

Causal Knowledge Graphs for Intelligent Manufacturing: A Review Based on the Causal Hierarchy

Yiyun Fei and Rojer J. Jiao

School of Mechanical Engineering

Georgia Institute of Technology, Atlanta, GA, USA

yfei38@gatech.edu

Abstract

Manufacturing systems today are heavily instrumented and connected, producing large amounts of diverse operational data and encouraging the use of AI for prediction and decision support. However, models based mainly on correlations often fall short in robustness, generalization across settings, and clear mechanistic understanding, which are essential in high-value and safety-critical production environments. This review examines the integration of causal reasoning and knowledge graph infrastructures as a structural foundation for explainable and decision-oriented manufacturing intelligence. Pearl's three-level causal hierarchy (association, intervention, and counterfactual) is adopted to classify recent research articles according to the depth of causal capability they operationalize. The analysis shows that most existing work remains concentrated at structured associational reasoning and interventional effect estimation, while fully operational counterfactual systems are comparatively rare. Knowledge graphs are widely employed to provide structured semantic representation and traceability, supporting human-centered and machine-interpretable reasoning. However, formal interventional and counterfactual semantics remain partially integrated. Recurring consensus and open challenges are identified, including hidden confounding, partial observability, scalability of expert-driven modeling, underdeveloped dynamic causality, and ambiguity in the use of "counterfactual" terminology. The review clarifies the structural role of causal knowledge graphs in advancing manufacturing intelligence from traceable association toward quantified intervention and counterfactual foresight and outlines future research directions.

Keywords

Causal Reasoning, Knowledge Graph, Intelligent Manufacturing, Decision Support, Explainability AI

1. Introduction

Manufacturing systems are increasingly instrumented and interconnected, generating massive volumes of sensor and operational data. This data-rich environment creates unprecedented opportunities to deepen understanding of manufacturing processes and to enhance human decision support. As a result, data-driven methods and artificial intelligence (AI) have rapidly advanced and achieved widespread adoption across industrial applications. However, despite strong predictive performance, many deployed AI systems continue to face limitations in robustness, transferability across contexts, and mechanistic interpretability. In critical and high-value production environments, decision-making requires more than accurate prediction; it requires transparency, traceability, and structural understanding. This need has driven the rapid development of explainable AI (XAI), aimed at mitigating the black-box nature of complex models and improving trust in AI-assisted decision-making. In parallel, knowledge graph (KG) modeling has emerged as a structured knowledge infrastructure for organizing domain semantics, integrating heterogeneous information sources, and enabling contextual reasoning. Knowledge graphs are widely adopted in manufacturing for representing process knowledge, defect relationships, and system dependencies. Increasingly, these two streams, AI and KG modeling, are being integrated. Knowledge graphs can inject structured domain knowledge into machine learning systems, thereby enhancing contextual awareness and interpretability.

Causal reasoning further extends this trajectory. Although causality is sometimes discussed in the XAI literature as a form of causal or counterfactual explanation (Madathil et al. 2025), it fundamentally refers to modeling structural mechanisms that support intervention and counterfactual reasoning. Because causal models are inherently graph-based representations, the integration of causal reasoning with knowledge graph infrastructures, also referred to as causal knowledge graphs, represents a promising direction for advancing explainable and decision-oriented manufacturing intelligence.

Existing research spans causal reasoning, knowledge graphs, XAI, and large models. However, there is no unified survey that jointly assesses the level of causal capability in manufacturing systems and clarifies the role of knowledge graph infrastructures in supporting intervention and counterfactual reasoning. This paper addresses this gap by examining the integration of causal reasoning and knowledge graph infrastructures for explainable and intelligent decision-making in manufacturing systems. The main contributions are:

- A systematic causal-level analysis of intelligent manufacturing literature based on Pearl's hierarchy, classifying recent works according to the depth of causal reasoning implemented.
- A synthesis of recurring themes and open challenges across causal levels.

2. Background and Survey Landscape

2.1 Core Concepts

2.1.1 Causal Hierarchy

The causal hierarchy, introduced by Judea Pearl in *Causality* (Pearl 2009), distinguishes three fundamentally different types of reasoning: prediction, intervention, and counterfactual reasoning. These levels are not merely alternative techniques; they correspond to qualitatively different kinds of questions. The hierarchy reflects increasing structural and mechanistic assumptions about the system under study. Moving from prediction to intervention and then to counterfactual reasoning requires progressively stronger knowledge about how variables are generated and how they interact.

L1: The first level is prediction/association (or “seeing”). It addresses questions based purely on passive observation, such as: “What is the expected demand when the price is observed to be p_0 ?” At this level, reasoning relies on the joint probability distribution of the variables. Conditional probabilities and correlations are sufficient, and no assumptions about causal direction or structural relationships are required. Consequently, this level cannot distinguish genuine causal influence from spurious association. It provides answers about what is expected to be observed, not about what would occur under deliberate change.

L2: The second level is intervention (or “doing”). The focus shifts from observation to manipulation: “What would the demand be if the price were set to p_0 ?” This type of reasoning is formalized using the do-operator, which represents an external intervention that replaces the original mechanism determining a variable. Answering such questions requires an explicit causal structure, commonly represented by a directed acyclic graph (DAG) or a structural causal model (SCM). The joint distribution alone is insufficient. The structural model encodes directional relationships among variables, enabling identification of causal effects and computation of the consequences of deliberate actions.

L3: The third level is counterfactual reasoning (or “imagining”). These questions concern alternative scenarios about specific instances in which certain facts are already known. For example: “Given that the price was p_0 and demand was observed, what would demand have been if the price had instead been set to p_1 ?” Counterfactual reasoning requires a structural causal model with explicit functional relationships and, often, assumptions about latent background factors. The reasoning process typically involves three steps: abduction (updating beliefs about unobserved factors based on observed evidence), action (modifying the model to reflect the hypothetical intervention), and prediction (computing the outcome under the modified model). By combining observed facts with structural assumptions about underlying mechanisms, counterfactual reasoning represents the most expressive and most demanding level in the hierarchy.

2.1.2 Knowledge Graph

A knowledge graph is defined as a directed, labeled graph in which edges are assigned explicit and semantically well-defined meanings. It provides a principled abstraction for representing structured knowledge and integrating information derived from heterogeneous data sources. By encoding entities and their relationships within a unified

graph structure, knowledge graphs enable the explicit representation of meaning through topology and labeling, rather than through implicit schema constraints.

The structural foundation of a knowledge graph consists of nodes, edges, labels, and triples. Nodes denote entities or concepts, which may include physical objects, organizations, processes, or abstract notions. Edges connect pairs of nodes to represent specific relationships between them. Each edge is associated with a label that defines the semantic interpretation of the relationship. Formally, a knowledge graph can be characterized as a subset of the Cartesian product $N \times L \times N$, where N represents the set of nodes and L the set of relation labels. Each element of this set is referred to as a triple, conventionally expressed as subject–predicate–object, and constitutes the fundamental unit of representation. Figure 1 presents an example data model of a knowledge graph designed for storing defect cases.




Figure 1. Illustrative Example: Data Model of the Knowledge Graph for Storing Defect Cases

2.2 Existing Survey Landscape

This section reviews existing survey papers at the intersection of causality, knowledge graphs (KG), artificial intelligence (AI), and explainable AI (XAI) in manufacturing, as summarized in Table 1. The goal is to understand how these dimensions have been covered in prior work and to identify structural gaps in the literature.

From the perspective of causality, Vuković and Thalmann (2022) conducts a structured literature review of data-driven causal discovery in manufacturing. The authors organize prior work into four application areas: fault detection/management, root cause analysis, causality as a facilitator, and domain/conceptual work. They also summarize commonly used methods such as Granger causality, transfer entropy, Bayesian networks, and SCM-related approaches. The review concludes that manufacturing applications remain scattered and assumption-sensitive, motivating a research agenda focused on integrating causality into continuous process improvement, improving user guidance, and developing methods that better handle non-ideal industrial data. Complementarily, Li et al. (2025) presents a comprehensive survey on causal inference in industrial fault diagnosis, systematically reviewing the evolution from physics-driven and data-driven approaches to causal-driven methods. The paper organizes the field around causal structure discovery and causal effect estimation, summarizing techniques such as Bayesian networks, structural causal models, transfer entropy, time series-based causal modeling, and deep causal networks. The survey highlights challenges such as confounding bias, nonlinear modeling, and data complexity.

From the perspective of knowledge graphs, Xiao et al. (2023) reviews knowledge graph-based manufacturing process planning, positioning process knowledge graphs (PKGs) as enablers for intelligent computer-aided process planning. The paper systematically analyzes the technical pipeline, including process knowledge representation (ontology-, rule-, and graph-based models), knowledge extraction from CAD/STEP files and unstructured process documents, schema/data layer construction, graph refinement, and KG-driven process generation. It further discusses the integration with LLMs for automated knowledge construction and reasoning. It identifies challenges such as automated construction, knowledge completion, and intelligent Q&A-based process generation as future research directions. Similarly, Wan et al. (2024) provides a broader review of KGs in smart manufacturing, covering knowledge extraction, KG construction, embedding, fusion, and reasoning across lifecycle-wide applications, such as engineering design, scheduling, quality control, supply chain management, and predictive maintenance. Emerging trends such as LLM–KG integration, multimodal KGs, digital twin coupling, and KG-enhanced decision support are identified. However, formal causal modeling remains on the margin.

From the perspective of explainable AI, Moosavi et al. (2024) surveys XAI in manufacturing and industrial cyber-physical systems, classifying techniques and their applications across use cases such as predictive maintenance, fault detection and diagnosis, process control, product development, inventory management, quality assurance, and cybersecurity. While highlighting the importance of transparency, trustworthiness, regulatory compliance, and human oversight in industrial AI systems, the survey mainly focuses on post hoc interpretability technique. Similarly, Madathil et al. (2025) provides a systematic review of XAI in smart manufacturing, examining both intrinsic and post-hoc interpretability techniques across a wide range of industrial applications. The paper classifies XAI methods by modeling stage, explanation type (feature relevance, rule-based, visual, causal, counterfactual), and manufacturing task (predictive maintenance, defect detection, process optimization, digital twins, supply chain, cybersecurity). It highlights persistent challenges including the explainability–performance trade-off, computational latency in real-time systems, lack of standardized evaluation metrics, scalability across production levels, and regulatory and ethical barriers. Both surveys primarily treat explainability at the feature-attribution or model-transparency level rather than through formal causal reasoning or structured knowledge representation.

Table 1. Related Survey Coverage Matrix

Paper	Year	Causality	KG	AI	XAI	Large Model	Focus
Vuković and Thalmann	2022	✓		△	△		Causal discovery in manufacturing
Li et al.	2025	✓		✓		△	Causal inference in industrial fault diagnosis
Xiao et al.	2023		✓	△		△	Knowledge graph-based manufacturing process planning
Wan et al.	2024	△	✓	✓	△	△	Knowledge graph for smart manufacturing
Moosavi et al.	2024			✓	✓		Explainable AI in manufacturing
Madathil et al.	2025	△		✓	✓		Explainable AI in manufacturing
Ours	2026	✓	✓	△	△	△	Causal knowledge graph for intelligent manufacturing

✓ = Addressed as a core focus of the survey

△ = Mentioned or partially integrated

In summary, Table 1 reveals a clear structural fragmentation in the existing survey landscape. Causality-focused surveys concentrate on discovery and inference without integrating knowledge graph infrastructures. KG-focused surveys emphasize semantic representation and reasoning but rarely formalize interventional or counterfactual semantics. XAI surveys prioritize interpretability techniques but largely remain at the associational or post-hoc explanation level. Large models are discussed only partially and typically in isolation from causal structure. This fragmentation motivates the need for a unifying perspective. In contrast to prior surveys, this paper focuses on the integration of causal reasoning and knowledge graph infrastructures to support explainability and intelligent decision-making in manufacturing systems.

2.3 Review Landscape

Excluding the six survey papers reviewed in Section 2.2, this study analyzes 47 primary research articles. These papers were retrieved from Google Scholar using combinations of the keywords “*causal reasoning*,” “*knowledge graph*,” and “*manufacturing*,” along with closely related synonyms. The search strategy emphasized intersections among causality, structured knowledge representation, and intelligent manufacturing applications. The publication window primarily focuses on 2024–2025 to reflect the most recent developments in the field. A limited number of works from 2022, 2023, and early 2026 were also included to broaden contextual coverage and ensure conceptual completeness. The corpus consists predominantly of journal articles, supplemented by selected conference papers and a small number of preprints.

3. Causal Level Analysis of Selected Literature

This chapter analyzes the selected literature through the lens of the 3-tier causal hierarchy (associational, interventional, and counterfactual). Rather than organizing papers by application domain or algorithmic family, this chapter classifies them according to the depth of causal reasoning they operationalize. The objective is to clarify what type of causal capability each work truly enables—whether structured correlation modeling (L1), quantified intervention analysis (L2), or twin-world counterfactual simulation (L3)—and to reveal the progression of manufacturing intelligence from traceable association to actionable intervention and finally to proactive causal foresight. The proposed analytical framework is structured according to the 3-tier causal hierarchy, as illustrated in Figure 2.

3.1 L1 Associational Manufacturing Intelligence

This subsection summarizes L1 associational manufacturing intelligence, where “causality” is encoded as structured relations (KGs/graphs/event links) and used for retrieval, link prediction, path reasoning, and explanation, without identifying causal effects under causal interventions. Across the 24 papers classified under this category, the common pattern is: make causal-like knowledge computable and produce traceable reasoning for engineers.

3.1.1 Graph Construction from Industrial Knowledge

The first category of research focuses on building the graph substrate from heterogeneous industrial sources. Works here transform documents, logs, and multimodal streams into machine-readable structures: temporal industrial KGs that bind time-series events to entities (Zhou et al. 2023), industrial temporal scene graphs from videos using VLM-based multi-agent prompting for annotation-free spatial-temporal relation extraction (Wang et al. 2025), causal information extraction from FMEA/slides (Razouk et al. 2025), assembly/process KGs from unstructured documentation (Li et al. 2025), two-layer causal networks from quality problem-solving records (Wang et al. 2025), supplier capability KGs from web pages (Li et al. 2025), and system-level frameworks/meta-triples that support dynamic instantiation and governance (Cai et al. 2023; Li et al. 2025). The shared contribution is turning dispersed operational knowledge into an explicit relational form that downstream reasoning can traverse and audit.




Figure 2. The Hierarchy of Causal Level Analysis.

3.1.2 Graph Reasoning & Explainability

The second category of research uses graphs primarily for associational reasoning and interpretability. Typical mechanisms include event-overlap weighting and graph-based ranking for bottleneck root cause analysis (Cen et al. 2026), graph-structural tracing in cognitive digital twins for maintenance workflows (Lim et al. 2024), causality-aware feature dependency modeling in image-based manufacturing tasks (Terziyan and Vitko 2023; Streltsov et al. 2025), and causal-graph-informed attribution propagation such as Shapley-flow explanations in battery production (Kim et al. 2025). In addition, several works adopt explanation-oriented “counterfactual” techniques that generate minimal feature changes to flip model predictions in plant operation or equipment diagnostics (Bhakte and Srinivasan 2024; Narayanan et al. 2024), while governance-oriented discussions emphasize embedding XAI into monitoring systems for trustworthy deployment (Johannssen et al. 2025). These approaches enhance transparency and diagnostic speed by providing interpretable paths, rankings, and feature-level narratives, yet remain observational because they do not estimate intervention effects.

3.1.3 Graph-enabled Decision Support with LM Systems

The third stream operationalizes L1 graphs into decision-support pipelines, frequently integrating LLM or VLM interfaces. Examples include KG-driven machining process recommendation systems (Su et al. 2025), knowledge-graph-enhanced LLM pipelines for CNC fault diagnosis (Wu et al. 2025) and Industry 5.0 maintenance reasoning (Ma et al. 2025), and KG-LLM hybrid bottleneck analysis frameworks in ship manufacturing (Ma et al. 2025). Other works use LLMs to ground process-planning knowledge retrieval (Hoang et al. 2025), orchestrate human-in-the-loop time-series fault diagnosis (Zhang et al. 2025), or support adaptive self-labeling and redeployment in cyber-manufacturing IoT systems (Ren et al. 2025). Graph-based recommendation mechanisms are also applied to supplier discovery (Li et al. 2025), process knowledge personalization (Wu et al. 2026), and constraint-based process planning with interactive refinement (Huang et al. 2025). Even when directly influencing operational choices, these systems rely on structured retrieval, path reasoning, ranking, and constraint checking over stored relations rather than identifying causal effects under explicit interventions; therefore, they remain within the L1 associational layer.

3.2 L2 Interventional Manufacturing Intelligence

Unlike L1, which represents and reasons over structured dependencies, L2 explicitly estimates or evaluates the effects of interventions. Interventions may occur at the physical variable level, the cross-process system level, or the operational decision layer. Across the literature, L2 approaches typically rely on SCMs, Bayesian networks, causal discovery under intervention assumptions, backdoor adjustment, domain shift modeling, or equivalent intervention analysis, often accompanied by refutation or robustness testing.

3.2.1 Physical Process/Mechanism Level

At this level, interventions occur on physical variables (e.g. process parameters, material properties, sensor data), and effect sizes are explicitly estimated or evaluated under confounding control. The objective is primarily mechanistic understanding, defect reduction, and process optimization through quantified intervention effects. Several works implement formal causal inference pipelines to estimate unbiased physical effects. Backdoor adjustment and refutation testing are used to validate identified relationships in machining contexts (Wang et al. 2025; Jeong et al. 2025). A two-stage approach constructs causal networks for physical machining processes and estimates effect sizes after controlling confounders (Wang et al. 2025). Electrical variables such as current and voltage are modeled causally to improve sustainability-oriented decision-making (Jeong et al. 2025). One of the most comprehensive pipelines is presented in (Fekete et al. 2025), integrating expert knowledge, causal discovery, SCM construction, identification, effect estimation, and robustness refutation to support energy-aware process optimization. The framework quantifies intervention effects explicitly (e.g., a 1 mm increase in outer diameter increases energy consumption by 0.42 kWh), demonstrating fully operational interventional inference. Other works combine digital-physical integration with structural causal modeling. Xia et al. (2025) models machining processes via SCM integrated with GNN over digital twins, enabling effect reasoning across physical and computational layers. Sandt and Martin (2025) identifies unbiased causal effects in foundry process variables under confounding control. Zhou et al. (2024) bridges physical-level causal inference and language reasoning by building a Causal Quality Knowledge Graph (CQKG) via interventional SCG-SPN modeling, estimating causal effects under confounder control, and integrating the resulting structured effects into an LLM (CausalKGPT) for aerospace defect analysis. Causal representation learning also contributes at this level. While not always constructing explicit causal graphs, disentangled latent models uncover data-generating mechanisms and mitigate bias beyond correlation-based explanations (Becker et al. 2025), emphasizing that manufacturing explainability must reveal hidden mechanisms rather than merely interpret predictions. From the vision inspection perspective, interventional strategies are used to remove background-induced confounding in defect detection tasks (Liang et al. 2025; Li et al. 2025). These approaches explicitly intervene on spurious visual correlations and integrate causal adjustment into neural networks to improve robustness and explainability under complex backgrounds.

3.2.2 System/Production Line Level

This level models cross-process propagation and upstream-downstream mechanisms across multiple units, machines, or workflow stages. Here, causal graphs represent structural dependencies between subsystems rather than single process parameters. Wehner et al. (2023) creates an interactive RCA system across manufacturing processes using a causal Bayesian network learned from observational sensor data through constraint-based and score-based structure learning under SEM-based assumptions. Lippi et al. (2023) constructed a hierarchical digital twin architecture for smart factory and use it as an active reasoning agent for causal model learning, such as performing intervention to evaluate the effect of upstream machine faults to downstream machines. Multi-domain causal discovery frameworks Yu et al. (2025) and Yu et al. (2025) was proposed to separate domain-invariant and domain-specific causal

mechanisms in manufacturing systems. By leveraging domain shifts as implicit interventions, the method first identifies stable shared causal structure and then refines individual domain mechanisms. This approach improves causal structure recovery compared to pooled-only or single-domain discovery strategies. Ji et al. (2025) proposes an interventional causal reasoning framework that combines equivalent intervention analysis, nonlinear time-lagged causal discovery, and causal-graph pruning to reconstruct fault propagation plant-wide systems. Conceptual and hybrid frameworks also emerge at this level. Youssef (2024) proposes an integrated architecture combining neural networks with causal reasoning modules in manufacturing pipelines. Shenoy et al. (2025) develops an LLM-augmented hybrid neuromyotonic model for semiconductor SMT lines, integrating rule-based ontology, physics-based causal reasoning, and deep learning to achieve system-wide defect diagnosis under zero-defect manufacturing objectives.

3.2.3 Operational Decision Level

At this level, interventions are applied directly to policies and operational configurations, such as scheduling, allocation, resource distribution, and execution plans. Xiao et al. (2025) proposes xFBCI that performs privacy-preserving federated Bayesian causal inference to estimate average treatment effects of operational variables across distributed manufacturing sites. It controls confounding and quantifies uncertainty without sharing raw data, enabling evidence-based policy decisions. Feng et al. (2025) formulates the disassembly line balancing problem under uncertain component states as a dynamic multi-objective optimization problem. Its causal inference-inspired evolutionary algorithm adaptively reallocates tasks, robots, and workstation assignments as component conditions are revealed. While it does not implement explicit do-operator effect estimation, environmental changes trigger interventions on operational policies, modifying system configurations and execution plans in real time.

3.3 L3 Counterfactual Manufacturing Intelligence

L3 systems instantiate parallel factual and counterfactual trajectories, enabling root-cause attribution, failure replay, and prospective validation before real actuation. This marks a transition from causal explanation to causal foresight: manufacturing intelligence becomes capable of simulating alternate process evolutions, comparing policy paths, and quantifying uncertainty around hypothetical interventions. Under strict Pearl-style definition, very few manufacturing papers truly reach this level. Therefore, the present classification adopts a slightly broader but principled interpretation: systems are considered L3 if they operationalize structural counterfactual reasoning within an explicit causal model, even if they do not implement full unit-level abduction–action–prediction pipelines.

A foundational enabler of this transition is lifecycle governance for counterfactual models. Maier et al. (2024) formalizes an industrial lifecycle for SCM-based systems, positioning counterfactual reasoning as a managed, versioned, and auditable engineering artifact rather than an isolated analytical tool. By defining documentation standards, validation stages, and integration patterns, it reframes counterfactual models as long-lived infrastructure embedded within safety, diagnostics, and manufacturing workflows. Importantly, this work treats counterfactual reasoning as a core industrial capability requiring governance and traceability.

Beyond governance, operational counterfactual agents demonstrate how L3 reasoning becomes actionable in manufacturing environments. Liang et al. (2025) integrates causal discovery, total effect computation, and interactive counterfactual validation to support root cause analysis and operator-facing “what-if” validation in real manufacturing data. Given a learned linear SEM, counterfactual predictions are computed via effect propagation and compared against observed outcomes. A more infrastructure-oriented realization is presented by Pothireddi (2025). This architecture embeds SCMs inside digital twin platforms, maintaining a DAG that supports real-time “do” and “what-if” queries prior to control deployment. The counterfactual module combines identifiable effect computation and bounded simulation to evaluate hypothetical interventions under uncertainty. It operationalizes structural counterfactual simulation within a live cyber-physical loop. It is one of the clearest examples of counterfactual reasoning embedded directly in industrial control workflows. Liu et al. (2025) operationalizes counterfactual generation through an explicit generative causal model. Although the counterfactual mechanism operates at the level of representation and generative modeling, it explicitly implements structural counterfactual generation grounded in a causal model. Counterfactual intelligence also extends beyond the shop floor to networked industrial systems. Gottimukkala and Bhuram (2025) integrates SCMs with digital twin simulation and optimization layers to compare factual and hypothetical policy trajectories under disruption scenarios. Counterfactual reasoning is implemented by modifying policy variables within SCM and propagating effects to performance metrics for scenario comparison. This constitutes policy-level counterfactual simulation.

4. Analysis and Discussion: Common Consensus and Open Challenges

4.1 Common Consensus in Existing Literature

Consensus 1: Correlation-Based ML Is Insufficient for Manufacturing

Across the selected literatures, a recurring concern is that traditional machine learning models lack mechanistic interpretability and may be vulnerable to spurious correlations. Manufacturing environments involve confounders, distribution shifts, and process interdependencies, which challenge purely associative modeling approaches. Predictive accuracy alone does not guarantee correct mechanistic understanding. As a results, correlation-based approaches may produce unstable or misleading conclusions when deployed in dynamic industrial settings. This recognition motivates an increasing methodological shift from black-box ML toward explainable AI (XAI) and, more fundamentally, toward causal reasoning framework.

Consensus 2: Causal Structure Improves Interpretability and Traceability

Structured representations such as knowledge graphs, structural causal models, and digital twins are widely motivated as mechanisms to enhance auditability, root-cause traceability, interpretability, and cross-domain integration in manufacturing systems. Even when full intervention semantics are absent at L1, structured relational representations enable faster diagnosis and more transparent decision support. Causal modeling works (Wehner et al. 2023; Wang et al. 2025; Fekete et al. 2025) and KG-centered approaches (Wan et al. 2024), Xiao et al. (2023) converge on the importance of explicit structural representation. At L2, there is a methodological consensus that unbiased effect estimation requires formal adjustment mechanisms, such as backdoor adjustment, instrumental variable methods, causal discovery under structural assumptions, and robustness/refutation testing.

Consensus 3: Human-in-the-Loop Is Essential

In many studies, domain expert involvement is frequently positioned as critical to the deployment of causal systems in manufacturing. Experts are not merely end users of AI outputs but active participants in refining the causal graph, validating discovered relationships, and governing model evolution (Wehner et al. 2023; Moosavi et al. 2024; Liang et al. 2025). In safety-critical industrial contexts, automated causal discovery is typically complemented by domain-informed constraints, structural priors, and institutional governance mechanisms. Collectively, these works suggest that industrial causal intelligence is predominantly implemented as a human-centered decision-support paradigm rather than a fully autonomous system.

Consensus 4: Causality Must Be Embedded in Operational Contexts

Standalone causal graphs provide explanatory value, but their operational impact increases when embedded within decision workflows. Causal models demonstrate enhanced value when integrated into digital twins, optimization engines, scheduling systems, maintenance pipelines, or plant-wide monitoring infrastructures. Several L2 and L3 works embed causal models within digital twin architectures (Xia et al. 2025; Lippi et al. 2023; Pothireddi 2025; Gottimukkala and Bhuram 2025). When coupled with structural causal reasoning, digital twins provide the computational substrate for testing interventions, simulating fault propagation, and performing counterfactual evaluation without physical risk.

Consensus 5: Causal Intelligence Progresses from Traceability to Foresight

The surveyed body of work implicitly reflects a layered progression: L1: Traceable structured association and explanation, L2: Quantified intervention and effect estimation, L3: Twin-world counterfactual foresight and policy comparison. The distribution of papers across levels indicates that most existing work remains concentrated at L1 and L2, with much fewer fully operational L3 implementations. This pattern suggests that manufacturing AI is still in transition from structured explainability toward proactive causal foresight. Across levels, many works frame manufacturing intelligence as evolving from descriptive explanation toward proactive decision validation. In parallel, large models are increasingly integrated across levels. At L1, they assist in knowledge extraction, graph construction, and explanation generation. In some L2 systems, they interface with structured causal models to support reasoning and interactive diagnosis. Early L3 prototypes begin to explore coupling language interfaces with counterfactual engines, enabling operators to query “what-if” scenarios through natural language. In the reviewed literature, large models primarily function as knowledge extraction and interaction interfaces, while structural causal models remain the computational backbone for interventional reasoning.

4.2 Common Challenges

Challenge 1: Hidden Confounding and Partial Observability

A recurring limitation across the reviewed studies concerns the assumptions required for causal identification. Many methods implicitly rely on causal sufficiency, faithfulness, and the absence of unobserved confounders. These assumptions are analytically convenient but difficult to justify in real production systems. Manufacturing environments are only partially observable, for example, important process states such as tool wear, material microstructural changes, or operator interventions are often not directly measured. In addition, interactions across upstream and downstream processes are incompletely captured. Several surveys explicitly note that industrial datasets frequently violate i.i.d. assumptions and exhibit domain shifts, missing interventions, and incomplete observability (Vukovic and Thalmann 2022; Wang et al. 2025). Empirical works further illustrate how hidden confounders can distort learned causal structures or bias estimated effects (Wang et al. 2025; Sandt and Martin 2025). In multi-domain settings, where processes differ across factories or product lines, the stability of causal mechanisms becomes uncertain. These conditions weaken the credibility of estimated causal effects and limit external validity. Addressing hidden confounding and partial observability therefore remains a central methodological challenge rather than a peripheral issue.

Challenge 2: Scalability of Expert-Driven Causal Modeling

A substantial portion of the literature relies on expert-informed causal structures. At the physical-process level (L2.1), such strategies can yield controlled and technically meaningful effect estimation, particularly when grounded in well-understood mechanisms. However, the reliance on expert input introduces scalability constraints. Extending curated models to multiple factories, new product lines, or evolving process conditions require repeated manual effort and domain-specific adaptation. As causal reasoning is expanded to system-level (L2.2) or operational decision level (L2.3) settings, the number of variables, dependencies, and temporal interactions grows rapidly. High-dimensional graphs and dynamic dependencies increase computational complexity and reduce transparency. In addition, real-time inference requirements and explainability constraints in production environments may restrict the practical deployment of computationally intensive causal pipelines. Balancing interpretability, automation, and scalability remains an open methodological challenge in current research.

Challenge 3: Temporal and Dynamic Causality Is Underdeveloped

Manufacturing systems operate under continuous state transitions, performance drift, and event-driven changes. However, much of the existing causal modeling literature still relies on static DAG representations or assumes time-invariant mechanisms. Temporal information is often incorporated through lagged variables or window-based aggregation, but the underlying structure is still assumed to be time-invariant. Such simplifications are difficult to defend in environments characterized by equipment degradation, maintenance interventions, shifting demand patterns, or policy adjustments. Several recent works demonstrate awareness of temporal variability and attempt to encode dynamic information within graph-based frameworks (Xia et al. 2025; Zhou et al. 2023; Li et al. 2025). Nevertheless, even in these approaches, temporal dynamics are primarily embedded as evolving attributes, event links, or state updates layered onto an underlying structural graph. The deeper question of time-varying causal mechanisms of where the causal structure itself changes across regimes remains less systematically addressed. Issues such as delayed multi-step propagation, structural adaptation under regime shifts, and feedback-driven causal evolution are not yet treated with the same rigor as static causal discovery. As a result, while temporal extensions are emerging, dynamic causal modeling in manufacturing remains methodologically underdeveloped relative to the inherently time-dependent nature of production systems.

Challenge 4: Ambiguity of “Counterfactual” in smart manufacturing

Across the filtered and reviewed literature and, the term counterfactual is used with varying meanings, creating conceptual ambiguity in smart manufacturing research. There are very few genuinely L3-leaning papers, and most “counterfactual” claims in industrial AI correspond either to explainable AI recourse methods or to interventional what-if simulations rather than formal structural counterfactual inference. Only a small minority of works implement counterfactual generation grounded in an explicit structural or generative causal model. Importantly, “counterfactual” does not imply L3 by default. First, many so-called counterfactual approaches generate minimally modified inputs that flip a model’s prediction. These are counterfactual explanations of the machine learning model, not counterfactuals of the underlying physical manufacturing process. Second, some works use “counterfactual generation” for data augmentation or synthetic sample creation (e.g., in open-set diagnosis). Third, certain counterfactual root cause analysis approaches rely on graph-guided matching or propagation scoring without performing explicit

intervention or structural re-simulation. This ambiguity creates a gap between claimed and implemented causal levels. Clarifying this distinction is essential for accurately assessing progress toward genuine L3 manufacturing intelligence.

5. Conclusion and Future Directions

This review organized the literature on intelligent manufacturing through the three-level causal hierarchy and knowledge graph infrastructures. The analysis shows that most work remains concentrated at L1 (structured association) and L2 (intervention estimation), while fully operational L3 counterfactual systems are still limited. Although causal reasoning is increasingly embedded into digital twins and decision pipelines, methodological and deployment gaps persist.

Future research should move from diagnosing partial observability toward developing identifiable and transferable causal models that remain valid across heterogeneous industrial contexts. Emphasis should be placed on strengthening cross-domain validation and effect generalization rather than reiterating assumption violations. Scalable causal modeling across production systems requires architectural strategies that allow local mechanism refinement while maintaining global coherence. Modular and hierarchically decomposed causal representations may enable plant-wide reasoning without sacrificing interpretability. Dynamic and time-varying causality demands formal extensions beyond static structural assumptions. Models capable of representing regime shifts, delayed propagation, and evolving mechanisms are necessary for longitudinal intervention analysis. Finally, the integration of language and multimodal foundation models presents additional opportunities. Closer structural alignment between large models and explicit causal frameworks may allow multimodal intelligence to support intervention and counterfactual reasoning rather than remaining at the semantic level.

Reference

- Becker, M., Niemietz, P. and Bergs, T., Explainable Neural Network for Time Series-Based Condition Monitoring in Sheet Metal Shearing, *Journal of Intelligent Manufacturing*, 2025, <https://doi.org/10.1007/s10845-025-02596-3>.
- Bhakte, A., and Srinivasan, R., Potential for Counterfactual Explanations to Support Digitalized Plant Operations, *Computer Aided Chemical Engineering*, vol. 53, pp. 3073–3078, 2024, <https://doi.org/10.1016/B978-0-443-28824-1.50513-5>.
- Cai, J., Ren, L., Qin, W., Zhang, L., and Lai, L., A Knowledge Graph Reasoning Framework for Manufacturing Multi-View Heterogeneous Data, *Proceedings of the 35th European Modeling & Simulation Symposium (EMSS)*, 2023, <https://doi.org/10.46354/i3m.2023.emss.012>.
- Cen, W., Su, C., Chen, G., and Xie, L., Event Driven Causal Knowledge Graphs for Bottleneck Root Cause Analysis in Manufacturing Systems, *Expert Systems with Applications*, vol. 296, 2026, <https://doi.org/10.1016/j.eswa.2025.129117>.
- Fang, Y., Li, Z., Huang, K., Chen, J., Gui, L., and Chen, X., Causal Inference-Based Dynamic Optimization for Disassembly Line Balancing with Uncertain Component States, *Journal of Manufacturing Systems*, vol. 83, pp. 470–482, 2025, <https://doi.org/10.1016/j.jmsy.2025.09.018>.
- Fekete, T., Petrone, I. M. and Wicaksono, H., A Comprehensive Causal AI Framework for Analysing Factors Affecting Energy Consumption and Costs in Customised Manufacturing, *International Journal of Production Research*, vol. 64, pp. 1–38, 2025, <https://doi.org/10.1080/00207543.2025.2580541>.
- Gottimukkala, S. R. and Bhuram, S. K., Prescriptive Analytics for Next-Gen Supply Chains: Integrating Causal AI with Digital Twin Technologies, *2025 IEEE International Conference on Pattern Recognition, Machine Vision and Artificial Intelligence (PRMVAI)*, pp. 1–6, June 2025.
- Hoang, D., Gorsich, D., Castanier, M.P., Imani, F., Knowledge Graph Fusion with Large Language Models for Accurate, Explainable Manufacturing Process Planning, *arXiv preprint arXiv:2506.13026*, 2025.
- Huang, Z., Guo, X., Jiang, C., Yang, M., Xue, H., Zhao, W., and Wang, J., mKGMPP: A Multi-Layer Knowledge Graph Integration Framework and Its Inference Method for Manufacturing Process Planning, *Advanced Engineering Informatics*, vol. 65, 2025, <https://doi.org/10.1016/j.aei.2025.103266>.
- Ji, S., Wang, J., Wang, Z., and Zhang, F., Explainable Causal Graph-based Method for Diagnosis and Root Cause Analysis in Process Systems, *Process Safety and Environmental Protection*, vol. 203, 2025.
- Jeong, H., Fekete, T., Bashyal, A., and Wicaksono, H., From Theory to Practice: Implementing Causal AI in Manufacturing for Sustainability, *Procedia Computer Science*, vol. 253, pp. 1495–1504, 2025.
- Johannssen, A., Qiu, P., Yeganeh, A., and Chukhrova, N., Explainable AI for Trustworthy Intelligent Process Monitoring, *Computers & Industrial Engineering*, vol. 209, 2025, <https://doi.org/10.1016/j.cie.2025.111407>.

- Kim, H., Yun, J.Y., Jung, E., Lee, B., Kim, H., and Lee, J.M., Decoding Industrial-Scale Battery Manufacturing Process through Integration of Causal Graphs into Explainable Artificial Intelligence, *Engineering Applications of Artificial Intelligence*, vol. 159, 2025, <https://doi.org/10.1016/j.engappai.2025.111657>.
- Li, B., Li, Q., Du, T., Liu, D., Yang, Q., Chen, T., Xiong, J., Peng, B., Ren, J., and Zhao, J., Research, Application, and Challenges of Causal Inference in Industrial Fault Diagnosis: A Survey, *Engineering Applications of Artificial Intelligence*, vol. 158, 2025, <https://doi.org/10.1016/j.engappai.2025.111376>.
- Li, K., Liu, K., Zhai, S., Zhuang, C., and Pei, F., Automatic Generation Method of Knowledge Graph for Complex Product Assembly Processes Based on Text Mining, *Chinese Journal of Mechanical Engineering*, vol. 38, no. 1, p. 133, 2025, <https://doi.org/10.1186/s10033-025-01284-w>.
- Li, L., Yang, M., Makanda, I.L.D., and Jiang, P., Meta-Knowledge Triple Driven Multi-Modal Knowledge Graph Construction Method and Application in Production Line Control with Gantt Charts, *Journal of Manufacturing Systems*, vol. 80, pp. 224–242, 2025, <https://doi.org/10.1016/j.jmsy.2025.03.002>.
- Li, Q., Huang, T., Liu, J., and Wang, S., Causality-Guided Fault Diagnosis under Visual Interference in Fused Deposition Modeling, *Engineering Applications of Artificial Intelligence*, vol. 143, 2025.
- Li, Y., Ko, H. and Ameri, F., Integrating Graph Retrieval-Augmented Generation With Large Language Models for Supplier Discovery, *Journal of Computing and Information Science in Engineering*, vol. 25, 2025.
- Liang, T., Liu, T., Wang, J., Zhang, J., and Zheng, P., Causal Deep Learning for Explainable Vision-Based Quality Inspection under Visual Interference, *Journal of Intelligent Manufacturing*, vol. 36, no. 2, pp. 1363–1384, 2025.
- Lim, K. Y. H., Yosai, T. S., Chen, C. H., Zheng, P., Wang, L. and Xu, X., Graph-enabled cognitive digital twins for causal inference in maintenance processes, *International Journal of Production Research*, vol. 62, no. 13, pp. 4717–4734, 2024, <https://doi.org/10.1080/00207543.2023.2274335>.
- Lippi, M., Martinelli, M., Picone, M., and Zambonelli, F., Enabling Causality Learning in Smart Factories with Hierarchical Digital Twins, *Computers in Industry*, vol. 148, p. 103892, 2023.
- Liu, R., Hu, P., Zhao, S., Lin, D., Zhang, W., and Ding, S. X., Causal Counterfactual Faithfulness Generation for Open-Set Fault Diagnosis of Complex Industrial Processes, *IEEE Transactions on Industrial Informatics*, vol. 21, no. 12, pp. 9538–9549, 2025, doi: 10.1109/TII.2025.3598419.
- Ma, Y., Wu, T., Zhou, B., Liang, X., Du, J., and Bao, J., Enhancing Bottleneck Analysis in Ship Manufacturing with Knowledge Graphs and Large Language Models, *Machines*, vol. 13, no. 3, 2025.
- Ma, Y., Zheng, S., Yang, Z., Pan, H., and Hong, J., Knowledge-Graph Enhanced Large Language Model-Based Fault Diagnostic Reasoning and Maintenance Decision Support Pipeline towards Industry 5.0, *International Journal of Production Research*, vol. 0.0, pp. 1–22, 2025, <https://doi.org/10.1080/00207543.2025.2472298>.
- Madathil, A. P., Luo, X., Liu, Q., Walker, C., Madarkar, R., and Qin, Y., A Review of Explainable Artificial Intelligence in Smart Manufacturing, *International Journal of Production Research*, vol. 63, no. 23, pp. 8654–8697, 2025, <https://doi.org/10.1080/00207543.2025.2513574>.
- Maier, R., Schlattl, A., Guess, T., and Mottok, J., CausalOps -- Towards an Industrial Lifecycle for Causal Probabilistic Graphical Models, *Information and Software Technology*, vol. 174, 2024.
- Moosavi, S., Farajzadeh-Zanjani, M., Razavi-Far, R., Palade, V., and Saif, M., Explainable AI in Manufacturing and Industrial Cyber-Physical Systems: A Survey, *Electronics*, vol. 13, no. 17, 2024.
- Narayanan, A., Kulkarni, K., Riswan, F., and Salehan, M.A.B., Machine Learning Based Explainable Fault Detection of Vacuum Pumps in PV Solar Cell Manufacturing, *2024 Tenth Indian Control Conference (ICC)*, pp. 379–384, 2024, <https://doi.org/10.1109/ICC64753.2024.10883739>.
- Pearl, J., *Causality: Models, Reasoning, and Inference*, 2nd Edition, Cambridge University Press, 2009.
- Pothireddi, S. C., Causal Digital Twins: Real-Time Counterfactuals for Industrial Process Optimization, *Journal of Computer Science and Technology Studies*, vol. 7, No. 7, 2025.
- Razouk, H., Benischke, L., Garber, D., and Kern, R., Increasing the Accessibility of Causal Domain Knowledge via Causal Information Extraction Methods: A Case Study in the Semiconductor Manufacturing Industry, *Applied Sciences*, vol. 15, no. 5, 2025, <https://doi.org/10.3390/app15052573>.
- Ren, Y., He, Y., Zhang, X., Yen, A., and Li, G.P., A Cyber Manufacturing IoT System for Adaptive Machine Learning Model Deployment by Interactive Causality-Enabled Self-Labeling, *Machines*, vol. 13, no. 4, p. 304, 2025.
- Sandt, M. D. and Martin, R. J., Molding Material Reclamation: A Case Study of Manufacturing Process Analysis Using a Causal-Based Graphical Model Approach, *International Journal of Metalcasting*, vol. 19, no. 2, pp. 905–916, 2025, <https://doi.org/10.1007/s40962-024-01334-4>.
- Shenoy, M., Sahasrabudhe, S. and Ameri, F., Neurosymbolic AI-Driven Zero-Defect Manufacturing in Semiconductor Assembly: A Hybrid Framework, *2025 36th Annual SEMI Advanced Semiconductor Manufacturing Conference (ASMC)*, pp. 1–6, May 2025, <https://doi.org/10.1109/ASMC64512.2025.11010401>.

- Shyalika, C., Sharma, A., Kalach, F.E., Jaimini, U., Henson, C., Harik, R., and Sheth, A., CausalTrace: A Neurosymbolic Causal Analysis Agent for Smart Manufacturing, *arXiv preprint arXiv:2510.12033*, 2025.
- Streltsov, O., Terziyan, V. and Vitko, O., Discover and Explore Weak Causality and Causal Disposition in Images for Smart Manufacturing Tasks, *Proceedings of the 6th International Conference on Industry 4.0 and Smart Manufacturing*, vol. 253, pp. 187–198, 2025, <https://doi.org/10.1016/j.procs.2025.01.082>.
- Su, C., Jiang, Q., Han, Y., Wang, T., and He, Q., Knowledge Graph-Driven Decision Support for Manufacturing Process: A Graph Neural Network-Based Knowledge Reasoning Approach, *Advanced Engineering Informatics*, vol. 64, 2025, <https://doi.org/10.1016/j.aei.2024.103098>.
- Terziyan, V. and Vitko, O., Causality-Aware Convolutional Neural Networks for Advanced Image Classification and Generation, *Procedia Computer Science*, vol. 217, pp. 495–506, 2023.
- Vuković, M. and Thalmann, S., Causal Discovery in Manufacturing: A Structured Literature Review, *Journal of Manufacturing and Materials Processing*, vol. 6, no. 1, 2022, <https://doi.org/10.3390/jmmp6010010>.
- Wan, Y., Liu, Y., Chen, Z., Chen, C., Li, X., Hu, F., and Packianather, M., Making Knowledge Graphs Work for Smart Manufacturing: Research Topics, Applications and Prospects, *Journal of Manufacturing Systems*, vol. 76, pp. 103–132, 2024, <https://doi.org/10.1016/j.jmsy.2024.07.009>.
- Wang, H., Xu, Y., Peng, T., Agbozo, R.S.K., Xu, K., Liu, W., and Tang, R., Two-Stage Approach to Causality Analysis-Based Quality Problem Solving for Discrete Manufacturing Systems, *Journal of Engineering Design*, vol. 36, no. 11, pp. 1879–1903, 2025, <https://doi.org/10.1080/09544828.2023.2247858>.
- Wang, Y., Qiang, S., Yue, X., Li, T., and Zhang, K., A Two-Layer Causal Knowledge Network Construction Method Based on Quality Problem-Solving Data, *Systems*, vol. 13, no. 3, 2025, <https://doi.org/10.3390/systems13030142>.
- Wang, Z., Yan, Z., Li, S., and Liu, J., IndVisSGG: VLM-Based Scene Graph Generation for Industrial Spatial Intelligence, *Advanced Engineering Informatics*, vol. 65, 2025, <https://doi.org/10.1016/j.aei.2024.103107>.
- Wehner, C., Kertel, M. and Wewerka, J., Interactive and Intelligent Root Cause Analysis in Manufacturing with Causal Bayesian Networks and Knowledge Graphs, *2023 IEEE 97th Vehicular Technology Conference (VTC2023-Spring)*, pp. 1–7, June 2023, <https://doi.org/10.1109/VTC2023-Spring57618.2023.10199563>.
- Wu, P., Mou, X., Gong, L., Tu, H., Qiu, L., and Yang, B., An Automatic Machine Fault Identification Method Using the Knowledge Graph-Embedded Large Language Model, *The International Journal of Advanced Manufacturing Technology*, vol. 138, no. 2, pp. 725–739, 2025.
- Wu, Z., Wu, H., Tan, J., Qian, Y., He, L., and Goh, M., Integrating Context Awareness and Knowledge Graphs for Enhanced Knowledge Recommendation in Manufacturing Process Planning, *Advanced Engineering Informatics*, vol. 71, 2026, <https://doi.org/10.1016/j.aei.2026.104324>.
- Xia, M., Mo, X., Zhang, Y., and Hu, X., A Knowledge Graph Construction and Causal Structure Mining Approach for Non-Stationary Manufacturing Systems, *Robotics and Computer-Integrated Manufacturing*, vol. 95, 2025.
- Xiao, X., Alharbi, K., Zhang, P., Qin, H., and Yue X., Bayesian Federated Causal Inference and Its Application in Manufacturing, *Journal of Intelligent Manufacturing*, 2025, <https://doi.org/10.1007/s10845-025-02665-7>.
- Xiao, Y., Zheng, S., Shi, J., Du, X., and Hong, J., Knowledge Graph-Based Manufacturing Process Planning: A State-of-the-Art Review, *Journal of Manufacturing Systems*, vol. 70, pp. 417–435, 2023.
- Youssef, S., Causality-Driven AI for Manufacturing Systems, *2024 Joint Workshop of Fraunhofer IOSB and Institute for Anthropomatics, Vision and Fusion Laboratory*, 2024.
- Yu, J., Pychynski, T. and Huber, M. F., Causal Mechanism Estimation in Multi-Sensor Systems Across Multiple Domains, *arXiv preprint arXiv:2507.17792*, 2025, <https://doi.org/10.48550/arXiv.2507.17792>.
- Yu, J., Pychynski, T. and Huber, M. F., Multi-Domain Causal Discovery and Downstream Applications in Manufacturing Systems, *2025 IEEE Sensor Data Fusion: Trends, Solutions, Applications (SDF)*, pp. 1–8, 2025.
- Zhang, Q., Xu, C., Li, J., Sun, Y., Bao, J., and Zhang, D., LLM-TSFD: An Industrial Time Series Human-in-the-Loop Fault Diagnosis Method Based on a Large Language Model, *Expert Systems with Applications*, vol. 264, 2025.
- Zhou, B., Li, X., Liu, T., Xu, K., Liu, W., and Bao, J., CausalKGPT: Industrial Structure Causal Knowledge-Enhanced Large Language Model for Cause Analysis of Quality Problems in Aerospace Product Manufacturing, *Advanced Engineering Informatics*, vol. 59, 2024, <https://doi.org/10.1016/j.aei.2023.102333>.
- Zhou, B., Shen X., Lu, Y., Li, X., Hua, B., Liu, T., and Bao, J., Semantic-Aware Event Link Reasoning over Industrial Knowledge Graph Embedding Time Series Data, *International Journal of Production Research*, vol. 61, no. 12, pp. 4117–4134, 2023, <https://doi.org/10.1080/00207543.2021.2022803>.

Biographies

Yiyun (Cindy) Fei is a Ph.D. candidate in Mechanical Engineering at Georgia Institute of Technology, Georgia, USA, where she also earned her bachelor's and master's degree in the same field. Her research focuses on the modeling,

analysis, and optimization of complex operational systems to support data-driven design and decision-making. Yiyun worked as a Manufacturing and Process Development Engineer at TE Connectivity, where she gained four years of hands-on industry experience. During her time there, she led and contributed to multiple projects involving AI-based machine vision, process automation, and discrete event simulation of manufacturing systems.

Roger Jiao is the editor-in-chief of Journal of Engineering Design and an associate professor of mechanical and industrial engineering at Georgia Tech, USA. Prior to joining the School of Mechanical Engineering at Georgia Tech in December 2008, he was an Assistant Professor and then Associate Professor in the School of Mechanical and Aerospace Engineering at Nanyang Technological University, Singapore. Before his career in Singapore, he was a Visiting Scholar in the Department of Industrial Engineering and Engineering Management at Hong Kong University of Science and Technology from 1998 to 1999. From 1993 to 1994, he was a Lecturer of Industrial Engineering in the School of Management at Tianjin University, China, and from 1988 to 1990, he worked as an Associate Lecturer in the Department of Industrial Design at Tianjin University of Science and Technology, China. More info about his research: <https://scholar.google.com/citations?user=9yikEHAAAAAJ&hl=en&oi=ao>.