Available online at www.sciencedirect.com

ScienceDirect

Procedia Computer Science 287 (2026) 144–152

Procedia Computer Science

www.elsevier.com/locate/procedia

The 13th International Conference on Information Technology and Quantitative Management

An Agentic AI Framework for Collaborative Manufacturing Environments

Stefan-Alexandru Precup[a,*], Gabriela Candea[b], Ciprian Candea[b], Constantin-Bala Zamfirescu[a], Bogdan-Constantin Pirvu[c]

[a]*Computer Science and Electrical Engineering Department, Lucian Blaga University of Sibiu, Emil Cioran 4, 550025 Sibiu, Romania*
[b]*Ropardo SRL, Reconstructiei 2A, 550129 Sibiu, Romania*
[c]*Industrial Engineering and Management Department, Lucian Blaga University of Sibiu, Emil Cioran 4, 550025 Sibiu, Romania*

Abstract

This work introduces an agentic AI framework to address the rigidity of traditional systems while enabling collaborative interaction between humans and intelligent agents. The framework resulted from the synthesis of two research-based use cases: a manual assembly process and a system for advanced material research. Traditional manual assembly support systems often rely on rigid, predefined decision logic that is both time-intensive to develop and difficult to scale. The proposed framework leverages the reasoning capabilities of Large Language Models to provide continuous guidance without requiring manual programming of discrete assembly steps or exhaustive edge-case mapping. When applied to the second use case, the integration of the agentic framework shifts the role of the human operator from a manual orchestrator to a high-level supervisor of a collaborative autonomous system. This approach can significantly reduce time-to-adoption, allowing more effort to be directed toward refining the system. The proposed framework acts as a flexible abstraction that can be adapted to a wide range of scenarios. Its modular design supports the seamless integration of new tools and services into the decision-making process. This paper focuses on the technical implementation of Large Language Models within collaborative production environments, while the broader impact on real-world production remains under evaluation.

Keywords: Agentic AI; Agents; Large Language Models; GPT; Industry 5.0

1. Introduction

The industrial manufacturing landscape is continuously evolving, driven by continuous digitalization and the adoption of new technologies. Industry 4.0, and even more so Industry 5.0, are reshaping traditional production systems by integrating cyber-physical systems, the Industrial Internet of Things, cloud computing, and artificial intelligence into

* Corresponding author.
E-mail address: stefan.precup@ulbsibiu.ro

everyday operations[5]. As a result, manufacturing is no longer limited to isolated production lines but has evolved into interconnected, data-driven ecosystems where humans still play a central role in ensuring operational flexibility and managing complex assembly tasks.

Digital transformation in manufacturing means more than introducing new digital tools. It involves rethinking how data, processes, and people interact across the entire production life cycle, moving beyond simple data collection toward the development of digital twins and integrated simulation platforms. However, initial implementations, such as DiMAT architecture [16], function primarily as a collection of microservices. While it succeeds in connecting experimental data, simulation tools, and domain knowledge within a unified infrastructure, it still relies on the human researcher to act as the orchestrator by manually transferring data between various services.

Technological transformation also significantly affects other areas of operation, including manual assembly processes. As product demand and customization requirements continue to change, operators are required to demonstrate high levels of adaptability to accommodate fluctuating production needs. This creates a significant challenge for the industrial sector, as maintaining agility now depends on the ability to rapidly reskill employees [15]. While employing human instructors to retrain workers is common, this approach has several limitations. It lacks scalability because the instructor-to-worker ratio fluctuates constantly, leading to either staffing shortages or underutilized resources. Human instruction could potentially introduce subjective bias, which can negatively affect process standardization. A recent study of the Ebbinghaus' forgetting curve [8] indicates that, without immediate application, humans forget most of the newly acquired information within 24 hours. These inefficiencies lead to increased downtime and higher production costs for every process or product modification.

Efforts to address the need for rapid reskilling include the development of smart, autonomous assembly stations that collaboratively assist users throughout the assembly process. The impact of such assistive systems was studied in [11], highlighting the importance of an adaptive system to guide the user throughout the assembly process. However, this approach faces two severe limitations: whenever the process description changes, the logic must be rewritten, and new edge cases must be addressed. This results in increased development time and potential downtime in the manufacturing process.

To address these challenges, this paper introduces a widely applicable agentic AI framework, designed to overcome the rigidity of traditional systems. The framework is validated through two distinct use-cases by transitioning the aforementioned applications from simple, reactive systems that rely on hard-coded decision logic and human input into proactive, agent-based systems. The remainder of the paper is structured as follows: Section 2 provides a literature review of existing agentic approaches and applications, Section 3 introduces the agentic framework and presents two agentic use cases, and Section 4 discusses the limitations of the current work.

2. Related Work

Continuous efforts are being made to reduce the cognitive load experienced by users during manual assembly operations, ranging from intention recognition approaches for robotic assistance [17] to methods for predicting the next assembly step [10]. In [11], the authors introduced an adaptive assembly station designed for manual operations that is capable of adapting to individual users. When a new user initiates an assembly process, the station retrieves the user's characteristics and automatically adjusts parameters such as workstation height, voice interface, and visual settings to better suit the operator. The system then provides instructions regarding the assembly process to be performed. During operation, the station continuously monitors the user and their actions and, through a cognitive control unit, determines the appropriate course of action. The system is capable of detecting more than 20 types of incorrect actions and, once such an action is identified, it provides guidance to the user on how to correct it. When a correct action is detected, an AI-based module within the control unit predicts the next assembly step based on previously observed actions and user characteristics. Throughout the process, the station continuously monitors the operation and provides feedback to support and assist the user. Experimental validation revealed that most participants achieved high knowledge retention regarding product configuration, highlighting the importance of assembly stations in the industrial environment.

Advances in computing power have enabled researchers to develop increasingly computationally intensive software. In the field of artificial intelligence, this progress has facilitated the emergence of resource-intensive large language models (LLM). In other domains, such as materials engineering, significant computational resources are required for the design and simulation of new materials with specified mechanical properties or with predefined behavior under particular conditions.

Advanced material processing involves the formulation, modeling, and simulation of complex material transformation processes using computational and data-driven methods. The authors of [16] introduce DiMAT, a framework aimed at digital modelling and simulation for design, processing, and manufacturing of advanced materials. The framework aims to provide a comprehensive suite of tools for the materials manufacturing sector. It includes a data and assessment suite that integrates databases, knowledge acquisition frameworks, and life cycle assessment capabilities. In addition, a modeling and design suite incorporates a materials design framework together with a modeling toolkit that leverages AI for analysis and prediction, as well as a materials designer intended for the engineering of new composites. Furthermore, the framework includes a materials simulation and optimization suite that integrates digital twin technology for process optimization. Overall, the framework is intended to support SMEs that require scalable and affordable open-source software solutions for their operations.

The emergence of LLMs has marked a significant shift in how businesses operate, influencing both decision-making processes and operational workflows [1]. To support this transition, new tools had to be developed. Two tools have gained prominence: Retrieval-Augmented Generation (RAG) and the Model Context Protocol (MCP), which together enable the development of agentic AI systems. RAG enables LLMs to retrieve relevant information from external knowledge bases and integrate it into the generated responses, while MCP provides a standardized protocol for agents to discover, access, and interact with external tools and services.

Agentic AI is currently regarded as the next step in artificial intelligence [9, 13]. As [9] notes, its key advantages over traditional AI include the ability to autonomously make decisions, continuously learn and adapt, and proactively act rather than merely respond to predefined inputs.

The authors of [13] discuss how agentic AI is transforming industrial manufacturing from knowledge retrieval and automated documentation to multi-modal agents capable of providing diagnostics or decision support based on sensor data and specific domain knowledge, highlighting the narrow scope of traditional AI. The paper notes that key challenges for successful deployment include alignment with manufacturing processes, interpretability and explainability, accountability, and, critically, resistance to adoption by the workforce.

In [6], the authors introduce a specialized agent framework tailored for the semiconductor manufacturing sector. The system is designed to respond to user queries by synthesizing information from eight distinct data sources, including real-time monitoring, process control, logs, maintenance data, and technical manuals. To generate insights, the system integrates this multi-source knowledge with semiconductor-specific analysis tools. The authors claim that through this integrated approach, the system is capable of real-time process optimization and task automation, leading to increased productivity, quality, and overall operational efficiency.

Another agentic AI approach, presented in [2], has been tested and validated on a real-world scenario in a ceramic tile manufacturing facility. The system employs specialized agents for each stage: sensing agents to acquire data, reasoning agents for analytics, and action agents responsible for resource allocation and maintenance scheduling. Coordination agents oversee inter-agent communication, ensuring efficient task distribution and system optimization. The proposed system improved data quality and achieved sub-30-second anomaly detection for 94% of extreme events, leading to increased equipment efficiency and reduced energy costs.

These studies showcase the potential impact of agentic AI integration in the two use cases addressed by this paper. Moreover, the framework introduced in the next section can also be employed to represent the applications analyzed in the related work.

3. A framework for agentic systems

3.1. The Framework

Many agentic solutions developed today involve a user interacting with a system, while an agent continuously monitors this system and responds to environmental changes. In Figure 1, we introduce a generalized framework for agentic AI workflows consisting of three core entities: the user, the system, and the agent. This model emphasizes the directional interactions and functional relationships shared between them.

The creation of an agent starts with the definition of the process description, a preliminary step in which all relevant information must be defined. This information must include the objective of the system and should be in an easily readable format for both humans and LLMs. Additional details are to be defined on a specific case-by-case basis and can include implementation details, constraints, or success metrics. The process description will serve as the blueprint from which the agent is instantiated, ensuring that its behavior aligns with the system's objectives. Optionally, a memory and learning module can be integrated to store additional information, including episodic data, for later use and improved agent performance.

Through MCP, the agent discovers the integrated services it can use to reach the objective. In instances where required services are not exposed via the MCP, they must be explicitly defined within the process description during the initialization phase. Integrated services available to the agent include internal core functions, external APIs, and other specialized agents. These components are not universally required and can be selected based on the specifics of the tasks. For instance, sensors and actuators facilitate the acquisition of real-time environmental data, while external systems provide analytics, predictions, or event-driven triggers. Additionally, specialized agents may be integrated to collaboratively tackle domain-specific problems and provide expert reasoning.

The physical/virtual system serves as the interface through which the user collaborates with the agentic framework. Physical configurations may encompass machinery, workstations, or complex cyber-physical systems, while virtual environments include web-based platforms, virtual reality, or specialized software applications. For the user, this layer facilitates interaction and provides dedicated channels for instructions and feedback, while for the agentic framework, it provides contextual data about the environment to the agent and the means to execute its commands.

The user must interact with the physical or virtual system to fulfill a specific task, where each user action alters the state of the environment. The system communicates these state changes to the agent, which then evaluates the updated context against the predefined process description. Through the available integrated services, the agent determines the necessary operations and issues the resulting commands back for execution. It should be noted that the system may alter the environment by itself, either through the commands issued by the agent or other internal events, without the need for an action from the user. This interaction flow is also showcased in Figure 1.

Although the proposed architecture is generic, the effectiveness of an agentic system depends on the quality of the process description and the services provided. To illustrate this aspect more clearly, the following subsections will present two real-world use cases that implement this architecture.

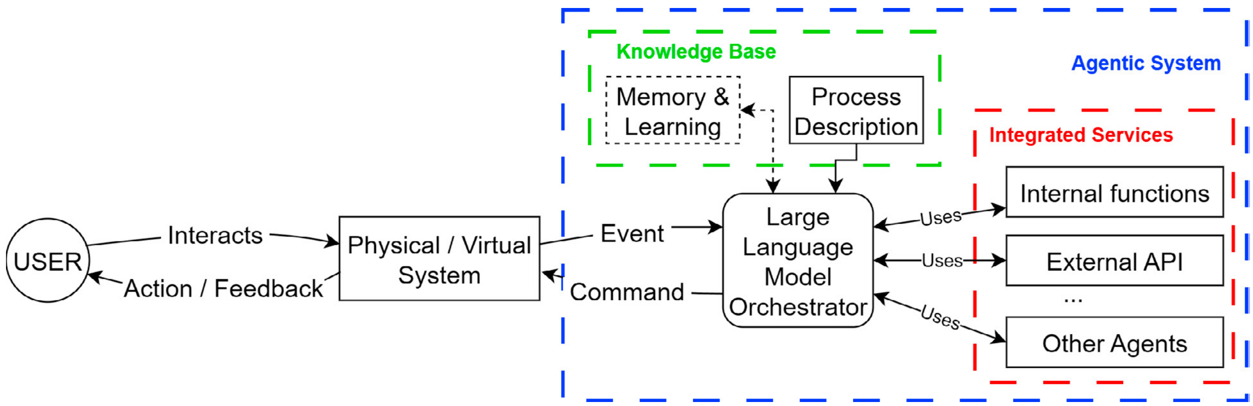

Fig. 1. General Framework for a collaborative agentic AI workflow in factories.

3.2. Usecase: Assembly Station

The first real-world use case is in manual assembly, where the worker must use the workstation (Figure 2 (b)) to assemble a modular product (Figure 2 (a)). The assembly station follows a human-centric design. It is capable of adjusting itself to the user currently interacting with it by adjusting the table height and by modifying how visual and auditory feedback are delivered. Through a touchscreen, the user can interact with the station to go back one step or

(a) (b)

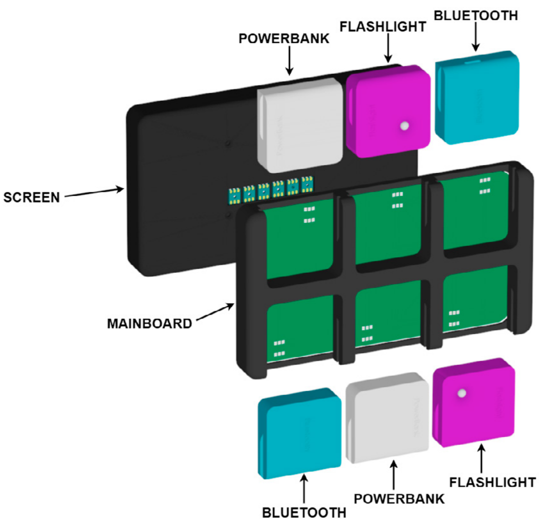

Fig. 2. (a) The target product[11]. (b) The training station[11].

repeat the instruction for the current assembly step. The cameras facilitate the acquisition of environmental data. The station was designed to be modular and, if needed, to allow for further integration of sensors and actuators.

The product was developed in alignment with Industry 5.0, which focuses on highly customizable, low-volume production batches[3]. The product is modular and allows for any configuration of the three available modules: the powerbank, the flashlight, and the Bluetooth module. These modules, along with the screen, are mounted onto a central mainboard. To showcase the complexity of even a simple configuration, a configuration using two modules of each type has been selected. Since there are no assembly constraints, the components can be installed in any order. This results in 90 possible ways to assemble the product. A more detailed description of the product and the assembly station can be found in [11].

The implementation of the framework for this use case is shown in Figure 3, highlighting the collaboration between the human worker and the agentic system through the assembly station. The agentic system uses four integrated services. Human feature extraction and object detection are built into the assembly station, using its cameras to collect data about the environment. The text-to-speech service relies entirely on external providers to customize the auditory feedback for the worker. To tailor the system to the user, a next-step prediction service was also included. By analyzing previously performed assembly steps and current environmental data, the service can predict the most suitable next step for the user. This service can run independently on either a separate computer or the assembly station itself.

To enable the agentic system to function as the brain of the station and utilize the integrated services to their full potential, a process description was required. This description follows the same specification introduced in [12], which allows for the easy definition of assembly steps and their associated constraints.

The workflow operates as a continuous loop, having the user assemble the product or interact with the assembly station while the agent monitors the assembly state. By analyzing data from the assembly station, the integrated services, and the process description, the agent can determine if the current assembly state is valid. If correct, the agent utilizes the next-step prediction service to identify the most suitable next assembly step. Finally, the agent sends this information to the station, which delivers visual and auditory instructions to the user.

If the assembly is incorrect, there are numerous ways the user could have made a mistake. As highlighted in [11], over 20 incorrect actions have been identified so far. This is where the agentic approach proves its value because the agent can identify and handle scenarios as they occur instead of requiring a programmer to manually code for every

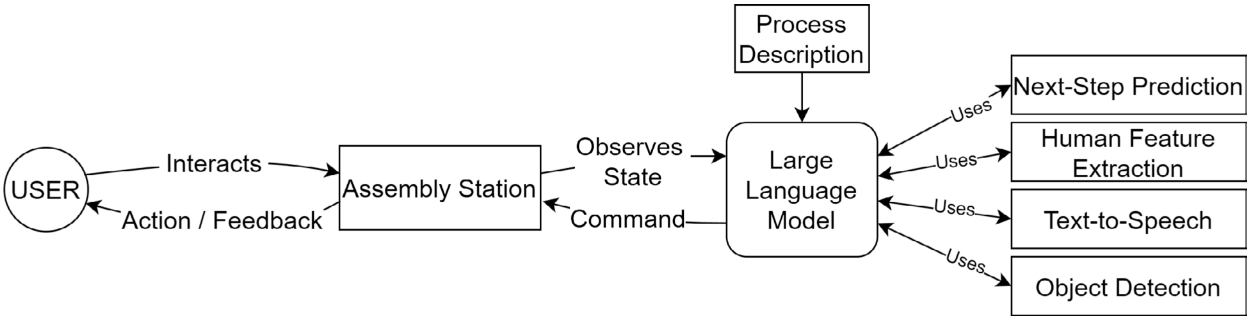

Fig. 3. Agentic workflow employed in the assembly station.

edge case, which is very time-consuming. Once a problem is detected, the agent sends the appropriate command to the assembly station to provide the user with corrective feedback.

3.3. Usecase: Agentic DiMAT

The framework's second application facilitates the architectural migration of the DiMAT platform from a microservices-oriented configuration toward an autonomous agentic system (Agentic DiMAT). The platform is organised into three functional suites: data and assessment; modeling and design; and simulation and optimization. The suites and their toolkits are thoroughly detailed in [16].

Because DiMAT already integrated LLMs within several toolkits, it provided a starting point that eased the modernization of the DiMAT platform from the existing toolkit-based architecture into a collaborative system of autonomous, LLM-orchestrated agents. In this design, the three functional suites retain their domain specialization, but each toolkit is wrapped as an autonomous agent capable of receiving goals, planning actions, and reporting results. A central LLM Orchestrator coordinates the agents, translating high-level user objectives, such as "design a polymer composite with specified mechanical properties", into executable sequences of agent invocations.

Figure 4 shows how the agentic framework works within the Agentic DiMAT platform by illustrating the interactions among the user, the platform, and the agentic system. The user sets a high-level goal through the platform, which is then converted into commands for the LLM Orchestrator. The Orchestrator breaks down the goal and assigns tasks to specialized agentic toolkits: Data & Assessment Agents, responsible for querying materials databases and knowledge graphs; Modelling & Design Agents, which support AI-driven predictions and materials engineering activities; and Simulation & Optimization Agents, which enable digital twin–based process control. This distribution of responsibilities reflects the three-part architectural structure of the DiMAT platform [16, 14]. A Memory & Learning module, powered by process descriptions, helps the system remember context across interactions and improve its strategies over time. Notably, edge devices running Small Language Models (SLMs) expand the agentic system's capabilities from the cloud to the manufacturing floor, enabling low-latency, privacy-protected decision support at the point of operation.

Table 1 highlights the evolutionary differences between DiMAT and Agentic DiMAT. While DiMAT relied primarily on loosely coupled microservices that exposed functionality via REST interfaces, the new architecture introduces autonomous agents as the primary unit of composition, capable of reasoning, planning, and interacting with external tools. Additionally, the architecture moves from stateless services and direct database queries to memory-augmented reasoning, where agents access contextual knowledge through retrieval mechanisms and ontologies. Tool integration is extended through the MCP, which standardizes how agents access external capabilities. Decision-making shifts toward agent autonomy with human-in-the-loop supervision, while deployment evolves from a cloud-only model to a hybrid edge–cloud infrastructure. Overall, Agentic DiMAT transforms the platform from a set of distributed services into a collaborative multi-agent system designed to support intelligent manufacturing assistance.

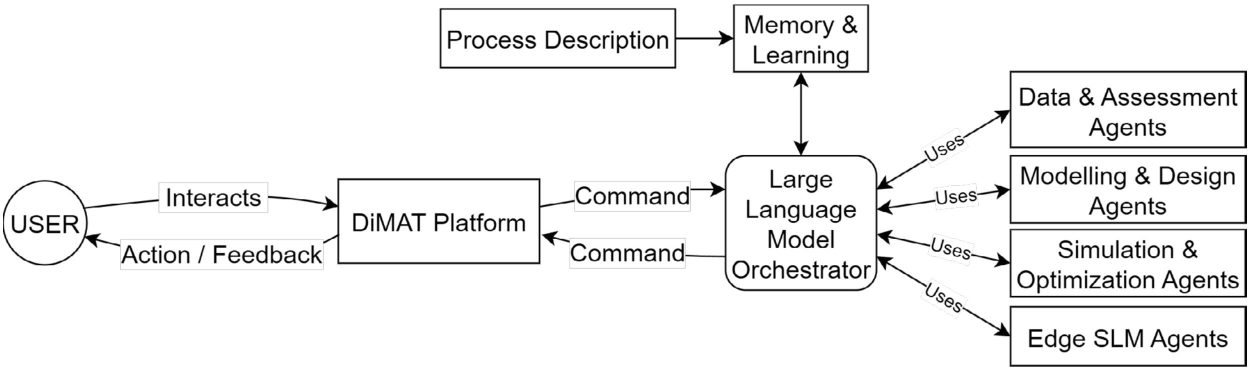

Fig. 4. Agentic DiMAT workflow.

Table 1. Evolution from DiMAT (microservices) to Agentic DiMAT Architecture.

Dimension	DiMAT	Agentic DiMAT
Unit of Composition	Toolkit (microservice)	Autonomous Agent
Coordination	Manual / scripted	LLM Orchestrator
State Management	Stateless REST APIs	Memory-augmented (RAG)
Knowledge Access	Direct DB queries	Ontology-guided reasoning
Tool Integration	REST / MQTT APIs	MCP
Decision Making	Human-driven	Agent-autonomous + HITL
Deployment Model	Cloud-only	Hybrid edge-cloud (LLM/SLM)
Safety Pattern	Input validation	Simulate-before-actuate

4. Discussion and Further work

This work introduced a general Agentic AI framework validated against two use-cases: manual assembly operations and a software system for material processing. The two widely different use cases demonstrate that this framework provides a domain-agnostic architecture capable of modeling a broad range of agent-based systems while facilitating collaborative interactions with human users.

Considering the manual assembly use case, psychological and empirical research highlight the critical role of autonomy within the user workspace [4]. The agentic approach addresses this requirement by providing users with continuous, noninvasive guidance throughout the assembly process. Another critical challenge was the transition from a static, rule-driven architecture to a dynamic system capable of facilitating diverse assembly scenarios and potentially supporting a broader range of products. While the system's performance on a wider variety of assembly tasks has not yet been validated, the elimination of complex decision logic has already reduced software development time. Edge cases that should have been previously treated by the developers are now handled by the agentic system, providing them with more time to test and refine the final system. By adopting the agentic approach, the system can now extend its capabilities by easily integrating new tools or services.

However, by employing LLMs as the brain of the decision system, the system effectively becomes a black box, and the transparency of the decision is lost, providing no explanation whatsoever on how the agent arrived at a certain decision [18]. Alongside transparency, the system also loses consistency, having no guarantee that the agent will make the same decisions when faced with similar or even the same inputs. Furthermore, as in any real-time system, the time from input to output is critical. As LLMs can take a long time to reason before outputting tokens, the use case, in its current form, can not transition from laboratory to industrial application. To address all these issues, we seek to merge this framework with the SOAR-based approach from [11], leveraging SOAR's explainability with the flexibility provided by LLMs.

The integration of agentic AI in the second use case highlights the transition from the traditional microservice architecture to an agentic-based one. It redefines the role of the human from a manual orchestrator to a high-level supervisor of a self-correcting system. In the previous implementation [16], users manually transferred data between the services. In the agentic approach, the LLM Orchestrator autonomously determines the required sequence of operations, dynamically routing specific sub-tasks to domain-specialized agents for execution.

A major bottleneck to be addressed is the scalability of the system. In complex workflows, the orchestrator must maintain a coherent task dependency graph while managing concurrent agent executions. The current design addresses this limitation through domain-specific prompt engineering and RAG-augmented context; however, a formal evaluation of orchestration accuracy across diverse workflow types has not yet been conducted.

As in any system modification, the transition requires significant changes in how researchers and engineers interact with the platform. Shifting from manual toolkit invocation, where users retain full control over each step, to autonomous agent orchestration necessitates establishing trust in the system's decision-making capabilities [7]. To address this, humans are still playing an important role in decision-making. However, identifying the correct approval thresholds remains an empirical issue that can only be resolved through sustained pilot deployment with real users across industrial pilots.

Future work will involve empirical validation through pilot deployments and benchmarking the autonomous workflow completion rates against manual orchestration baselines. More thorough studies are required to further assess the safety and robustness of the presented use cases.

References

[1] Arslan, M., Munawar, S., Cruz, C., 2024. Sustainable digitalization of business with multi-agent rag and llm. Procedia Computer Science 246, 4722–4731. doi:https://doi.org/10.1016/j.procs.2024.09.337. 28th International Conference on Knowledge Based and Intelligent information and Engineering Systems (KES 2024).

[2] Fernández-Miguel, A., Ortíz-Marcos, S., Jiménez-Calzado, M., Fernández del Hoyo, A.P., García-Muiña, F.E., Settembre-Blundo, D., 2025. Agentic ai in smart manufacturing: Enabling human-centric predictive maintenance ecosystems. Applied Sciences 15. doi:10.3390/app152111414.

[3] Hasani, N., Hosseini, A., Ashjazadeh, Y., Diederichs, V., Ghotb, S., Riggio, M., Hansen, E., Nasir, V., 2025. Outlook on human-centred design in industry 5.0: towards mass customisation, personalisation, co-creation, and co-production. International Journal of Sustainable Engineering 18, 2486343. doi:10.1080/19397038.2025.2486343, arXiv:https://doi.org/10.1080/19397038.2025.2486343.

[4] Johannsen, R., Zak, P.J., 2020. Autonomy raises productivity: An experiment measuring neurophysiology. Frontiers in Psychology Volume 11 - 2020. doi:10.3389/fpsyg.2020.00963.

[5] Karic, H., Arockiadoss, A.S., Schilberg, D., 2025. Smart factory and industry 4.0: A survey on advancements, technologies, methods and perspectives of digital transformation in manufacturing. Human Aspects of Advanced Manufacturing, Production Management and Process Control 184, 1–11. doi:10.54941/ahfe1006444.

[6] Lin, C.Y., Tsai, T.H., Tseng, T.L., 2025. Generative ai for intelligent manufacturing virtual assistants in the semiconductor industry. IEEE Robotics and Automation Letters 10, 4132–4139. doi:10.1109/LRA.2025.3544506.

[7] Liu, B., 2021. In ai we trust? effects of agency locus and transparency on uncertainty reduction in human–ai interaction. Journal of Computer-Mediated Communication 26, 384–402. doi:10.1093/jcmc/zmab013, arXiv:https://academic.oup.com/jcmc/article-pdf/26/6/384/41139653/zmab013.pdf.

[8] Murre, J.M.J., Dros, J., 2015. Replication and analysis of ebbinghaus' forgetting curve. PLOS ONE 10, 1–23. doi:10.1371/journal.pone.0120644.

[9] Pati, A.K., 2025. Agentic ai: A comprehensive survey of technologies, applications, and societal implications. IEEE Access 13, 151824–151837. doi:10.1109/ACCESS.2025.3585609.

[10] Precup, S.A., Matei, A., Walunj, S., Gellert, A., Plociennik, C., Zamfirescu, C.B., 2023a. Collaborative exploitation of various ai methods in adaptive assembly assistance systems. Procedia Computer Science 221, 1170–1177. doi:https://doi.org/10.1016/j.procs.2023.08.103. tenth International Conference on Information Technology and Quantitative Management (ITQM 2023).

[11] Precup, S.A., Pirvu, B.C., Gellert, A., Zamfirescu, C.B., 2025. Cognitive control units in smart manufacturing: Insights from a soar-based assistance system. IEEE Access 13, 196167–196180. doi:10.1109/ACCESS.2025.3633609.

[12] Precup, S.A., Walunj, S., Gellert, A., Plociennik, C., Antony, J., Zamfirescu, C.B., Ruskowski, M., 2023b. Recognising worker intentions by assembly step prediction, in: 2023 IEEE 28th International Conference on Emerging Technologies and Factory Automation (ETFA), pp. 1–8. doi:10.1109/ETFA54631.2023.10275423.

[13] Ren, Y., Liu, Y., Ji, T., Xu, X., 2025. Ai agents and agentic ai–navigating a plethora of concepts for future manufacturing. Journal of Manufacturing Systems 83, 126–133. doi:https://doi.org/10.1016/j.jmsy.2025.08.017.

[14] Stavropoulou, G., Tsitseklis, K., Mavraidi, L., Chang, K.I., Zafeiropoulos, A., Karyotis, V., Papavassiliou, S., 2024. Digital twin meets knowledge graph for intelligent manufacturing processes. Sensors 24. doi:10.3390/s24082618.

[15] Sundaram, S., 2025. Innovating for the future: Sustainable corporate strategies for workforce upskilling and reskilling. International Journal of Research in Entrepreneurship & Business Studies 6, 33–46.

[16] Tsitseklis, K., Morand, L., Mateo-Casali, M.A., Stavropoulou, G., Mavraidi, L., Voikos, S., Nahshon, Y., Büschelberger, M., De Andres, P., Lavasa, A., Karakostas, A., Papadimitriou, I., Cosma, L., Gomez, J., De La Rosa Ramírez, H., Zafeiropoulos, A., Fraile, F., Boza, A.,

Candea, C., Papavassiliou, S., 2025. A reference architecture for digital transformation of smes in the manufacturing domain, in: 2025 IEEE International Conference on Engineering, Technology, and Innovation (ICE/ITMC), pp. 1–10. doi:10.1109/ICE/ITMC65658.2025.11106592.

[17] Walunj, S., Pahlevannejad, P., Sintek, M., Saborio, J.C., Garg, P., Tavakoli, H., Plociennik, C., Grossmann, N., Bernardi, A., Ruskowski, M., Hertzberg, J., Dengel, A., 2025. Context-aware robotic assistance for workers using intention recognition and semantic digital twin, in: Singh, D., van 't Klooster, J.W., Tiwary, U.S. (Eds.), Intelligent Human Computer Interaction, Springer Nature Switzerland, Cham. pp. 92–104.

[18] Yin, F., Zhong, M., Ru, Z., 2025. Exploring explainability in large language models. Preprints doi:10.20944/preprints202503.2318.v1.