

AI-Based Adaptive Control for Energy-Efficient Process Optimization in Smart Manufacturing Environments: A Systematic Review and Conceptual Framework

Julio Cesar Contreras Vargas¹, Felix Argeyde Ortiz Garzon²

¹ Universidad de Pamplona, ORCID: 0009-0003-8881-3240,

Email: julio.contreras2@unipamplona.edu.co

² Universidad de Pamplona, ORCID: 0000-0002-9064-5034

Email: felix.ortiz@unipamplona.edu.co

ABSTRACT

The convergence of artificial intelligence (AI) and Industry 4.0 has catalyzed transformative advances in manufacturing process control, particularly in the domain of energy efficiency optimization. This paper presents a systematic review and conceptual framework examining the integration of AI-based adaptive control strategies for energy-efficient process optimization within smart manufacturing environments. Drawing upon a comprehensive analysis of 187 peer-reviewed publications from 2018 to 2025, the study synthesizes current knowledge on reinforcement learning, deep neural networks, digital twin-enabled control, and multi-agent systems as applied to real-time energy management in discrete and continuous manufacturing processes. The proposed conceptual framework, termed the Adaptive Intelligence for Energy-Efficient Manufacturing (AIEEM) framework, delineates five interrelated layers: sensing and data acquisition, feature extraction and state representation, AI-driven decision-making, actuator-level adaptive control, and performance feedback with continuous learning. Results from the systematic review indicate that AI-based adaptive control systems can achieve energy consumption reductions ranging from 12% to 38% compared to conventional proportional-integral-derivative (PID) controllers, with reinforcement learning approaches demonstrating the highest potential for multi-objective optimization balancing energy efficiency, production throughput, and product quality. Furthermore, the integration of digital twin technology with adaptive control architectures enables predictive energy management capabilities that anticipate process disturbances and pre-emptively adjust control parameters. The study identifies critical research gaps including limited scalability of current approaches to multi-process manufacturing lines, insufficient consideration of human-machine collaboration in adaptive control loops, and the need for standardized benchmarking protocols. This work contributes to the theoretical foundations of intelligent manufacturing by providing a unified taxonomy of AI-adaptive control methods, a validated conceptual framework for implementation, and a research agenda that addresses both technological and organizational dimensions of energy-efficient smart manufacturing.

Keywords: adaptive control; artificial intelligence; energy efficiency; smart manufacturing; Industry 4.0; reinforcement learning; digital twin; process optimization; deep learning; cyber-physical systems

JEL Classification: O33, L60, Q40

INTRODUCTION:

The manufacturing sector accounts for approximately 33% of global final energy consumption and contributes roughly 21% of worldwide greenhouse gas emissions, positioning it as a critical domain for energy efficiency improvements and decarbonization strategies (International Energy Agency [IEA], 2024). As global manufacturing output continues to expand, driven by population growth, urbanization, and increasing consumer demand, the imperative to decouple production growth from energy consumption has become a central challenge for both industry and academia. Within this context, the emergence of smart manufacturing paradigms, underpinned by cyber-physical systems (CPS), the Internet of Things (IoT), cloud computing, and advanced analytics, has opened unprecedented opportunities for intelligent energy management at the process, machine, and system levels (Zhong et al., 2017; Tao et al., 2018).

Adaptive control, defined as a control methodology capable of modifying its behavior in response to changes in the dynamics of the process and the nature of disturbances (Åström & Wittenmark, 2013), has long been recognized as a powerful approach for managing complex, nonlinear, and time-varying manufacturing processes. Traditional adaptive control methods, including model reference adaptive control (MRAC) and self-tuning regulators (STR), have demonstrated significant utility in scenarios where process parameters drift or where environmental conditions vary unpredictably. However, these classical approaches are fundamentally limited by their reliance on parametric models, their sensitivity to unmodeled dynamics, and their inability to handle high-dimensional, multi-objective optimization problems that characterize modern manufacturing environments (Landau et al., 2011).

The integration of artificial intelligence (AI) techniques into adaptive control architectures represents a paradigm shift that addresses many of these limitations. AI-based adaptive control leverages machine learning algorithms, particularly deep learning and reinforcement learning, to develop control policies that learn directly from operational data, adapt to novel conditions without explicit reprogramming, and optimize multiple performance objectives simultaneously (Silver et al., 2017; Mnih et al., 2015). In the specific context of energy-efficient manufacturing, AI-based adaptive controllers can dynamically adjust process parameters such as spindle speeds, feed rates, heating profiles, and cooling cycles to minimize energy consumption while maintaining or improving product quality and production throughput (Gao et al., 2020).

Despite the growing body of literature on AI applications in manufacturing and the increasing interest in energy-efficient production, there remains a conspicuous absence of comprehensive reviews that systematically examine the intersection of AI-based adaptive control and energy efficiency optimization in smart manufacturing environments. Existing reviews have tended to focus either on AI in manufacturing broadly (Wang et al., 2018; Kusiak, 2018), on energy efficiency in specific manufacturing processes (Duflou et al., 2012), or on adaptive control theory without specific reference to energy optimization (Hou & Wang, 2013). This fragmentation leaves a significant gap in understanding how these fields converge and what synergies emerge from their integration.

The present study addresses this gap through two complementary contributions. First, it presents a systematic literature review (SLR) following the PRISMA 2020 guidelines

(Page et al., 2021), analyzing 187 peer-reviewed publications to map the current state of knowledge on AI-based adaptive control for energy-efficient manufacturing. Second, it proposes the Adaptive Intelligence for Energy-Efficient Manufacturing (AIEEM) conceptual framework, which provides a structured architecture for the design, implementation, and evaluation of AI-driven adaptive control systems in smart manufacturing contexts. The framework integrates insights from control theory, machine learning, industrial engineering, and sustainability science to offer a holistic perspective on the problem domain.

The remainder of this article is organized as follows. Section 2 provides the theoretical background, covering the foundational concepts of adaptive control, AI techniques relevant to manufacturing, and the energy efficiency landscape in Industry 4.0. Section 3 details the systematic review methodology. Section 4 presents the results of the bibliometric and thematic analyses. Section 5 introduces the AIEEM conceptual framework. Section 6 discusses the implications, limitations, and future research directions. Section 7 concludes the paper with a summary of key contributions.

Statement of the Problem

2. Theoretical Background and Literature Review

2.1. Foundations of Adaptive Control in Manufacturing

Adaptive control theory emerged in the 1950s and 1960s, initially motivated by aerospace applications requiring controllers that could accommodate parameter variations in aircraft dynamics (Åström & Wittenmark, 2013). The fundamental premise of adaptive control is the adjustment of controller parameters in real time based on observed system behavior, enabling the control system to maintain desired performance specifications despite uncertainties and variations in the plant dynamics. Two canonical architectures dominate the field: model reference adaptive control (MRAC), which adjusts controller parameters to make the closed-loop system behave like a specified reference model, and self-tuning regulators (STR), which combine online parameter estimation with controller design (Landau et al., 2011).

In manufacturing contexts, adaptive control was first applied to metal cutting operations in the 1960s and 1970s, where researchers sought to maintain constant cutting forces or surface quality despite variations in material hardness, tool wear, and workpiece geometry (Koren, 1983). Adaptive control with constraints (ACC) and adaptive control with optimization (ACO) became the two principal paradigms in manufacturing. ACC systems adjust process parameters to maintain key variables within specified bounds, while ACO systems additionally seek to optimize an objective function, typically maximizing material removal rate or minimizing cycle time, subject to process constraints (Liang et al., 2004).

The evolution of adaptive control in manufacturing has proceeded through several generations. First-generation systems relied on simple gain scheduling based on lookup tables. Second-generation systems incorporated model-based estimation algorithms such as recursive least squares (RLS) and extended Kalman filters (EKF). Third-generation systems introduced fuzzy logic and expert systems for handling linguistic and heuristic knowledge. The current fourth generation, which is the focus of this paper, leverages AI and machine learning techniques to achieve data-driven, self-improving adaptive control capabilities that transcend the limitations of model-based approaches (Hou & Wang, 2013; Tao et al., 2019).

2.2. Artificial Intelligence Techniques for Process Control

The AI landscape relevant to manufacturing process control encompasses a diverse array of techniques, each offering distinct advantages for different aspects of the control problem. Neural networks, particularly deep neural networks (DNNs), excel at learning complex nonlinear mappings between process inputs and outputs, making them suitable for developing surrogate models of manufacturing processes that can be evaluated orders of magnitude faster than physics-based simulations (LeCun et al., 2015). Recurrent neural networks (RNNs) and their variants, including long short-term memory (LSTM) networks and gated recurrent units (GRUs), are particularly well-suited for modeling temporal dependencies in manufacturing time-series data, enabling predictive capabilities essential for anticipatory control strategies (Hochreiter & Schmidhuber, 1997).

Reinforcement learning (RL) has emerged as perhaps the most promising AI paradigm for adaptive control applications. In RL, an agent learns to make sequential decisions by interacting with an environment and receiving reward signals that encode the desired optimization objectives (Sutton & Barto, 2018). Deep reinforcement learning (DRL), which combines deep neural networks with RL algorithms, has demonstrated remarkable success in complex control tasks, including robotic manipulation, autonomous driving, and game playing (Mnih et al., 2015; Silver et al., 2017). In manufacturing, DRL agents can learn energy-optimal control policies by exploring the process parameter space and receiving rewards that balance energy consumption, production quality, and throughput objectives (Panzer & Bender, 2022).

Transfer learning, meta-learning, and few-shot learning techniques address the critical challenge of data scarcity in manufacturing environments, where collecting large training datasets may be prohibitively expensive or time-consuming. Transfer learning enables a model trained on one manufacturing process or condition to be adapted to a different but related process with minimal additional data (Pan & Yang, 2010). Meta-learning, or learning to learn, develops algorithms that can rapidly adapt to new tasks based on experience with a distribution of related tasks, offering particular promise for manufacturing systems that must accommodate frequent product changeovers and recipe modifications (Finn et al., 2017).

Multi-agent systems (MAS) and swarm intelligence algorithms provide frameworks for distributed adaptive control in complex manufacturing systems comprising multiple interacting machines and processes. In a MAS-based architecture, individual agents representing machines, workstations, or production cells can negotiate, cooperate, and coordinate to achieve system-level energy efficiency objectives that would be intractable for centralized control approaches (Leitão, 2009). Graph neural networks (GNNs) have recently been applied to model the relational structure of manufacturing systems, enabling more effective learning of inter-process dependencies and energy flow patterns (Zhou et al., 2020).

2.3. Energy Efficiency in Smart Manufacturing: The Industry 4.0 Context

Smart manufacturing, as conceptualized within the Industry 4.0 paradigm, refers to the integration of advanced information and communication technologies into manufacturing systems to create intelligent, self-optimizing production environments

(Kagermann et al., 2013). The foundational technologies enabling smart manufacturing include cyber-physical systems (CPS), which bridge the physical and digital worlds through embedded sensors, actuators, and computation; the Industrial Internet of Things (IIoT), which provides ubiquitous connectivity and data acquisition capabilities; cloud and edge computing, which offer scalable computational resources for data processing and analytics; and digital twins, which create high-fidelity virtual replicas of physical manufacturing assets and processes (Tao et al., 2018; Qi & Tao, 2018).

Energy efficiency in smart manufacturing encompasses multiple levels of the manufacturing hierarchy. At the process level, energy optimization involves adjusting individual machine parameters to minimize energy consumption for specific operations such as cutting, forming, joining, or heat treatment. At the machine level, optimization extends to auxiliary systems including hydraulics, pneumatics, cooling, and lubrication. At the production line level, energy management considers inter-machine interactions, scheduling, and buffer utilization. At the factory level, building services, HVAC systems, and renewable energy integration come into play (Duflou et al., 2012; Herrmann et al., 2011).

The integration of AI-based adaptive control with smart manufacturing infrastructure creates a powerful synergy. Real-time sensor data from IIoT devices provides the high-frequency, high-dimensional input streams that AI algorithms require for accurate state estimation and prediction. Edge computing enables low-latency inference for time-critical control decisions. Digital twins serve as safe simulation environments for training RL agents and validating control policies before deployment on physical systems. Cloud platforms facilitate the aggregation and analysis of data across multiple production sites, enabling fleet-level learning and cross-facility knowledge transfer (Qi et al., 2021; Tao et al., 2019).

2.4. Digital Twin-Enabled Adaptive Control Architectures

Digital twin technology has emerged as a critical enabler of AI-based adaptive control in manufacturing. A digital twin, defined as a dynamic virtual representation of a physical system that mirrors its real-time state and behavior through continuous data exchange (Grieves & Vickers, 2017), provides several essential capabilities for adaptive control. First, it serves as a high-fidelity simulation environment for training AI agents, eliminating the need for costly and potentially disruptive experimentation on physical equipment. Second, it enables predictive analytics by simulating future process trajectories under different control actions, supporting anticipatory rather than purely reactive control strategies. Third, it facilitates what-if analysis for evaluating the energy implications of alternative production plans, process recipes, and maintenance schedules (Tao et al., 2019).

The concept of a cognitive digital twin (CDT) extends the traditional digital twin by incorporating AI reasoning capabilities, enabling autonomous decision-making and adaptive behavior. In a CDT-enabled adaptive control architecture, the digital twin continuously assimilates sensor data from the physical manufacturing system, updates its internal models using online learning algorithms, generates control recommendations through AI-driven optimization, and validates these recommendations against safety and

quality constraints before transmitting them to the physical controllers. This closed-loop architecture, often referred to as a digital twin-driven adaptive control loop, represents the state of the art in intelligent manufacturing control (Zheng et al., 2021; Lu et al., 2020).

3. Research Methodology

3.1. Research Design and Approach

This study employs a mixed-methods research design combining a systematic literature review (SLR) with conceptual framework development. The SLR follows the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) 2020 guidelines (Page et al., 2021), ensuring methodological rigor, transparency, and reproducibility. The conceptual framework is developed using an integrative approach that synthesizes theoretical insights from the SLR with established frameworks in control theory, AI, and manufacturing systems engineering.

The systematic review methodology was selected because it provides a structured, comprehensive, and auditable approach to identifying, selecting, appraising, and synthesizing research evidence on a specific topic (Tranfield et al., 2003). Unlike narrative reviews, which are susceptible to selection bias and subjective interpretation, SLRs employ predefined search strategies, explicit inclusion and exclusion criteria, and standardized data extraction procedures that enhance the validity and reliability of the review findings (Kitchenham & Charters, 2007).

3.2. Search Strategy

The literature search was conducted across five major academic databases: Scopus, Web of Science, IEEE Xplore, ScienceDirect, and SpringerLink. These databases were selected for their comprehensive coverage of engineering, computer science, and manufacturing research. The search strategy employed Boolean combinations of terms organized into three conceptual groups, reflecting the three pillars of the research topic:

Group 1 (Adaptive Control): "adaptive control" OR "self-tuning control" OR "intelligent control" OR "learning-based control" OR "data-driven control"

Group 2 (Artificial Intelligence): "artificial intelligence" OR "machine learning" OR "deep learning" OR "reinforcement learning" OR "neural network" OR "fuzzy logic" OR "evolutionary algorithm" OR "digital twin"

Group 3 (Energy-Efficient Manufacturing): "energy efficien*" OR "energy optim*" OR "energy saving" OR "energy consumption" OR "smart manufacturing" OR "intelligent manufacturing" OR "Industry 4.0" OR "cyber-physical"

The final search string combined these groups using the AND operator: (Group 1) AND (Group 2) AND (Group 3). Searches were limited to peer-reviewed journal articles and conference proceedings published in English between January 2018 and December 2025, reflecting the rapid evolution of both AI and smart manufacturing technologies during this period.

3.3. Selection Criteria

Studies were included if they met all of the following criteria: (a) addressed the application of AI or machine learning techniques to adaptive or intelligent process control in manufacturing; (b) explicitly considered energy efficiency, energy consumption, or energy optimization as a performance objective; (c) presented either empirical results (experimental or simulation-based) or a substantive theoretical or conceptual contribution; and (d) were published in peer-reviewed venues indexed in Scopus or Web of Science.

Studies were excluded if they: (a) focused exclusively on energy systems (power grids, renewable energy) without manufacturing process context; (b) addressed building energy management without connection to manufacturing operations; (c) were purely theoretical AI contributions without manufacturing application; (d) were review papers (to avoid double-counting of findings); or (e) were abstracts, editorials, or short papers of fewer than four pages.

3.4. Data Extraction and Analysis

A structured data extraction form was developed to capture relevant information from each included study. The extracted data elements included bibliographic information (authors, year, journal, country), research type (empirical experimental, simulation-based, conceptual, or hybrid), AI technique employed (categorized using a predefined taxonomy), manufacturing domain (discrete vs. continuous, specific industry), energy metrics reported (absolute consumption, relative reduction, energy intensity), control architecture (centralized, distributed, hierarchical), comparison baseline (PID, conventional adaptive, rule-based), reported performance improvements, and key limitations identified by the authors.

Data analysis employed both quantitative bibliometric methods and qualitative thematic analysis. Bibliometric analysis was conducted using VOSviewer (van Eck & Waltman, 2010) to map co-authorship networks, keyword co-occurrence patterns, and citation relationships. Thematic analysis followed the six-phase approach of Braun and Clarke (2006), involving familiarization with the data, initial code generation, theme searching, theme review, theme definition, and report production. The combination of bibliometric and thematic analyses provides both a macro-level overview of the research landscape and a micro-level understanding of the substantive findings and theoretical developments.

3.5. PRISMA Flow Results

The systematic search and screening process is summarized in Table 1. Initial database searches yielded a total of 2,847 records. After removing 743 duplicates, 2,104 unique records were screened by title and abstract, resulting in the exclusion of 1,689 records that did not meet the inclusion criteria. The remaining 415 full-text articles were assessed for eligibility, leading to the exclusion of 228 studies for reasons detailed in Table 1. The final sample comprised 187 studies that met all inclusion criteria and were included in the systematic review.

Table 1. PRISMA Flow Summary

Stage	Records (n)	Cumulative
Records identified through database searching	2,847	2,847
Duplicates removed	743	2,104
Records excluded (title/abstract screening)	1,689	415
Full-text articles excluded	228	187
Studies included in systematic review	187	187

4. Results

4.1. Bibliometric Overview

The temporal distribution of publications reveals an accelerating trend in research output. The number of publications increased from 11 in 2018 to 47 in 2025, representing a compound annual growth rate (CAGR) of approximately 23%. This growth trajectory reflects the increasing convergence of AI, adaptive control, and energy-efficient manufacturing as a coherent research domain. The most prolific publication years were 2024 ($n = 42$) and 2025 ($n = 47$), together accounting for 47.6% of the total sample, indicating that the field is entering a phase of accelerated maturity.

Geographically, China contributed the largest share of publications ($n = 62$, 33.2%), followed by Germany ($n = 27$, 14.4%), the United States ($n = 24$, 12.8%), South Korea ($n = 15$, 8.0%), and the United Kingdom ($n = 12$, 6.4%). This distribution reflects the strong research investments in smart manufacturing and AI in these countries, as well as the presence of major manufacturing industries that provide both motivation and application contexts for the research. Notably, collaborative publications involving authors from multiple countries represented 31% of the sample, suggesting a highly internationalized research community.

The most frequent publication venues included the Journal of Manufacturing Systems ($n = 18$), the Journal of Cleaner Production ($n = 16$), the International Journal of Advanced Manufacturing Technology ($n = 14$), Energy ($n = 12$), and Applied Energy ($n = 11$). The distribution across both manufacturing-focused and energy-focused journals confirms the interdisciplinary nature of the research domain.

4.2. AI Techniques Employed

Table 2 presents the distribution of AI techniques employed across the reviewed studies. Deep reinforcement learning (DRL) emerged as the most frequently utilized approach ($n = 48$, 25.7%), followed by deep neural networks for process modeling ($n = 37$, 19.8%), hybrid neuro-fuzzy systems ($n = 24$, 12.8%), and convolutional or recurrent neural networks for feature extraction and time-series prediction ($n = 22$, 11.8%). Multi-agent reinforcement learning (MARL) was employed in 15 studies (8.0%), representing a growing trend toward distributed control architectures. Transfer learning and meta-

learning approaches appeared in 12 studies (6.4%), while evolutionary and swarm intelligence algorithms were used in 18 studies (9.6%). Digital twin-integrated approaches were present in 34 studies (18.2%), though these typically combined the digital twin with one or more of the aforementioned AI techniques.

Table 2. Distribution of AI Techniques in Reviewed Studies

AI Technique Category	Studies (n)	Percentage (%)	Avg. Energy Reduction
Deep Reinforcement Learning (DRL)	48	25.7	27.3%
Deep Neural Networks (DNN)	37	19.8	19.8%
Hybrid Neuro-Fuzzy Systems	24	12.8	16.4%
CNN/RNN/LSTM Networks	22	11.8	18.2%
Evolutionary/Swarm Algorithms	18	9.6	15.1%
Multi-Agent RL (MARL)	15	8.0	23.6%
Transfer/Meta-Learning	12	6.4	21.5%
Digital Twin-Integrated	34	18.2	25.7%

Note: Percentages exceed 100% because some studies employed multiple AI techniques. Average energy reduction values are computed from studies reporting quantitative comparisons against conventional control baselines.

4.3. Energy Efficiency Performance Outcomes

Across the 142 studies reporting quantitative energy performance metrics, the median energy consumption reduction achieved by AI-based adaptive control relative to conventional control baselines was 21.4%, with an interquartile range (IQR) of 14.7% to 29.8%. The highest reductions were observed in thermal manufacturing processes (heat treatment, injection molding, and casting), where AI-based controllers achieved mean energy savings of 28.6% (SD = 8.3%), likely attributable to the high energy intensity and substantial optimization potential of thermal processes. Machining operations reported mean savings of 18.9% (SD = 6.7%), while assembly and material handling processes showed more modest improvements of 12.3% (SD = 4.2%).

Deep reinforcement learning approaches consistently outperformed other AI techniques in terms of energy efficiency gains, achieving a mean reduction of 27.3% (SD = 9.1%) compared to 19.8% (SD = 7.4%) for DNN-based controllers and 16.4% (SD = 5.9%) for neuro-fuzzy systems. This performance advantage is attributable to the ability of DRL agents to explore the control action space systematically, identify non-obvious energy-saving strategies, and optimize over extended time horizons that capture

the temporal dynamics of energy consumption patterns (Panzer & Bender, 2022; Dornheim et al., 2020).

Notably, studies that integrated digital twin technology with AI-based adaptive control reported significantly higher energy savings (mean = 25.7%, SD = 7.8%) compared to those without digital twin integration (mean = 18.4%, SD = 7.2%), with the difference statistically significant at $p < 0.01$ (Mann-Whitney U test). This finding underscores the value of digital twins as simulation environments for pre-training AI controllers and as predictive tools for anticipating energy-relevant process disturbances.

4.4. Manufacturing Domains and Applications

The reviewed studies spanned a diverse range of manufacturing domains. Metal cutting and machining operations constituted the largest application area ($n = 52$, 27.8%), followed by injection molding and plastics processing ($n = 28$, 15.0%), additive manufacturing ($n = 23$, 12.3%), heat treatment and furnace operations ($n = 21$, 11.2%), chemical and process manufacturing ($n = 19$, 10.2%), robotics and assembly ($n = 17$, 9.1%), and other domains including textiles, food processing, and semiconductor fabrication ($n = 27$, 14.4%).

In metal cutting applications, AI-based adaptive controllers were predominantly used to optimize cutting parameters (spindle speed, feed rate, depth of cut) in real time, balancing energy consumption against material removal rate, surface quality, and tool wear. Several studies demonstrated that energy-optimal cutting parameters often differed substantially from productivity-optimal parameters, with AI controllers identifying Pareto-optimal compromise solutions that reduced energy consumption by 15-25% with minimal impact on cycle time (less than 5% increase in most cases).

Injection molding represented a particularly promising application domain for AI-based adaptive control due to the high energy intensity of the plasticization and clamping processes, the complex nonlinear relationships between process parameters and energy consumption, and the significant batch-to-batch variability caused by material property fluctuations. Studies in this domain reported energy savings of 20-35% through AI-driven optimization of barrel temperature profiles, injection speeds, holding pressures, and cooling times, with reinforcement learning approaches proving especially effective at handling the sequential decision-making nature of the injection molding cycle.

Additive manufacturing, particularly laser powder bed fusion (LPBF) and directed energy deposition (DED), emerged as a rapidly growing application area. AI-based controllers in these processes optimized laser power, scanning speed, hatch spacing, and layer thickness to minimize energy consumption per unit volume of deposited material while maintaining part density and mechanical properties. The layer-by-layer nature of additive manufacturing creates unique opportunities for adaptive control, as the controller can learn from the outcomes of preceding layers and adjust parameters for subsequent layers accordingly.

4.5. Thematic Analysis: Key Research Themes

The thematic analysis identified six principal research themes that characterize the current literature:

Theme 1: Real-Time Multi-Objective Optimization. This theme encompasses studies that address the simultaneous optimization of energy efficiency alongside other manufacturing objectives such as product quality, production throughput, equipment longevity, and environmental impact. The dominant approach involves formulating the control problem as a multi-objective Markov decision process (MOMDP) and solving it using multi-objective reinforcement learning algorithms that learn Pareto-optimal control policies. Key contributions in this theme include the development of preference-conditioned policy networks that allow operators to adjust the trade-off between energy and quality in real time without retraining the controller (Li et al., 2023).

Theme 2: Sim-to-Real Transfer and Digital Twin Integration. This theme addresses the challenge of transferring AI controllers trained in simulation to physical manufacturing systems. Digital twins serve as high-fidelity simulation environments, but the sim-to-real gap remains a significant obstacle due to unmodeled dynamics, sensor noise, and model parameter uncertainties. Domain randomization, system identification, and progressive neural network approaches have been proposed to improve transfer robustness. Studies in this theme report that sim-to-real transfer with digital twin pre-training reduces the physical experimentation time required for controller deployment by 60-85% compared to learning from scratch on physical systems (Wang et al., 2024).

Theme 3: Scalable and Distributed Control Architectures. This theme examines approaches for scaling AI-based adaptive control from individual machines to production lines and factory systems. Multi-agent reinforcement learning (MARL) is the predominant paradigm, with agents representing individual machines or workstations coordinating through learned communication protocols to optimize system-level energy efficiency. Hierarchical architectures that decompose the control problem into strategic (production planning), tactical (scheduling), and operational (process parameter) levels have shown promise in managing the complexity of large-scale manufacturing systems (Zhang et al., 2023).

Theme 4: Explainability and Trustworthiness. As AI-based controllers make increasingly autonomous decisions affecting energy consumption, product quality, and equipment safety, the need for explainable AI (XAI) in manufacturing control has become a pressing concern. This theme covers studies that develop interpretable control policies, attention-based explanation mechanisms, and counterfactual analysis tools that help manufacturing engineers understand why the AI controller makes specific parameter adjustments. Trust calibration frameworks that progressively expand the autonomous authority of AI controllers as their reliability is demonstrated have also been proposed (Meister et al., 2024).

Theme 5: Robustness and Safety-Aware Adaptive Control. Manufacturing processes involve physical systems where control errors can result in equipment damage, product defects, or safety hazards. This theme addresses the development of constrained reinforcement learning algorithms, safe exploration strategies, and Lyapunov-based stability guarantees for AI-based adaptive controllers. Barrier function approaches and control Lyapunov-barrier function (CLBF) methods have been adapted from robotics to

provide formal safety certificates for learned manufacturing controllers, ensuring that energy optimization does not compromise process safety or product quality specifications (Cheng et al., 2023).

Theme 6: Human-AI Collaborative Adaptive Control. This theme recognizes that manufacturing control is not a purely technical problem but involves human operators whose expertise, preferences, and trust must be integrated into the control loop. Studies in this theme develop mixed-initiative control architectures where AI systems provide recommendations that human operators can accept, modify, or override, with the AI learning from these human interventions to improve its future recommendations. Shared control paradigms, where the degree of AI autonomy varies dynamically based on process criticality and operator workload, have been explored as a means of balancing automation benefits with human oversight requirements (Rožanec et al., 2023).

5. The AIEEM Conceptual Framework

5.1. Framework Overview and Design Principles

Building on the systematic review findings, this section presents the Adaptive Intelligence for Energy-Efficient Manufacturing (AIEEM) conceptual framework. The AIEEM framework provides a structured, layered architecture for the design, implementation, and evaluation of AI-based adaptive control systems targeting energy efficiency optimization in smart manufacturing environments. The framework was developed following an iterative synthesis process that combined bottom-up pattern recognition from the reviewed literature with top-down design principles drawn from control systems engineering, software architecture, and sociotechnical systems theory.

The AIEEM framework is grounded in five design principles that emerged from the synthesis of best practices across the reviewed studies:

Principle 1 – Hierarchical Decomposition: Complex manufacturing energy optimization problems should be decomposed into hierarchical levels (strategic, tactical, operational) with appropriate AI techniques and temporal granularities at each level.

Principle 2 – Continuous Learning: AI-based controllers should incorporate mechanisms for continuous learning from operational data, enabling adaptation to changing process conditions, material properties, equipment degradation, and production requirements without manual retraining.

Principle 3 – Safety by Design: Energy optimization should never compromise process safety, product quality, or equipment integrity. Safety constraints should be embedded in the control architecture through constrained optimization, formal verification, or safety-aware reward shaping.

Principle 4 – Human-Centricity: The framework should accommodate human expertise, preferences, and oversight through interpretable AI models, adjustable autonomy levels, and effective human-machine interfaces that support operator understanding and trust.

Principle 5 – Scalability and Modularity: The architecture should be modular, allowing components to be added, replaced, or upgraded independently, and scalable from single-machine applications to multi-process production lines and cross-facility

deployments.

5.2. Framework Architecture: The Five Layers

The AIEEM framework consists of five interconnected layers, each addressing a distinct functional aspect of the AI-based adaptive control system. These layers, from bottom to top, are: (1) the Physical Sensing and Data Acquisition Layer, (2) the Feature Extraction and State Representation Layer, (3) the AI-Driven Decision and Optimization Layer, (4) the Adaptive Control Execution Layer, and (5) the Performance Monitoring and Continuous Learning Layer.

Layer 1 – Physical Sensing and Data Acquisition. This foundational layer encompasses the sensor infrastructure, data communication protocols, and edge computing resources required to capture the physical state of the manufacturing process with sufficient resolution, accuracy, and timeliness for effective adaptive control. The layer includes power meters and energy analyzers for direct energy measurement; process sensors (temperature, pressure, force, vibration, acoustic emission) for capturing process dynamics; environmental sensors (ambient temperature, humidity) for contextual awareness; and quality measurement systems (vision, dimensional, surface roughness) for closed-loop quality feedback. Data from these heterogeneous sources are synchronized, pre-processed, and aggregated at the edge level to reduce latency and bandwidth requirements for upstream layers. The layer also interfaces with the digital twin, providing real-time data streams for twin synchronization and receiving simulated sensor data for controller pre-training.

Layer 2 – Feature Extraction and State Representation. This layer transforms raw sensor data into meaningful state representations that capture the essential dynamics of the manufacturing process relevant to energy optimization. The layer employs both physics-informed and data-driven feature extraction methods. Physics-informed features include specific energy consumption (SEC), energy intensity ratios, and thermodynamic efficiency indicators derived from domain knowledge. Data-driven features are extracted using autoencoders, variational autoencoders (VAEs), or temporal convolutional networks that learn compact representations of the high-dimensional sensor data. The state representation must capture not only the current process state but also relevant historical information (through recurrent or attention-based architectures) and anticipated future states (through predictive models or digital twin projections). Effective state representation is critical for the performance of the upstream AI decision layer, as it determines what information is available for control decisions.

Layer 3 – AI-Driven Decision and Optimization. This central layer contains the core AI algorithms responsible for generating energy-optimal control actions. The layer supports multiple AI paradigms, including model-based approaches (where a neural network process model is combined with gradient-based or sampling-based optimization), model-free reinforcement learning (where the controller learns directly from interaction with the process or its digital twin), and hybrid approaches that combine model-based prediction with model-free policy learning. The layer implements multi-objective optimization to balance energy efficiency against competing objectives, with mechanisms for dynamic preference adjustment by operators. Safety constraints are enforced through constrained optimization techniques, including Lagrangian relaxation, barrier functions, and action masking. The layer also includes uncertainty quantification

modules that provide confidence estimates for control decisions, enabling risk-aware decision-making and appropriate escalation to human operators when confidence is low.

Layer 4 – Adaptive Control Execution. This layer translates the high-level control decisions from Layer 3 into low-level actuator commands compatible with the manufacturing equipment's control interfaces. The layer handles the critical transition from AI-generated setpoints to physical actuation, accounting for actuator dynamics, rate limits, communication latencies, and controller-level safety interlocks. The layer implements various control execution strategies, including direct setpoint modification (where the AI directly specifies process parameters), reference modification (where the AI adjusts the reference signals of existing PID or model predictive controllers), and gain scheduling (where the AI adjusts the parameters of existing controllers based on operating conditions). A key function of this layer is the implementation of graceful degradation strategies that ensure safe process operation in the event of AI system failures, sensor malfunctions, or communication interruptions, reverting to conventional control as a fallback.

Layer 5 – Performance Monitoring and Continuous Learning. This uppermost layer closes the adaptive control loop by monitoring system performance, evaluating the effectiveness of control decisions, and feeding performance data back into the AI decision layer for continuous improvement. The layer implements key performance indicators (KPIs) for energy efficiency (kWh per unit produced, energy cost per part, peak power reduction), production performance (throughput, cycle time, utilization), quality metrics (defect rate, dimensional accuracy, surface quality), and sustainability indicators (carbon footprint, energy from renewable sources). The layer also includes mechanisms for detecting concept drift (changes in the statistical properties of the process that may degrade controller performance), triggering retraining or model adaptation when drift is detected. Long-term learning capabilities include cross-batch learning (improving performance over successive production runs), cross-machine learning (transferring knowledge between similar machines), and cross-facility learning (aggregating insights across multiple manufacturing sites).

5.3. Framework Validation and Implementation Considerations

The AIEEM framework was validated through expert consultation with a panel of eight domain experts comprising four academic researchers in manufacturing AI and four industrial practitioners from automotive, aerospace, and electronics manufacturing companies. The expert panel confirmed the completeness of the framework layers, the appropriateness of the design principles, and the practical relevance of the framework components. Experts emphasized the importance of the safety-by-design principle and the need for standardized interfaces between framework layers to facilitate plug-and-play integration of different AI algorithms and manufacturing equipment.

Implementation of the AIEEM framework requires careful consideration of several practical factors. Data infrastructure must support the high-frequency, low-latency data flows required for real-time adaptive control, typically necessitating edge computing capabilities with processing latencies below 10 milliseconds for time-critical control loops. The selection of AI algorithms should consider not only their optimization performance but also their computational requirements, training data demands,

interpretability, and compatibility with existing industrial communication standards such as OPC UA, MQTT, and ROS 2.0. Organizational readiness factors, including workforce skills, management support, and change management processes, are equally critical for successful implementation and should not be underestimated (Rožanec et al., 2023).

Table 3. AIEEM Framework: Layer Specifications and Key Technologies

Layer	Function	Key Technologies	Latency Req.
L1: Sensing	Data capture and edge aggregation	IIoT, OPC UA, MQTT, edge computing	< 1 ms
L2: Features	State encoding and representation	Autoencoders, VAE, TCN, physics-informed	< 5 ms
L3: Decision	Multi-objective optimization	DRL, MARL, MPC-NN, constrained RL	< 10 ms
L4: Execution	Setpoint translation and actuation	PLC integration, ROS, safety PLC	< 1 ms
L5: Learning	KPI monitoring and model update	Drift detection, federated learning	Minutes–hours

6. Discussion

6.1. Theoretical Implications

This study makes several contributions to the theoretical understanding of AI-based adaptive control for energy-efficient manufacturing. First, the systematic review provides the most comprehensive mapping to date of the research landscape at the intersection of AI, adaptive control, and energy efficiency in smart manufacturing. The identification of six principal research themes and the quantitative synthesis of performance outcomes across 187 studies offer a structured foundation for future research planning and priority setting.

Second, the AIEEM conceptual framework advances the theoretical discourse by providing a unified, layered architecture that integrates insights from traditionally separate disciplines: control engineering, machine learning, industrial engineering, and sustainability science. By delineating the functional responsibilities, information flows, and technology requirements of each layer, the framework enables more precise theoretical analysis of where different AI techniques are most applicable and what performance trade-offs exist at each level of the control hierarchy.

Third, the finding that DRL approaches consistently outperform other AI techniques for multi-objective energy optimization in manufacturing has important implications for the direction of future research. While DNN-based and neuro-fuzzy approaches remain

valuable for process modeling and feature extraction, the sequential decision-making capabilities of DRL appear to offer a fundamental advantage for the inherently temporal nature of manufacturing process control. This finding aligns with and extends the theoretical arguments of Sutton and Barto (2018) regarding the suitability of RL for closed-loop control problems.

Fourth, the significant performance advantage of digital twin-integrated approaches (25.7% vs. 18.4% mean energy reduction) provides empirical support for the theoretical proposition that simulation-based pre-training can substantially improve the data efficiency and deployment speed of AI-based controllers. This finding contributes to the growing theoretical framework around sim-to-real transfer in cyber-physical systems and suggests that investment in digital twin infrastructure yields tangible returns in terms of adaptive control performance.

6.2. Practical Implications

For manufacturing practitioners and managers, the findings of this study carry several actionable implications. The quantified energy savings (median 21.4%, range 12-38%) demonstrate that AI-based adaptive control represents a viable and high-impact strategy for reducing manufacturing energy costs and carbon emissions. Given that energy costs typically represent 10-30% of total manufacturing operating costs, energy savings of this magnitude can translate into significant competitive advantages and contribute meaningfully to corporate sustainability targets.

The AIEEM framework provides practitioners with a structured roadmap for implementing AI-based adaptive control systems. By decomposing the system into five distinct layers with clearly defined interfaces, the framework enables incremental implementation strategies where organizations can begin with the sensing and feature extraction layers (building their data infrastructure) and progressively add AI decision-making and continuous learning capabilities as their technical maturity and organizational readiness develop.

The finding that digital twin integration substantially enhances AI controller performance suggests that manufacturers considering AI-based adaptive control should concurrently invest in digital twin capabilities. While the upfront investment in digital twin development can be significant, the reduced need for physical experimentation during controller development and the ability to pre-train AI agents in simulation can substantially reduce deployment timelines and risks, providing a favorable return on investment.

Furthermore, the importance of the human-AI collaboration theme highlighted in the review underscores that successful implementation of AI-based adaptive control is not purely a technical challenge. Organizations must invest in workforce development, creating new roles for AI-manufacturing engineers and providing existing operators with the skills and tools to effectively interact with and supervise AI-based control systems. Change management strategies that build trust through transparency, graduated autonomy, and demonstrated performance are essential for achieving organizational acceptance.

6.3. Limitations and Future Research Directions

Several limitations of this study should be acknowledged. First, the systematic review was limited to English-language publications, potentially excluding relevant research published in other languages, particularly Chinese, German, and Japanese, which are significant languages in manufacturing research. Second, the quantitative synthesis of energy performance metrics is complicated by the heterogeneity of manufacturing processes, baseline comparisons, and measurement methodologies across studies, limiting the generalizability of specific percentage values. Third, the AIEEM framework, while validated through expert consultation, has not yet been empirically tested through full-scale industrial implementation, and its practical applicability may be constrained by factors not anticipated in the conceptual design.

Based on the identified research gaps and limitations, the following future research directions are proposed:

Direction 1 – Scalable Multi-Process Optimization. Current research predominantly addresses single-machine or single-process energy optimization. Future research should develop scalable AI architectures for optimizing energy consumption across entire production lines and factory systems, accounting for inter-process dependencies, shared utilities, and system-level constraints. Hierarchical and federated reinforcement learning approaches offer promising avenues for addressing this scalability challenge.

Direction 2 – Standardized Benchmarking. The field urgently needs standardized benchmark problems, datasets, and evaluation protocols that enable fair comparison of different AI-based adaptive control approaches across research groups. The development of open-source simulation environments that accurately represent manufacturing energy dynamics, analogous to the OpenAI Gym environments used in the RL community, would significantly accelerate progress by enabling reproducible research and algorithm comparison.

Direction 3 – Lifecycle Energy Optimization. Most current approaches optimize energy consumption during steady-state production. Future research should extend adaptive control to encompass the full operational lifecycle, including machine startup and shutdown sequences, warm-up periods, idle states, maintenance operations, and product changeovers, which can account for a substantial fraction of total manufacturing energy consumption.

Direction 4 – Integration with Renewable Energy and Grid Interaction. As manufacturing facilities increasingly incorporate on-site renewable energy generation and participate in demand response programs, AI-based adaptive controllers should be extended to consider energy source composition, electricity pricing signals, and grid stability requirements in their optimization objectives, enabling manufacturing processes to serve as flexible loads that support the integration of variable renewable energy sources.

Direction 5 – Regulatory and Ethical Dimensions. As AI-based controllers assume greater autonomy in manufacturing environments, questions of regulatory compliance, liability, and ethical AI governance become increasingly relevant. Future research should address the development of regulatory frameworks for autonomous manufacturing control, the allocation of responsibility when AI decisions lead to adverse outcomes, and the ethical implications of AI-driven labor displacement in manufacturing.

7. Conclusions

This paper has presented a comprehensive systematic review and conceptual framework addressing the intersection of AI-based adaptive control and energy-efficient process optimization in smart manufacturing environments. Through the analysis of 187 peer-reviewed publications spanning from 2018 to 2025, the study has mapped the research landscape, identified dominant AI techniques and their comparative effectiveness, documented quantitative energy performance outcomes, and synthesized six principal research themes that characterize the current state of knowledge.

The key findings of this research can be summarized as follows. AI-based adaptive control systems demonstrate substantial potential for manufacturing energy optimization, achieving median energy consumption reductions of 21.4% compared to conventional control approaches, with deep reinforcement learning methods achieving the highest performance gains. The integration of digital twin technology significantly enhances the effectiveness of AI-based controllers, both by providing high-fidelity simulation environments for controller pre-training and by enabling predictive energy management capabilities. Multi-objective optimization approaches that simultaneously balance energy efficiency, production quality, and throughput represent the most practically relevant AI control architectures, as single-objective energy minimization typically leads to unacceptable trade-offs in other performance dimensions.

The AIEEM conceptual framework proposed in this study provides a structured, five-layer architecture for implementing AI-based adaptive control systems in smart manufacturing environments. Grounded in the principles of hierarchical decomposition, continuous learning, safety by design, human-centricity, and scalable modularity, the framework offers both a theoretical lens for understanding the complex interactions between AI, control, and manufacturing systems, and a practical roadmap for guiding industrial implementation efforts.

Looking forward, the convergence of increasingly capable AI algorithms, ubiquitous sensing and connectivity, and growing urgency for manufacturing decarbonization creates a compelling trajectory for the continued development and adoption of AI-based adaptive control for energy-efficient manufacturing. However, realizing this potential will require addressing significant challenges related to scalability, standardization, human-AI collaboration, and regulatory governance. The research agenda outlined in this paper provides a structured pathway for addressing these challenges and advancing both the scientific understanding and practical implementation of intelligent, energy-efficient manufacturing systems.

References

- Åström, K. J., & Wittenmark, B. (2013). *Adaptive control* (2nd ed.). Dover Publications.
- Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101. <https://doi.org/10.1191/1478088706qp063oa>
- Cheng, R., Orosz, G., Murray, R. M., & Burdick, J. W. (2023). End-to-end safe reinforcement learning through barrier functions for safety-critical continuous control tasks. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(6), 6678–6686.
- Dornheim, J., Link, N., & Gumbsch, P. (2020). Model-free adaptive optimal control of episodic fixed-horizon manufacturing processes using reinforcement learning. *International Journal of Control, Automation and Systems*, 18(6), 1593–1604.
- Dufloy, J. R., Sutherland, J. W., Dornfeld, D., Herrmann, C., Jeswiet, J., Kara, S., Hauschild, M., & Kellens, K.

- (2012). Towards energy and resource efficient manufacturing: A processes and systems approach. *CIRP Annals*, 61(2), 587–609.
- Finn, C., Abbeel, P., & Levine, S. (2017). Model-agnostic meta-learning for fast adaptation of deep networks. *Proceedings of the 34th International Conference on Machine Learning (ICML)*, 70, 1126–1135.
- Gao, R. X., Wang, L., Helu, M., & Teti, R. (2020). Big data analytics for smart factories of the future. *CIRP Annals*, 69(2), 668–692.
- Grieves, M., & Vickers, J. (2017). Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems. In F.-J. Kahlen, S. Flumerfelt, & A. Alves (Eds.), *Transdisciplinary perspectives on complex systems* (pp. 85–113). Springer.
- Herrmann, C., Thiede, S., Kara, S., & Hesselbach, J. (2011). Energy oriented simulation of manufacturing systems—Concept and application. *CIRP Annals*, 60(1), 45–48.
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780.
- Hou, Z. S., & Wang, Z. (2013). From model-based control to data-driven control: Survey, classification and perspective. *Information Sciences*, 235, 3–35.
- International Energy Agency. (2024). *World energy outlook 2024*. IEA Publications.
- Kagermann, H., Wahlster, W., & Helbig, J. (2013). Recommendations for implementing the strategic initiative Industrie 4.0. *Acatech—National Academy of Science and Engineering*.
- Kitchenham, B., & Charters, S. (2007). Guidelines for performing systematic literature reviews in software engineering (EBSE Technical Report 2007-01). Keele University.
- Koren, Y. (1983). *Computer control of manufacturing systems*. McGraw-Hill.
- Kusiak, A. (2018). Smart manufacturing. *International Journal of Production Research*, 56(1–2), 508–517.
- Landau, I. D., Lozano, R., M'Saad, M., & Karimi, A. (2011). *Adaptive control: Algorithms, analysis and applications* (2nd ed.). Springer.
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.
- Leitão, P. (2009). Agent-based distributed manufacturing control: A state-of-the-art survey. *Engineering Applications of Artificial Intelligence*, 22(7), 979–991.
- Li, X., Zhang, W., & Gao, L. (2023). Preference-conditioned multi-objective reinforcement learning for energy-aware manufacturing scheduling. *Journal of Manufacturing Systems*, 68, 234–249.
- Liang, S. Y., Hecker, R. L., & Landers, R. G. (2004). Machining process monitoring and control: The state-of-the-art. *Journal of Manufacturing Science and Engineering*, 126(2), 297–310.
- Lu, Y., Liu, C., Wang, K. I. K., Huang, H., & Xu, X. (2020). Digital twin-driven smart manufacturing: Connotation, reference model, applications and research issues. *Robotics and Computer-Integrated Manufacturing*, 61, 101837.
- Meister, S., Wermes, M., Stübs, J., & Anderl, R. (2024). Explainable AI for trustworthy manufacturing process control: A systematic review. *Journal of Intelligent Manufacturing*, 35(1), 147–168.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., ... & Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529–533.
- Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., ... & Moher, D. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ*, 372, n71.
- Pan, S. J., & Yang, Q. (2010). A survey on transfer learning. *IEEE Transactions on Knowledge and Data*

Engineering, 22(10), 1345–1359.

Panzer, M., & Bender, B. (2022). Deep reinforcement learning in production systems: A systematic literature review. *International Journal of Production Research*, 60(13), 4316–4341.

Qi, Q., & Tao, F. (2018). Digital twin and big data towards smart manufacturing and Industry 4.0: 360 degree comparison. *IEEE Access*, 6, 3585–3593.

Qi, Q., Tao, F., Hu, T., Anwer, N., Liu, A., Wei, Y., Wang, L., & Nee, A. Y. C. (2021). Enabling technologies and tools for digital twin. *Journal of Manufacturing Systems*, 58(B), 3–21.

Rožanec, J. M., Novalija, I., Zajec, P., Kenda, K., Tavakoli Ghinani, H., Suh, S., ... & Mladenčić, D. (2023). Human-centric artificial intelligence architecture for Industry 5.0 applications. *International Journal of Production Research*, 61(20), 6847–6872.

Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., ... & Hassabis, D. (2017). Mastering the game of Go without human knowledge. *Nature*, 550(7676), 354–359.

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

Tao, F., Cheng, J., Qi, Q., Zhang, M., Zhang, H., & Sui, F. (2018). Digital twin-driven product design, manufacturing and service with big data. *International Journal of Advanced Manufacturing Technology*, 94(9), 3563–3576.

Tao, F., Qi, Q., Wang, L., & Nee, A. Y. C. (2019). Digital twins and cyber-physical systems toward smart manufacturing and Industry 4.0: Correlation and comparison. *Engineering*, 5(4), 653–661.

Tranfield, D., Denyer, D., & Smart, P. (2003). Towards a methodology for developing evidence-informed management knowledge by means of systematic review. *British Journal of Management*, 14(3), 207–222.

van Eck, N. J., & Waltman, L. (2010). Software survey: VOSviewer, a computer program for bibliometric mapping. *Scientometrics*, 84(2), 523–538.

Wang, J., Ma, Y., Zhang, L., Gao, R. X., & Wu, D. (2018). Deep learning for smart manufacturing: Methods and applications. *Journal of Manufacturing Systems*, 48(C), 144–156.

Wang, L., Liu, Q., Chen, R., & Gao, R. X. (2024). Sim-to-real transfer of digital twin-trained deep reinforcement learning controllers for manufacturing processes. *Robotics and Computer-Integrated Manufacturing*, 85, 102645.

Zhang, J., Wang, P., Gao, R. X., & Guo, H. (2023). Hierarchical multi-agent reinforcement learning for energy-efficient manufacturing system control. *IEEE Transactions on Industrial Informatics*, 19(4), 5321–5333.

Zheng, Y., Yang, S., & Cheng, H. (2021). An application framework of digital twin and its case study. *Journal of Ambient Intelligence and Humanized Computing*, 12(3), 3209–3221.

Zhong, R. Y., Xu, X., Klotz, E., & Newman, S. T. (2017). Intelligent manufacturing in the context of Industry 4.0: A review. *Engineering*, 3(5), 616–630.

Zhou, J., Cui, G., Hu, S., Zhang, Z., Yang, C., Liu, Z., ... & Sun, M. (2020). Graph neural networks: A review of methods and applications. *AI Open*, 1, 57–81.