out production: CNC machines, robotic work cells, conveyors, and material-handling equipment. These assets are treated as the ultimate actuation targets of every agent decision; no agent is permitted to act directly on this layer without passing through the edge and communication layers described below, which enforce safety interlocks and rate limits.

B. Edge Computing and IoT Layer

Sensors, programmable logic controllers, SCADA systems, and dedicated energy and emission meters continuously stream telemetry to edge gateways. Edge processing performs lightweight filtering, unit normalisation, and anomaly flagging before data is forwarded upstream, which reduces the bandwidth and latency burden on the agentic layer and allows time-critical safety responses to be handled locally rather than waiting for a remote agent decision.

C. Knowledge and Digital-Twin Layer

A communal knowledge repository and digital twin of processes provide each agent with a unified, updated perspective of plant condition, past performance, and sustainability metrics. The digital twin additionally acts as a low-risk simulation platform: before an agent commits to a physical action, candidate plans can be evaluated against the twin to estimate their energy, throughput, and quality consequences. A publish–subscribe communication bus connects this shared layer to the agentic decision layer so that agents remain loosely coupled and can be added, removed, or updated without redesigning the whole system.

D. Agentic AI Decision Layer

Four cooperating agent roles form the core of the decision layer. A perception agent fuses sensor and digital-twin data into a structured situational summary. A planning and reasoning agent, typically backed by a large language model, generates and ranks candidate courses of action against production goals. A negotiation and orchestration agent resolves conflicts between competing local objectives — for example when a maintenance agent wants a machine paused while a scheduling agent wants it running — and sequences the agreed actions for execution. A dedicated sustainability optimiser agent evaluates every candidate plan against energy, emission, and waste criteria and can veto or modify a plan that would breach a sustainability threshold even if it is otherwise attractive operationally.

E. Sustainability Governance Layer

The top layer aggregates plant-wide sustainability indicators into a live dashboard, produces the data needed for environmental, social, and governance (ESG) reporting, and encodes circular-economy policies — such as preferential routing of reworkable material or restrictions on high-emission process paths during peak grid-carbon-intensity periods — that all lower-layer agents must respect. Because this layer sits structurally above the decision layer rather than being consulted only for periodic reporting, sustainability constraints are treated with the same operational priority as safety interlocks.

IV. OPERATIONAL METHODOLOGY: THE MAPE-S AGENT CYCLE

Autonomic-computing systems have long used the Monitor–Analyse–Plan–Execute loop, coordinated through a shared Knowledge base (MAPE-K), to give software systems a degree of self-management. This paper extends that loop with an explicit sustainability checkpoint, producing a five-stage Monitor–Analyze–Plan–Execute–Sustain (MAPE-S) cycle, illustrated in Fig. 2.

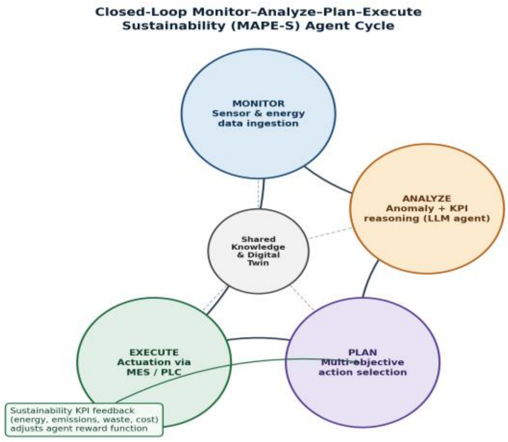

Fig. 2. The MAPE-S closed-loop agent cycle with an embedded sustainability feedback path.

In the Monitor stage, the perception agent ingests sensor, energy-meter, and order-book data and writes a normalised state snapshot to the shared knowledge base. In the Analyze stage, the reasoning agent compares the current snapshot against expected operating envelopes to identify deviations, emerging bottlenecks, or maintenance risks. In the Plan stage, the reasoning agent proposes one or more candidate action sequences, each scored by the sustainability optimiser agent before being passed to the orchestrator. In the Execute stage, the orchestrator sequences approved actions and issues them to the edge layer for actuation, subject to the safety interlocks described in Section III-B. The Sustain stage closes the loop: realised energy, emission, and waste outcomes are measured and fed back into the shared knowledge base, where they adjust the weighting the sustainability optimiser applies in future planning cycles.

Formally, at each planning cycle the system evaluates a candidate action a from a feasible set A and selects the action that maximises a weighted multi-objective score:

$$S(a) = w_1 \cdot Throughput(a) + w_2 \cdot Quality(a) - w_3 \cdot Energy(a) - w_4 \cdot Emissions(a) - w_5 \cdot Waste(a)$$

where the weights w_1 through w_5 are set by plant policy and can be adjusted dynamically by the sustainability governance layer — for instance, the weight on emissions can be increased automatically during periods of high grid-carbon-intensity, causing the agents to defer non-urgent, energy-intensive jobs without any change to the underlying planning algorithm. Actions whose sustainability score falls below a hard threshold are excluded from A entirely, regardless of their operational score, which is what allows the sustainability layer to function as a governing constraint rather than a soft preference.

V. REPRESENTATIVE USE-CASE SCENARIOS

To ground the architecture in concrete operational terms, Table II outlines five manufacturing scenarios in which the layered agent structure is expected to yield measurable sustainability benefit alongside operational benefit. The example KPI ranges are based on trends found in the literature concerning agentic predictive maintenance and AI-enhanced Industry 5.0 optimization [2], [8], [13]; they should be interpreted as indicative expectations rather than assured results, as actual performance is significantly influenced by site-specific circumstances.

TABLE II
Illustrative Sustainability-Oriented Use Cases

Use Case	Primary Agentic Function	Sustainability Benefit	Illustrative KPI Trend
Predictive maintenance	Perception + planning agents flag early degradation	Fewer emergency repairs and scrapped parts	Downtime reduced; fewer false alarms
Adaptive energy scheduling	Sustainability optimiser reshapes job timing to grid-carbon signal	Lower peak load and carbon intensity per unit	Energy cost per unit reduced
Dynamic production scheduling	Orchestrator re-sequences jobs under multi-objective score	Reduced idle-running and machine warm-up waste	Throughput steadier; less idle energy
Defect detection & rework routing	Perception agent + digital twin isolate root cause	Less scrap material and rework energy	Scrap rate reduced
Circular material routing	Governance layer enforces reuse-first routing policy	Higher recovered-material utilisation	Virgin-material demand reduced

VI. COMPARATIVE ANALYSIS

Table III compares the suggested sustainable agentic AI method with conventional fixed-logic automation and with interconnected but predominantly reactive Industry 4.0 systems, across seven operational dimensions. The comparison is qualitative and aims to elucidate the architectural contribution of this paper instead of asserting a performance advantage based on benchmarks.

TABLE III
Traditional Automation vs. Industry 4.0 vs. Sustainable Agentic AI

Dimension	Traditional Automation	Industry 4.0 (Connected)	Sustainable Agentic AI (Proposed)
Decision-making	Fixed rule / PLC logic	Rule logic with human dashboards	Autonomous multi-agent reasoning
Adaptability	Low; requires re-programming	Moderate; manual reconfiguration	High; agents replan continuously
Energy optimisation	Not addressed	Monitored, manually acted upon	Embedded in every planning cycle
Human role	Operator executes fixed steps	Operator interprets and decides	Operator supervises and overrides
Explainability	High (deterministic logic)	Moderate (dashboards, alerts)	Requires dedicated explanation layer
Sustainability integration	External audit only	Reported after the fact	Governing constraint on agent action
Scalability across plants	Low; site-specific logic	Moderate	High; agents and policies are modular

The comparison underscores a constant trend: conventional automation provides reliability but lacks flexibility, Industry 4.0 enhances transparency yet retains the human interpretation barrier, and the suggested agentic method sacrifices some certainty and clarity for adaptability and integrated sustainability enhancement. This balancing act is exactly why Section III establishes a specific sustainability governance framework above the decision layer and why Section VII views explainability as a research challenge rather than a resolved issue.

VII. CHALLENGES AND OPEN RESEARCH ISSUES

Although it holds great potential, implementing sustainable agentic AI in actual production environments presents various practical challenges, as outlined in Table IV. These challenges align with the worries highlighted in the examined agentic-AI and Industry 5.0 literature concerning data management, workforce adjustment, and accountability in AI [4], [6], [9].

TABLE IV
Key Deployment Challenges and Potential Mitigations

Challenge	Description	Potential Mitigation
Trust & explainability	LLM-based reasoning agents can produce plans that are difficult for operators to audit	Structured reasoning traces; mandatory human sign-off for high-impact actions
Legacy OT integration	Older PLCs and SCADA systems were not designed for agent-issued commands	Edge adapters and protocol gateways that translate agent intents into legacy control signals

Challenge	Description	Potential Mitigation
Agent energy footprint	Running large reasoning models continuously can itself consume significant energy	Hybrid design: small/rule-based agents for routine cycles, LLM agents only for exceptions
Data privacy & security	Shared knowledge layer aggregates sensitive process and supply-chain data	Federated learning, role-based access control, on-premise model hosting
Standardisation	No common protocol yet exists for cross-vendor agent interoperability	Adoption of open messaging schemas and industry reference architectures
Workforce readiness	Operators and engineers require new skills to supervise autonomous agents	Structured upskilling programmes and human-in-the-loop operating procedures

VIII. FUTURE SCOPE

Multiple paths arise organically from the framework suggested here. Federated agentic learning enables agents across various facilities to exchange acquired scheduling and maintenance strategies while safeguarding raw production data, reflecting the privacy-preserving techniques previously investigated in predictive-maintenance studies [2]. Lightweight, edge-deployable language models may lower the energy consumption of the reasoning agent, tackling one of the issues highlighted in Section VII. Tighter integration with digital product passports and circular-economy data standards would enable the sustainability governance layer to consider material origin and end-of-life management, along with in-plant energy and emissions. Ultimately, human-in-the-loop governance protocols — definitive escalation routes, explanation mandates, and override privileges — must develop concurrently with the technical framework if autonomous agents are to be entrusted with significant production power.

IX. CONCLUSION

This document introduced a tiered reference framework and a functioning MAPE-S cycle for integrating sustainability directly into AI agent decision-making within manufacturing settings, instead of considering environmental performance solely as a metric assessed after production choices are finalized. By establishing a distinct sustainability governance layer positioned above the multi-agent decision layer and integrating energy, emission, and waste components directly into the multi-objective action-selection score, the framework regards sustainability as a governing constraint on par with safety. A qualitative evaluation contrasting conventional automation with Industry 4.0 practices indicates that the main benefit of this method is in ongoing, integrated optimization instead of occasional reporting. Meanwhile, the key challenges involve explainability, integrating legacy systems, and the energy expenses of the agents themselves

REFERENCES

[1] M. A. Farahani et al., "Hybrid Agentic AI and Multi-Agent Systems in Smart Manufacturing," arXiv:2511.18258, Nov. 2025.

[2] "Agentic AI in Smart Manufacturing: Enabling Human-Centric Predictive Maintenance Ecosystems," Applied Sciences, vol. 15, no. 21, art. 11414, Oct. 2025.

[3] J. Lee and H. Su, "Agentic AI for Smart Manufacturing," SSRN Electronic Journal, Aug. 2025.

[4] J. Lee and H. Su, "Agentic AI for Smart Manufacturing," ScienceDirect (Journal of Manufacturing Systems, in press), 2025.

[5] "AI Agents and Agentic AI — Navigating a Plethora of Concepts for Future Manufacturing," ScienceDirect, Sept. 2025.

[6] "Agent Technologies in Smart Manufacturing: A Comprehensive Review of Evolution, Architectures, and Edge–Cloud Deployment," ScienceDirect, Dec. 2025.

[7] P. Ponce et al., "A Collaborative and Sustainable Engineering Design Framework Enabled by Agentic AI," Int. J. of Sustainable Engineering, art. 2624166, Feb. 2026.

[8] "AI-Driven Optimization Techniques for Smart Sustainable Manufacturing in Industry 5.0 Ecosystem: A Comprehensive Review," ScienceDirect, Jan. 2026.

[9] "From Automation to Collaboration: Exploring the Impact of Industry 5.0 on Sustainable Manufacturing," Discover Sustainability, Springer Nature, Apr. 2025.

[10] "A Systematic Review of Industry 5.0 and Sustainability in Manufacturing: Clustering of Concepts, Characteristics, and Technological Enablers," Discover Sustainability, Springer Nature, Jan. 2026.

[11] "Industry 5.0 and Sustainability: An Overview of Emerging Trends and Challenges for a Green Future," ScienceDirect, Aug. 2024.

[12] "Automation and Sustainability — The Impact of AI on Energy Consumption and Other Key Features of Industry 4.0/5.0 Technologies," Applied Sciences (MDPI), vol. 16, no. 5, art. 2550, Mar. 2026.

[13] "Sustainability Awareness in Manufacturing: A Review of IoT Audio Sensor Applications in the Industry 5.0 Era," PMC, National Library of Medicine, 2025.

[14] "Optimizing Sustainable Manufacturing with AI: Exploring Multi-Objective Optimization in Industry 5.0," EELET Journal of Engineering Research, 2025.

[15] "The 2025 AI Agent Index: Documenting Technical and Safety Features of Deployed Agentic AI Systems," MIT AI Agent Index, arXiv, 2025/2026.