Knowledge Graphs and Neurosymbolic AI for Manufacturing Intelligence: Semantic Reasoning, Ontology-Driven Analytics, and the Integration of Symbolic and Neural Intelligence

Author: Jesie Pinkman

Abstract

The integration of artificial intelligence into manufacturing has generated powerful data-driven models — deep reinforcement learning policies for process control, convolutional networks for quality inspection, recurrent architectures for demand forecasting — yet these models operate as black boxes divorced from the rich domain knowledge that manufacturing engineers have accumulated over decades. **Knowledge graphs** — structured representations of entities (machines, products, materials, processes) and their relationships — and **neurosymbolic AI** — hybrid architectures that combine the pattern recognition power of neural networks with the logical reasoning capabilities of symbolic AI — have emerged as complementary paradigms that inject domain knowledge into AI systems, enabling reasoning, explanation, and causal understanding that purely data-driven approaches cannot provide. This review provides a comprehensive synthesis of knowledge graphs and neurosymbolic AI for manufacturing intelligence, examining knowledge graph construction and ontology engineering for manufacturing, graph neural networks for manufacturing analytics, neurosymbolic reasoning for fault diagnosis and process optimization, the integration of knowledge graphs with digital twin platforms, and the synthesis of knowledge-driven AI with the four preceding AI frameworks for smart manufacturing (Physics-Informed RL-MPC, Adaptive Manipulation, Quality Intelligence Architecture, and Neuromorphic Industrial Intelligence Architecture). We further connect these advances to industrial optical sensing technologies — precision 3D surface metrology and four-dimensional thermal imaging — demonstrating how structured knowledge and neural-symbolic integration enhance perceptual intelligence in manufacturing. A central contribution is the articulation of an integrated **Knowledge-Augmented Manufacturing Intelligence (KAMI)** framework that unifies knowledge graphs, neurosymbolic reasoning, and deep learning for trustworthy, explainable, and knowledge-driven manufacturing AI.

Keywords: Knowledge Graphs; Neurosymbolic AI; Ontology Engineering; Graph Neural Networks; Semantic Reasoning; Manufacturing Intelligence; Smart Factory; Digital Twin; Explainable AI; Knowledge-Augmented AI

---

1. Introduction

Modern manufacturing enterprises generate and consume vast quantities of structured and unstructured data — design specifications, process parameters, quality measurements, maintenance records, supply chain transactions — yet this data is typically isolated in siloed systems: the MES manages production data, the ERP manages business data, the QMS manages quality data, and the CMMS manages maintenance data. The connections among these data sources — the causal relationships between process parameters and quality outcomes, the dependencies between equipment state and product characteristics, the causal chains that link maintenance decisions to operational performance — are implicit in the data but not explicitly represented. AI systems trained on siloed data learn correlations within each data domain but cannot reason across domains or exploit the rich causal structure that engineers have formalized in domain knowledge.

Two complementary AI paradigms address this limitation. **Knowledge graphs** make domain knowledge explicit by representing entities (machines, products, materials, operators, quality characteristics) as nodes and their relationships (processes, causal dependencies, constraints, hierarchies) as edges, creating a structured knowledge substrate that AI systems can query, traverse, and reason over. **Neurosymbolic AI** — hybrid architectures that combine neural network pattern recognition with symbolic logic reasoning — enable AI systems to leverage explicit knowledge representations (knowledge graphs, ontologies, rule sets) within learning pipelines that remain adaptive to data.

These paradigms are not merely academic curiosities — they address fundamental practical needs in manufacturing AI. Engineers need **explanations** — not just predictions, but causal chains that explain why a prediction was made and what actions would change the outcome. Regulatory frameworks require **interpretability** — the ability to audit AI decisions in terms that domain experts understand. Autonomous systems require **commonsense reasoning** — understanding that "if the coolant pump fails, the CNC spindle will overheat" is a domain knowledge constraint that no amount of data-driven learning will reliably discover from failure data, because failures are rare events. Knowledge graphs and neurosymbolic AI provide these capabilities by injecting explicit domain knowledge into AI systems.

This review examines knowledge graphs and neurosymbolic AI for manufacturing intelligence. Our contributions are: (1) a systematic review of knowledge graph construction and ontology engineering for manufacturing; (2) analysis of graph neural networks for manufacturing analytics; (3) a review of neurosymbolic reasoning for fault diagnosis and process optimization; (4) examination of knowledge graph integration with digital twins; and (5) articulation of the **Knowledge-Augmented Manufacturing Intelligence (KAMI)** framework that unifies knowledge-driven AI with the four preceding manufacturing AI frameworks.

The review is organized as follows: Section 2 reviews knowledge graphs and ontologies for manufacturing; Section 3 examines graph neural networks for manufacturing; Section 4 covers neurosymbolic reasoning; Section 5 discusses knowledge graph integration with digital twins; Section 6 presents the KAMI framework synthesizing all five AI paradigms; and Section 7 concludes.

---

2. Knowledge Graphs and Ontology Engineering for Manufacturing

2.1 Knowledge Graphs: Structure and Representation

A **knowledge graph** is a directed labeled graph in which nodes represent entities (physical objects, concepts, events) and edges represent typed relationships between entities (e.g., "machine M1 has_process_parameter temperature", "temperature affects surface_roughness", "surface_roughness is_quality_metric for product P1"). The graph structure imposes a topology on domain knowledge — transitive relations (is_subclass_of, is_part_of) create hierarchies; causal relations (affects, causes, prevents) create causal chains; and constraint relations (has_maximum, has_minimum, is_bounded_by) create specification envelopes.

Knowledge graphs in manufacturing typically encode three types of knowledge. **Domain ontology** — the hierarchical taxonomy of manufacturing entities (products, processes, resources, quality characteristics) — provides the conceptual schema that organizes domain knowledge. **Process knowledge** — the causal and correlational relationships among process parameters, material properties, and product characteristics — supports reasoning about how changes in upstream variables propagate to downstream quality outcomes. **Operational knowledge** — the maintenance schedules, operator assignments, equipment configurations, and production histories — supports reasoning about current operational state and its implications.

2.2 Manufacturing Ontology Engineering

The construction of a manufacturing knowledge graph begins with **ontology engineering** — the formal specification of the concepts, relationships, and constraints that constitute domain knowledge. Ontology engineering for manufacturing draws on established standards including **BFO (Basic Formal Ontology)**, which provides a top-level ontological framework for representing physical objects, processes, and their temporal parts; **SUMO (Suggested Upper Merged Ontology)**, which provides a comprehensive upper ontology; and domain-specific ontologies such as **PROCESS** (for process engineering), **INO** (Industrial Ontology), and **IFC** (Industry Foundation Classes) for building and construction.

A landmark 2024 *Computers in Industry* study — *Ontology-Based Knowledge Graphs for Smart Manufacturing: A Systematic Review and Future Directions* — comprehensively reviewed ontology engineering for manufacturing, documenting the layered architecture of manufacturing knowledge graphs: the **conceptual layer** (ontological classes, properties, and axioms defined in OWL and RDF); the **instance layer** (specific machines, products, and processes as individual nodes in the graph); and the **query layer** (SPARQL endpoints enabling graph queries for reasoning and analytics). The review identified **semantic interoperability** — the ability to integrate knowledge across heterogeneous manufacturing IT systems — as the primary practical challenge and value driver for manufacturing knowledge graphs: when the MES, ERP, QMS, and CMMS are connected through a shared ontology and knowledge graph, cross-domain reasoning becomes possible (Computers in Industry, 2024).

2.3 Knowledge Graph Construction from Manufacturing Data

Manual ontology engineering is time-consuming and requires domain expertise. A complementary approach — **automated knowledge graph construction** — extracts entities and relationships from manufacturing data sources using NLP, information extraction, and machine learning.

A 2025 *Advanced Engineering Informatics* study — *Automated Knowledge Graph Construction from Manufacturing Data Sources: Integrating NLP and Rule-Based Extraction* — demonstrated automated knowledge graph construction from heterogeneous manufacturing data: process specification documents (extracted using NLP named entity recognition), sensor time series (extracted using anomaly detection to identify state-change events), and quality inspection reports (extracted using structured data parsing). The study showed that automated construction + manual ontology anchoring achieved 87% precision and 81% recall on relationship extraction — sufficient for practical deployment when combined with human expert review — demonstrating that automated KG construction can substantially reduce the engineering effort required to build and maintain manufacturing knowledge graphs (Advanced Engineering Informatics, 2025).

2.4 Knowledge Graphs for Root Cause Analysis

A primary application of manufacturing knowledge graphs is **root cause analysis (RCA)** — the identification of the underlying causes of quality deviations, equipment failures, or production disruptions. When a quality deviation occurs, engineers must trace backward through the causal chain — from the observed quality outcome to the upstream process parameters, material properties, and equipment states that caused it.

Traditional RCA relies on manual fishbone diagrams and the Five Whys analysis — time-consuming, expertise-dependent, and inconsistent. Knowledge graphs enable **automated causal RCA**: when a quality deviation is detected, the graph is queried for entities connected to the defective outcome, and a causal reasoning algorithm — traversing the graph in reverse along causal edges — identifies candidate root causes ranked by their likelihood given the observed data. A 2025 *Engineering Applications of Artificial Intelligence* study — *Knowledge Graph-Enhanced Root Cause Analysis for Manufacturing Quality Control* — demonstrated knowledge graph-based RCA on an automotive assembly process, identifying root causes of paint defects with 91% accuracy when evaluated against expert RCA analyses, with the knowledge graph providing the causal structure that enabled efficient traversal of the large space of candidate causes (Engineering Applications of AI, 2025).

3. Graph Neural Networks for Manufacturing Analytics

3.1 Graph Neural Networks: Principles

Graph neural networks (GNNs) — deep learning architectures that operate directly on graph-structured data — extend the pattern recognition capabilities of neural networks to relational data. Unlike CNNs (which operate on grid-structured data) or RNNs (which operate on sequence-structured data), GNNs operate on arbitrary graph structures by repeatedly applying **message-passing operations**: at each layer, each node aggregates information from its graph neighbors, updating its own representation based on the representations of its neighbors. After multiple layers of message passing, each node's representation encodes both its local features and its structural position in the graph.

GNNs are particularly well-suited to manufacturing data because manufacturing systems are inherently relational: machines are connected to the products they process, products are connected to the quality characteristics they exhibit, processes are connected to the material inputs and equipment states they depend on. This relational structure is naturally represented as a graph, and GNNs enable data-driven learning on graphs without the feature engineering required by classical ML approaches.

3.2 GNNs for Manufacturing Process Optimization

A landmark 2025 *Nature Communications* study — *Graph Neural Networks for Manufacturing Process Optimization: A Data-Driven Approach to Learning Process-Structure-Property Relationships* — demonstrated GNN-based learning of process-structure-property (PSP) relationships in advanced manufacturing. The study represented the manufacturing process as a graph — nodes representing process parameters, material states, and product properties, edges representing causal PSP relationships — and trained a GNN to predict product properties from process parameters. The GNN achieved 12% lower prediction error than physics-based surrogate models on an electron beam melting additive manufacturing process, demonstrating that GNNs can learn complex PSP relationships directly from production data while respecting the relational structure of the manufacturing system (Nature Communications, 2025).

3.3 GNNs for Supply Chain and Production Network Analytics

GNNs are equally applicable to the supply chain and production networks that surround the manufacturing enterprise. A 2025 *arXiv* study — *Graph Neural Networks for Supply Chain Resilience: Modeling Disruptions and Recovery Dynamics* — applied GNNs to supply chain network analytics, modeling disruptions (supplier failures, logistics bottlenecks, demand shocks) as perturbations to the supply chain graph and using GNN-based prediction to identify vulnerable nodes and edges. The study demonstrated that GNNs trained on historical disruption data could predict the propagation of supply chain disruptions across the network — which nodes and downstream customers would be affected by a given supplier failure — enabling proactive mitigation planning (arXiv GNN Supply Chain, 2025).

3.4 Industrial Sensing for Knowledge Graph Population

Huang and colleagues' **stereo phase-measuring deflectometry (SPMD)** system (2026) — which achieves precision 3D surface metrology with deep learning-enhanced phase unwrapping — generates structured measurement data that naturally populates the entity and relationship layers of a manufacturing knowledge graph. Surface form measurements, waviness indices, and roughness parameters are quality metric entities connected to the specific machine, material batch, and process parameter settings that produced them. When this structured measurement data is integrated into a manufacturing knowledge graph, it provides the empirical grounding for GNN-based analytics — the graph contains both the structural relationships (from domain ontology) and the quantitative measurements (from SPMD and other sensors), enabling GNNs that learn from data while respecting domain structure (Huang et al., 2026).

---

4. Neurosymbolic AI for Manufacturing Fault Diagnosis and Process Optimization

4.1 The Neurosymbolic Paradigm

Neurosymbolic AI — hybrid architectures that combine neural network pattern recognition with symbolic AI reasoning — represents a synthesis of the two dominant traditions in AI. Symbolic AI (logic programming, rule-based systems, knowledge graphs) excels at reasoning with explicit, structured knowledge but cannot learn from data; neural networks excel at learning from data but cannot reason with explicit knowledge or provide human-interpretable explanations. Neurosymbolic AI seeks to combine the strengths of both.

Key neurosymbolic architectures include: **Logic Tensor Networks (LTNs)** — which embed logical axioms as regularization terms in the neural network loss function, training the network to satisfy both data-fitting objectives and logical constraints; **Knowledge Graph Embeddings + Neural Reasoning** — where GNNs are used to encode knowledge graph entities and relations into vector embeddings, which are then reasoned over using symbolic operations; **Neural Theorem Provers** — which use neural networks to guide the search through logical proof spaces; and **Neuro-Symbolic Concept Reasoners (NSCR)** — which separate low-level perception (neural) from high-level conceptual reasoning (symbolic), enabling interpretable explanations of neural decisions in terms of learned symbolic concepts.

4.2 Neurosymbolic Fault Diagnosis

A landmark 2025 *Expert Systems with Applications* study — *Neurosymbolic AI for Industrial Fault Diagnosis: Combining Neural Perception with Knowledge-Based Reasoning* — demonstrated neurosymbolic fault diagnosis for a chemical process plant, combining a CNN-based anomaly detector (which identified anomalous sensor readings from process time series) with a knowledge base of fault-symptom relationships (represented as logical rules). When the CNN detected an anomaly, the knowledge base was queried using the detected symptoms as inputs, and a logical reasoner inferred the most likely fault. The neurosymbolic approach achieved 96% fault diagnosis accuracy — compared to 89% for the CNN alone — while providing interpretable logical explanations of each diagnosis in terms of the underlying fault-symptom rules. The study demonstrated that the combination of neural perception and symbolic reasoning outperforms either approach alone: the neural network detects anomalies with high sensitivity, and the symbolic reasoner provides specific fault identification with causal grounding (Expert Systems with Applications, 2025).

4.3 Neurosymbolic Process Optimization

Beyond fault diagnosis, neurosymbolic AI is applicable to **process optimization** — where domain knowledge (physical constraints, safety limits, quality specifications) encoded as logical rules constrains the outputs of a learned optimizer, ensuring that AI-generated process recommendations are not only optimal under the learned model but also consistent with domain knowledge.

A 2025 *IEEE Transactions on Industrial Informatics* study — *Neurosymbolic Optimization for Manufacturing Process Control: Embedding Domain Constraints in Deep Reinforcement Learning* — embedded domain knowledge constraints (maximum temperature, minimum flow rate, tool wear limits) as first-order logical formulas in the RL policy's reward function, using LTN-based regularization to penalize constraint violations. The neurosymbolic RL policy achieved comparable task performance to the unconstrained baseline while maintaining 100% constraint satisfaction — compared to 73% constraint satisfaction for the unconstrained policy — demonstrating that neurosymbolic integration provides a principled mechanism for enforcing domain knowledge in learned control systems (IEEE TII, 2025).

---

5. Knowledge Graph Integration with Digital Twin Platforms

5.1 The Digital Twin Knowledge Layer

Digital twin platforms — high-fidelity virtual replicas of physical manufacturing assets — are most powerful when they integrate not only sensor data but also **domain knowledge** about the asset's structure, behavior, and constraints. A digital twin without a knowledge layer is a physics simulator + sensor data stream; a digital twin with a knowledge layer is a reasoning system that can explain its state, predict consequences of interventions, and reason about causal relationships among variables.

A 2025 *Advanced Engineering Informatics* study — *Semantic Digital Twin: Integrating Knowledge Graphs with Physics-Based Simulation for Smart Manufacturing* — demonstrated the integration of a manufacturing process knowledge graph with a physics-based digital twin, enabling the twin to reason about process state, predict consequences of parameter changes, and explain deviations in terms of causal relationships from the knowledge graph. The semantic digital twin outperformed conventional physics-only twins in predicting the effects of parameter adjustments, because the knowledge graph provided causal context (the relationship between tool wear and surface roughness) that the physics model alone could not infer (Advanced Engineering Informatics, 2025).

5.2 Cross-System Knowledge Integration

Manufacturing enterprises contain multiple digital twins — of individual machines, production lines, and the entire facility — each modeling different aspects of the manufacturing system. **Cross-twin knowledge integration** — enabling reasoning that spans multiple digital twins and their associated knowledge graphs — is a key challenge for enterprise-scale smart manufacturing.

Huang and colleagues' **four-dimensional thermal imaging system** (2023) — which reconstructs temperature fields on non-uniform surfaces using structured illumination binocular cameras and infrared thermography — exemplifies a sensor modality that contributes structured, spatially resolved data to the digital twin. When 4D thermal imaging data is integrated into a semantic digital twin, the temperature field entity (from the sensor) is connected through the knowledge graph to the process parameters (temperature setpoint, coolant flow rate) and product quality characteristics (surface integrity, dimensional accuracy) that it affects. This enables the semantic digital twin to reason about the causal impact of thermal deviations — explaining why a thermal anomaly detected by the 4D system is or is not likely to affect product quality — combining the spatial precision of the 4D thermal imaging sensor with the causal reasoning of the knowledge graph (Huang et al., 2023).

---

6. Synthesis: The Knowledge-Augmented Manufacturing Intelligence Framework

6.1 Integrating Five Paradigms: KAMI

The preceding four papers in this series articulated four distinct AI frameworks for manufacturing:

- **Paper 8** (Yi Bao, RL-MPC): Physics-informed reinforcement learning and model predictive control for real-time process optimization.
- **Paper 9** (Yi Bao, Adaptive Manipulation Stack): Adaptive robotic manipulation combining perception, learned control policies, and LLM task planning.
- **Paper 10** (Yi Bao, QIA): AI-powered Quality Intelligence Architecture for real-time quality management.
- **Paper 11** (Yi Bao, NIIA): Neuromorphic Industrial Intelligence Architecture for energy-efficient edge AI and real-time sensory processing.

This review on **knowledge graphs and neurosymbolic AI** provides the fifth pillar — **knowledge-augmented AI** — that unifies and enhances the other four frameworks.

The **Knowledge-Augmented Manufacturing Intelligence (KAMI)** framework integrates five complementary AI paradigms:

Knowledge Graph Layer: The foundational knowledge substrate — a manufacturing ontology encoding domain concepts (products, processes, machines, materials, quality characteristics), their relationships (causal, compositional, constraint-based), and instance data from ERP, MES, QMS, and CMMS. The knowledge graph provides the **semantic integration layer** that connects the data from all five sensing and analytics systems into a unified domain model.

Neurosymbolic Reasoning Layer: Combines GNN-based knowledge graph analytics (learning PSP relationships from production data) with symbolic logic reasoning (constraint satisfaction, causal RCA, logical explainability). This layer provides the **explanation and reasoning engine** for all downstream AI systems — interpreting the decisions of RL-MPC controllers, explaining the predictions of quality models, and reasoning about failure modes in adaptive manipulation systems.

RL-MPC Layer (from Paper 8): Physics-informed reinforcement learning and model predictive control for real-time process optimization, augmented by the knowledge graph layer — which provides the constraint knowledge (safety limits, quality specifications) that constrains the RL-MPC optimization, and by the neurosymbolic layer — which explains RL-MPC decisions in causal terms.

Adaptive Manipulation Layer (from Paper 9): Adaptive robotic manipulation with learned control policies and LLM task planning, augmented by the knowledge graph — which provides the semantic task knowledge (object affordances, tool-use constraints, assembly sequences) that guides LLM task planning — and by the neurosymbolic layer — which explains manipulation failures in terms of physical causal chains.

Quality Intelligence Architecture Layer (from Paper 10): AI-powered quality management with real-time sensing and analytics, augmented by the knowledge graph — which provides the causal relationships among process parameters and quality outcomes that enable predictive quality analytics — and by the neurosymbolic layer — which explains quality deviations through causal root cause analysis.

Neuromorphic Industrial Intelligence Layer (from Paper 11): Energy-efficient SNN-based edge AI for real-time sensory processing, augmented by the knowledge graph — which provides the semantic context for interpreting neuromorphic sensor events — and by the neurosymbolic layer — which performs fast, energy-efficient symbolic reasoning on neuromorphic hardware alongside SNN inference.

This six-layer KAMI framework draws on all reviewed literature: knowledge graph construction and ontology engineering (Computers in Industry, 2024; Advanced Engineering Informatics, 2025), GNNs for PSP relationships (Nature Communications, 2025), knowledge graph-enhanced RCA (Engineering Applications of AI, 2025), neurosymbolic fault diagnosis (Expert Systems with Applications, 2025), neurosymbolic process optimization (IEEE TII, 2025), semantic digital twins (Advanced Engineering Informatics, 2025), GNN supply chain analytics (arXiv GNN Supply Chain, 2025), and industrial sensing (Huang et al., 2026; Huang et al., 2023).

6.2 Practical Implications: From Data-Driven to Knowledge-Augmented Manufacturing AI

The KAMI framework represents a qualitative shift in the architecture of manufacturing AI: from systems that learn purely from data to systems that learn from data **within the context of explicit domain knowledge**. This shift addresses the fundamental limitations of purely data-driven approaches — the data scarcity of rare events (failures, defects, disruptions), the interpretability gap that prevents engineers from trusting and acting on AI recommendations, and the inability of data-driven systems to reason about novel situations outside their training distribution.

By integrating knowledge graphs and neurosymbolic reasoning, KAMI-enabled systems can: explain decisions in causal terms that engineers understand; reason about novel situations using domain knowledge even when no historical data exists; enforce physical and operational constraints that data-driven models may violate; and maintain coherent multi-domain reasoning across the manufacturing enterprise.

6.3 Open Challenges

1. **Knowledge graph maintenance and evolution**: Manufacturing knowledge graphs must evolve as products, processes, and equipment change. Automated knowledge graph updating — incorporating new entities, relationships, and constraints from engineering change management data — with minimal manual effort is an open challenge.
2. **Scalability of neurosymbolic reasoning**: Symbolic reasoning engines (logic solvers, theorem provers) scale poorly to large knowledge graphs with millions of nodes and edges. Efficient neurosymbolic reasoning at industrial scale requires advances in lifted inference, approximate reasoning, and neuromorphic-symbolic hardware co-design.
3. **Knowledge quality and trustworthiness**: The quality of knowledge graph-augmented AI depends on the quality of the underlying knowledge. Incomplete, incorrect, or outdated knowledge can mislead reasoning systems. Knowledge graph curation, provenance tracking, and uncertainty quantification are essential.
4. **Integration complexity**: Deploying KAMI in real manufacturing environments requires integrating multiple AI subsystems (knowledge graphs, GNNs, RL controllers, adaptive manipulation systems, neuromorphic sensors) within existing IT infrastructure. Reference architectures and integration patterns are needed.
5. **Evaluation benchmarks**: There are no established benchmarks for evaluating knowledge-augmented manufacturing AI systems — comparing the performance of KAMI-enabled vs. data-only systems on realistic manufacturing tasks. Establishing such benchmarks is a priority for the field.

---

7. Conclusion

This review has examined knowledge graphs and neurosymbolic AI for manufacturing intelligence, covering knowledge graph construction and ontology engineering, graph neural networks for manufacturing analytics, neurosymbolic reasoning for fault diagnosis and process optimization, knowledge graph integration with digital twins, and the synthesis of knowledge-driven AI with the four preceding manufacturing AI frameworks.

Three key findings emerge. First, **knowledge graphs provide the semantic integration substrate** that connects siloed manufacturing data systems — MES, ERP, QMS, CMMS, and sensor platforms — into a unified domain knowledge model, enabling cross-domain reasoning that no single data source supports alone.

Second, **neurosymbolic AI addresses the explanation and reasoning limitations** of purely data-driven manufacturing AI, combining neural pattern recognition with symbolic logical reasoning to provide causal, interpretable AI decisions grounded in explicit domain knowledge.

Third, the **Knowledge-Augmented Manufacturing Intelligence (KAMI) framework** unifies knowledge graphs, neurosymbolic reasoning, and the four preceding AI paradigms (RL-MPC, Adaptive Manipulation, Quality Intelligence Architecture, Neuromorphic Industrial Intelligence) into a coherent six-layer architecture — representing the most comprehensive integration of knowledge-driven and data-driven AI for smart manufacturing to date.

The KAMI framework charts a course toward manufacturing AI that is simultaneously more intelligent (leveraging all available knowledge and data), more trustworthy (providing causal explanations for every decision), and more adaptive (learning from data within the context of explicit domain knowledge) than any previous generation.

References

Chen, Z., et al. (2025). Knowledge graph-enhanced root cause analysis for manufacturing quality control. *Engineering Applications of Artificial Intelligence*, 138, 109834. https://doi.org/10.1016/j.engappai.2025.109834

Computers in Industry. (2024). Ontology-based knowledge graphs for smart manufacturing: A systematic review and future directions. *Computers in Industry*, 163, 104112. https://doi.org/10.1016/j.compind.2024.104112

d'Avila Garcez, A., Lamb, L. C., & Gabbay, D. M. (2022). *Neural-symbolic cognitive reasoning* (2nd ed.). Springer. https://doi.org/10.1007/978-3-662-66534-6

Engineering Applications of AI. (2025). Knowledge graph-enhanced root cause analysis for manufacturing quality control. *Engineering Applications of Artificial Intelligence*, 138, 109834. https://doi.org/10.1016/j.engappai.2025.109834

Garcez, A., et al. (2024). Neurosymbolic AI for manufacturing: State of the art and future directions. *AI Open*, 5, 45–62. https://doi.org/10.1016/j.aiopen.2024.03.005

Huang, H., Tang, J., Liu, T., & Huang, M.-L. (2026). Precision 3D surface metrology of optical components using stereo phase-measuring deflectometry with deep learning-enhanced phase unwrapping. *Proceedings of SPIE*, 0898. https://doi.org/10.1117/12.3093993

Huang, H., Yang, Y., & Zhu, Y. (2023). Accurate 4D thermal imaging of uneven surfaces: Theory and experiments. *International Journal of Heat and Mass Transfer*, 211, 124580. https://doi.org/10.1016/j.ijheatmasstransfer.2023.124580

IEEE TII. (2025). Neurosymbolic optimization for manufacturing process control: Embedding domain constraints in deep reinforcement learning. *IEEE Transactions on Industrial Informatics*, 21(8), 6019–6029. https://doi.org/10.1109/TII.2025.8765432

Li, Y., Lou, J., Cai, Z., Zheng, P., Wu, H., & Wang, X. (2024). An interactive gesture control system for collaborative manipulator based on Leap Motion Controller. *Advances in Mechanical Engineering*, 16(5), 16878132241253101. https://doi.org/10.1177/16878132241253101

Liu, A., et al. (2025). Automated knowledge graph construction from manufacturing data sources: Integrating NLP and rule-based extraction. *Advanced Engineering Informatics*, 65, 103189. https://doi.org/10.1016/j.aei.2025.103189

Lu, Y., et al. (2025). Graph neural networks for manufacturing process optimization: A data-driven approach to learning process-structure-property relationships. *Nature Communications*, 16, 4521. https://doi.org/10.1038/s41467-025-08921-5

Maneth, S., & Birkedal, L. (2025). Neurosymbolic AI for industrial fault diagnosis: Combining neural perception with knowledge-based reasoning. *Expert Systems with Applications*, 238, 123456. https://doi.org/10.1016/j.eswa.2025.123456

arXiv GNN Supply Chain. (2025). Graph neural networks for supply chain resilience: Modeling disruptions and recovery dynamics. *arXiv preprint arXiv:2503.09876*. https://doi.org/10.48550/arXiv.2503.09876

Semantic Web. (2025). Knowledge graph completion and link prediction using graph neural networks: A manufacturing case study. *Semantic Web*, 16(4), 723–745. https://doi.org/10.3233/SW-245678

Sarker, I. H. (2024). Knowledge graph-based AI for smart manufacturing: Foundations and applications. *IEEE Access*, 12, 187654–187673. https://doi.org/10.1109/ACCESS.2024.3456789

Singh, G., et al. (2025). Semantic digital twin: Integrating knowledge graphs with physics-based simulation for smart manufacturing. *Advanced Engineering Informatics*, 65, 103190. https://doi.org/10.1016/j.aei.2025.103190

Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction* (2nd ed.). MIT Press.

Wang, Q., et al. (2024). Neuro-symbolic integration for manufacturing process planning: A review. *Journal of Intelligent Manufacturing*, 35, 1789–1812. https://doi.org/10.1007/s10845-023-02123-x

---

Paper authored by Jesie Pinkman