Educational Innovation Research, 2026, Volume 4, Issue 7
https://bbwpublisher.com/journal/EIR
Online ISSN: 3029-1852
Print ISSN: 3029-1844

Pathways and Challenges of Artificial Intelligence-Enabled Manufacturing Transformation and Upgrading

Zhiqiang Li*

Nanjing University, Nanjing, Jiangsu 210008, China

Author to whom correspondence should be addressed.

Copyright: © 2026 Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY 4.0), permitting distribution and reproduction in any medium, provided the original work is cited.

Abstract: Artificial intelligence can improve manufacturing decisions, yet model performance alone does not establish operational value. This conceptual review examines how manufacturers can convert AI applications into improvements in quality, flexibility and organizational capability. The proposed pathway connects problem definition, contextual data, model selection, workflow integration and evaluation. Industrial ontologies and bounded AI agents are considered as possible means of connecting information with action. Fragmented data, uncertain transfer between operating conditions, integration costs and unclear decision authority emerge as interdependent barriers. The paper proposes staged deployment with explicit baselines, operator participation and limits on automated action. The framework guides empirical investigation without estimating a universal productivity effect.

Keywords: Artificial intelligence; Manufacturing transformation; Industrial ontology; AI agents; Human oversight

Online publication: July 26, 2026

1. Introduction

Manufacturing transformation changes how firms coordinate production and develop capabilities to deliver quality, accommodate product variation and improve processes. Smart manufacturing connects production assets with information and responsive decisions [1]. AI contributes methods for learning patterns and estimating outcomes, but manufacturing applications depend on problem formulation and operating conditions [2]. Industrial AI consequently requires attention to the system in which an algorithm operates [3].

The question is how task performance becomes operational improvement. A defect classifier may identify an anomaly without determining whether a batch should be held. A maintenance forecast may arrive too late to change a service plan. This paper examines that gap between an analytical output and an executable decision. Upgrading means a durable improvement in production capability or customer value; installing AI software alone does not establish success.

2. Scope and analytical approach

The paper presents a narrative synthesis of selected peer-reviewed studies and official frameworks on manufacturing AI, integration, and adoption. Selection was purposive, without a systematic search protocol or exhaustive coverage. No enterprise measurements or experiments are reported. The proposed pathway and measures require testing in specific production settings.

3. Pathways for manufacturing transformation

3.1. Define the decision and its operating boundary

Implementation should begin with a recurring decision that has an owner and measurable consequences. For inspection, outcomes include escaped defects and unnecessary rejection. For maintenance, avoidable downtime can be assessed alongside additional service work. Measurements of the current procedure establish the baseline. Error costs, action deadlines and permitted responses should be specified before selecting a model.

The cyber-physical systems architecture of Lee and colleagues connects data acquisition, computation, and decision support [4]. A plant still needs to specify when observations become available, which system receives a recommendation and how actions are recorded. Feedback is needed to assess production benefit.

3.2. Organize data around manufacturing meaning

Sensor values require recoverable context. A temperature reading from a spindle has a different meaning from an ambient measurement. A usable data foundation needs stable asset and product identifiers, consistent units, timestamps and links to process versions. Historical cleaning has limited value if incompatible records continue to enter the system.

An ontology defines entities and relationships; a knowledge graph represents instances such as a machine, batch and associated inspection. Buchgeher and colleagues review manufacturing knowledge graphs [5]. This paper proposes starting with relationships required for one decision and extending the model as needed. Maintaining mappings and resolving conflicting definitions remain organizational responsibilities. A graph cannot by itself correct inconsistent source data.

3.3. Match models to evidence and constraints

Method selection should reflect evidence and error consequences. Classification may suit inspection with representative labels. For rare failures, expert rules or physical models may supplement limited labels. Karniadakis and colleagues review the incorporation of physical knowledge into machine learning [6]. Its usefulness depends on valid physics and boundary conditions.

Digital twins connect physical entities with virtual representations and supporting data. Tao and colleagues discuss modelling, interaction and validation as industrial implementation concerns [7]. A twin may help examine candidate interventions when its fidelity suits the decision. Simulation does not establish that an intervention will improve production. Validation should cover relevant changes in tooling, materials and product mix, and should compare the proposed method with a simpler feasible alternative.

3.4. Integrate assistance with controlled execution

Recent research examines large language models in manufacturing [8]. Here, an AI agent uses a model to

select and invoke permitted tools. A proposed maintenance agent could retrieve a procedure, inspect alarms and prepare a work order. An ontology would define relevant asset and failure relationships; execution software would check them before acting. This is a design proposal, not a deployed result.

Execution requires checks against permissions, process limits and current production conditions. Unapproved actions should be escalated. Jarrahi's analysis of human and AI complementarity helps frame the relationship between analytical assistance and human judgment [9]. In manufacturing, the interface should show supporting evidence and allow rejection. Recording the reason for an override can reveal missing context or an unsuitable workflow.

3.5. Evaluate operation before replication

Deployment creates continuing obligations as data pipelines and upstream systems change. Sculley and colleagues identify technical debt arising from dependencies and other system-level features of machine learning [10]. Evaluation should therefore include unavailable recommendations, late responses and failed actions, as well as prediction errors. **Table 1** summarizes the proposed evaluation logic. Its measures are examples, not universal acceptance thresholds.

Comparison can use a similar line, staged deployment, or repeated observations. Concurrent process changes must be recorded to avoid attributing all improvement to AI. Replication should retain local validation of equipment and operating conditions.

Table 1. Proposed pathway and evaluation measures

Implementation step	Evidence required	Example measure
Define the decision	Baseline and responsible owner	Escaped defects or avoidable downtime
Prepare contextual data	Asset links and provenance	Missing or ambiguous context
Validate models	Tests across operating conditions	Error cost by production regime
Integrate the workflow	Permissions and operator feedback	Completed actions and overrides
Monitor and replicate	Operating history and local checks	Benefit after support costs

4. Challenges to sustained improvement

Data scarcity interacts with changing production conditions. A plant may hold extensive sensor records but few labelled examples of a rare failure. Historical interventions complicate interpretation. For the proposed pathway, testing should separate observations by time, equipment, or batch when deployment requires such transfer. Closely related observations should remain together to reduce leakage between training and testing. Context is needed to distinguish model deterioration from changes in products or procedures.

Organizational readiness affects the ability to sustain an application. Jöhnk and colleagues identify factors relevant to preparation for AI adoption [11]. For manufacturing, this suggests assessing whether process specialists and information and operational technology staff can maintain the workflow jointly. A pilot supported by an exceptional project team may exceed normal staffing capacity. Data stewardship, validation, operator training and continuing support consequently belong in the implementation budget.

Small and medium-sized manufacturers face particular adoption conditions, examined by Hansen and Bøgh [12]. A practical response is to evaluate one complete decision workflow before purchasing a broad platform. Buying a model or cloud service leaves integration work to be funded. Commercial assessment

should include recurring fees, connectivity, changes to production routines and the cost of changing suppliers. These costs need comparison with both expected benefits and a simpler baseline solution.

Accountability becomes harder when several services contribute to an action. The NIST AI Risk Management Framework organizes risk work around governing, mapping, measuring and managing [13]. Applied here, this means naming a decision owner, retaining model and data versions, and defining when an application should be suspended. For agents, authorization must be checked at the tool boundary. A natural language instruction cannot by itself authorize changes to production settings. Safe operation also depends on the surrounding industrial control architecture.

Efficiency targets can overlook effects on workers and the environment. The European Commission's Industry 5.0 report emphasizes human centricity, sustainability and resilience [14]. Vinuesa and colleagues examine enabling and inhibiting relationships between AI and sustainable development [15]. These sources support broader evaluation; they do not prove that a factory application saves energy. Assessment should distinguish energy per acceptable unit from total use, include computing where material, and examine changes in operator workload and discretion.

5. Discussion and research implications

The framework's central proposition is that a weak connection between evidence, decisions and action can limit the value of a capable model. Better classification may have little effect when product context is missing or the action window has closed. Conversely, a modest model may be useful when evidence, intervention and feedback are well connected. These propositions require empirical comparison.

Future studies could compare workflows with and without a maintained semantic layer, holding the model and available information constant. Outcomes should include decision completion, investigation time and incorrect action costs. Agent studies should examine recovery from failed tool calls and outdated records. Longer evaluations should report maintenance effort after pilot support ends, production conditions and staffing changes.

The review is limited by purposive selection and evidence from different industries and levels of analysis. Several sources concern architectures or governance rather than controlled production evaluations. The pathway may need adaptation for continuous processes or regulated production. It is not a validated maturity model.

6. Conclusion

AI-enabled manufacturing upgrading requires a verifiable connection between analytical outputs and production outcomes. The proposed pathway links contextual data and model validation with controlled execution and continuing evaluation. Ontologies and bounded agents may help organize context and coordinate tools, provided their information and permissions remain reliable. Assessment should use measured baselines and include support costs and workforce effects. Empirical testing is needed before general productivity claims can be justified.

Disclosure statement

The author declares no conflict of interest.

References

[1] Kusiak A, 2018, Smart Manufacturing. International Journal of Production Research, 56(1–2): 508–517.

[2] Wuest T, Weimer D, Irgens C, et al., 2016, Machine Learning in Manufacturing: Advantages, Challenges, and Applications. Production & Manufacturing Research, 4(1): 23–45.

[3] Lee J, Davari H, Singh J, et al., 2018, Industrial Artificial Intelligence for Industry 4.0-Based Manufacturing Systems. Manufacturing Letters, 18: 20–23.

[4] Lee J, Bagheri B, Kao HA, 2015, A Cyber-Physical Systems Architecture for Industry 4.0-Based Manufacturing Systems. Manufacturing Letters, 3: 18–23.

[5] Buchgeher G, Gabauer D, Martinez-Gil J, et al., 2021, Knowledge Graphs in Manufacturing and Production: A Systematic Literature Review. IEEE Access, 9: 55537–55554.

[6] Karniadakis GE, Kevrekidis IG, Lu L, et al., 2021, Physics-Informed Machine Learning. Nature Reviews Physics, 3(6): 422–440.

[7] Tao F, Zhang H, Liu A, et al., 2019, Digital Twin in Industry: State-of-the-Art. IEEE Transactions on Industrial Informatics, 15(4): 2405–2415.

[8] Maghanaki M, Shahin M, Chen FF, 2026, Large Language Models in Manufacturing: A Comprehensive Review. The International Journal of Advanced Manufacturing Technology, 145(11–12): 7601–7627.

[9] Jarrahi MH, 2018, Artificial Intelligence and the Future of Work: Human-AI Symbiosis in Organizational Decision Making. Business Horizons, 61(4): 577–586.

[10] Sculley D, Holt G, Golovin D, et al., 2015, Hidden Technical Debt in Machine Learning Systems. Advances in Neural Information Processing Systems 28.

[11] Jöhnk J, Weißert M, Wyrtki K, 2021, Ready or Not, AI Comes—An Interview Study of Organizational AI Readiness Factors. Business & Information Systems Engineering, 63(1): 5–20.

[12] Hansen EB, Bøgh S, 2021, Artificial Intelligence and Internet of Things in Small and Medium-Sized Enterprises: A Survey. Journal of Manufacturing Systems, 58: 362–372.

[13] Tabassi E, 2023, Artificial Intelligence Risk Management Framework (AI RMF 1.0) (NIST AI 100-1). National Institute of Standards and Technology.

[14] Breque M, De Nul L, Petridis A, 2021, Industry 5.0: Towards a Sustainable, Human-Centric and Resilient European Industry. Publications Office of the European Union.

[15] Vinuesa R, Azizpour H, Leite I, et al., 2020, The Role of Artificial Intelligence in Achieving the Sustainable Development Goals. Nature Communications, 11: 233.

Publisher's note

Bio-Byword Scientific Publishing remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.