

Multi-Objective Evolutionary Optimization of Interpretable Rules for Smart Manufacturing Process Control

Kevin L. Fernandez

Department of Computer Science, University of Alabama at Birmingham, Birmingham, AL, USA.

kevinf@uab.edu

Isaac Woods

Department of Computer Science, University of New Hampshire, Durham, NH, USA.

contactisaac@unh.edu

Rowan Phellips

School of Information Technology, University of Cincinnati, Cincinnati, OH, USA.

rphillips@uc.edu

Abstract

The increasing complexity of modern manufacturing systems demands process control strategies that simultaneously reconcile multiple conflicting objectives while maintaining transparency for human decision-makers. This paper presents a comprehensive examination of multi-objective evolutionary optimization for the generation of interpretable rule sets in smart manufacturing process control. We articulate a systems-level framework that integrates concepts from multi-objective evolutionary computation, fuzzy rule-based systems, and explainable artificial intelligence to produce control policies that are both efficient and auditable. The discussion foregrounds the structural trade-offs among accuracy, interpretability, fairness, robustness, and sustainability, and delineates the architectural choices that influence the deployability of such models in real-world industrial infrastructures. By analyzing the governance implications of embedding evolutionary rule discovery into production environments, we highlight how algorithmic transparency can address regulatory and ethical requirements while enabling continuous adaptation to shifting operational conditions. The paper further explores the socio-technical dimensions of the approach, including the allocation of decision authority between autonomous controllers and human operators, the role of standards in ensuring interoperability, and the lifecycle concerns that arise when learned rules must be maintained across equipment generations. In doing so, we offer a forward-looking perspective on the design of intelligent, accountable, and sustainable manufacturing control systems that are capable of evolving in concert with their underlying production environments.

Keywords

multi-objective optimization, evolutionary algorithms, interpretable rules, smart manufacturing, process control, explainable AI, governance.

1. Introduction

The transformation of manufacturing under the banner of Industry 4.0 has been driven by the convergence of cyber-physical systems, ubiquitous sensing, and data-driven decision making. In this landscape, process control functions are increasingly performed by algorithmic agents that must satisfy a diverse array of performance criteria. Traditional single-objective optimization methods, which reduce all preferences to a scalarized cost function, often obscure the inherent tensions among competing goals such as productivity, energy efficiency, product quality, and equipment longevity. Multi-objective evolutionary algorithms have emerged as a compelling alternative paradigm, able to approximate entire Pareto-optimal fronts and thus make explicit the trade-offs that remain hidden in scalarized formulations [1, 2]. While these algorithms were originally demonstrated on benchmark problems, their potential to reimagine manufacturing control is now being recognized, especially when the learned control policies are expressed in symbolic, human-readable forms that facilitate operator understanding and regulatory inspection [3].

The demand for interpretability in manufacturing process control is not merely an academic curiosity; it arises from practical imperatives rooted in safety, accountability, and continuous improvement. In critical production environments, operators must trust that an automated controller will not initiate actions that endanger the process, violate environmental regulations, or induce unforeseen wear on capital equipment. Black-box machine learning models, regardless of their predictive accuracy, introduce a layer of opacity that can undermine this trust and complicate root-cause analysis when anomalies occur. Recent scholarship has forcefully argued that high-stakes decision contexts require inherently interpretable models rather than post-hoc explanations of opaque systems [4]. In manufacturing, this translates into a preference for rule-based controllers whose logic can be read, challenged, and incrementally refined by domain experts. The intersection of multi-objective optimization and interpretable rule induction therefore represents a fertile meeting ground between algorithmic rigor and industrial pragmatism.

Smart manufacturing environments are increasingly characterized by the presence of digital twins that mirror physical assets in real time, enabling a continuous loop of simulation, optimization, and deployment [5]. The data streams generated by such twins provide rich training signals for evolutionary rule discovery, yet they also introduce new system-level challenges. Optimized rule sets must not only perform well on historical data but must also generalize across varying raw material properties, ambient conditions, and production schedules. Moreover, the dynamic nature of manufacturing ecosystems implies that the objectives themselves may shift in response to market pressures, energy tariffs, or environmental reporting mandates, requiring the optimization framework to be inherently reconfigurable. Early surveys of machine learning in manufacturing have documented both the promise and the fragmentation of current approaches, noting a gap between laboratory demonstrations and production-hardened solutions [6]. Addressing this gap calls for a holistic design philosophy that bridges optimization theory, software architecture, and organizational governance.

From a systems perspective, the challenge of deploying evolutionary optimization for interpretable process control can be decomposed into several intersecting concerns. The first is the design of the rule representation and the corresponding search operators: should rules be crisp, fuzzy, or probabilistic, and how should their structure be constrained to ensure readability without sacrificing expressiveness? The second is the selection of a multi-objective algorithm and the articulation of objectives that go beyond simple error metrics to

encompass rule complexity, fairness across product lines, and robustness to sensor noise. The third is the integration of the optimization pipeline into the broader manufacturing information technology infrastructure, including edge computing nodes, supervisory control and data acquisition systems, and manufacturing execution systems. Each of these dimensions will be explored in the sections that follow, with an emphasis on structural trade-offs and long-term sustainability rather than on narrow algorithmic improvements.

2. Background and Foundational Concepts

The intellectual heritage of the present work draws from several distinct yet interlocking research communities. Multi-objective evolutionary optimization has matured into a well-established discipline, with landmark algorithms such as NSGA-II and SPEA2 providing robust baselines for approximating Pareto sets in continuous and combinatorial domains [1, 2]. These algorithms operate by maintaining a population of candidate solutions that are evolved through variation and selection mechanisms targeting both convergence toward the true Pareto front and diversity along it. Their application to the discovery of rule-based systems introduces the additional challenge that the genotype-to-phenotype mapping must be carefully crafted so that the resulting rules remain linguistically coherent and cognitively manageable for human users [3]. In parallel, the fuzzy systems community has developed a rich taxonomy of interpretability criteria, distinguishing between low-level readability of individual rules and high-level comprehensibility of the entire rule base as a knowledge structure [9, 10]. The interplay between these two bodies of work provides the conceptual foundation on which a multi-objective evolutionary framework for interpretable manufacturing control can be constructed.

Industrial interest in smart manufacturing has accelerated the adoption of data-driven techniques for process monitoring, diagnosis, and control. Foundational reviews in this area have highlighted the importance of transforming vast operational datasets into actionable knowledge while preserving the contextual understanding that experienced operators bring to production environments [6]. The concept of smart manufacturing itself has been elaborated as a socio-technical paradigm in which distributed intelligence, human-machine collaboration, and lifecycle-wide data integration are key enablers of agility and resilience [7, 8]. Within this paradigm, the notion of a rule-based controller takes on a renewed significance: it becomes a boundary object that mediates between the algorithmic and the human, encoding domain expertise in a format that can be jointly developed by data scientists, process engineers, and shop-floor personnel. Early evolutionary computation applications to job-shop scheduling demonstrated that genetic representations could capture manufacturing constraints and produce rule-like dispatching heuristics, although these early efforts often prioritized a single objective such as makespan or tardiness [11].

The interpretability imperative has been amplified by a broader societal turn toward accountable artificial intelligence. Surveys of bias and fairness in machine learning have documented the ways in which data-driven models can inadvertently encode and perpetuate inequities, a concern that extends to manufacturing whenever control policies differentiate among product types, production lines, or geographic facilities [12]. While the discourse on fairness has been most prominent in consumer-facing domains, its relevance to industrial settings is growing as automated quality grading, predictive maintenance scheduling, and resource allocation decisions come under scrutiny. Robust optimization, in turn, provides a complementary lens through which to evaluate the stability of control policies under uncertainty, an essential property when process measurements are subject to sensor drift and

environmental perturbations [13]. The aspiration to simultaneously optimize fairness and robustness alongside traditional performance metrics exemplifies the multi-objective nature of the problem and underscores the inadequacy of single-criterion approaches.

Sustainability considerations have increasingly become a first-class design objective in manufacturing systems engineering, driven by both regulatory pressure and corporate environmental commitments [14]. When process control rules are automatically generated from operational data, it is essential that the optimization objectives explicitly encode energy consumption, material waste, and emission footprints. Evolutionary multi-objective optimization is particularly well suited to this task, as it can reveal the cost of achieving sustainability improvements in terms of throughput or quality degradation, thereby enabling informed trade-off decisions by plant management. The challenge is compounded by the fact that sustainability metrics often interact nonlinearly with production parameters and may exhibit temporal dynamics that are not captured in steady-state data snapshots. Designing an optimization framework that can internalize these dynamics while producing interpretable rules remains an open research frontier.

The relationship between accuracy and interpretability has been a perennial theme in fuzzy modeling and has motivated a family of multi-objective genetic fuzzy systems that explicitly treat these two properties as competing objectives [10, 15]. These approaches have shown that modest concessions in numerical precision can yield dramatic gains in the transparency of the resulting rule bases, and that interactive visualization tools can help domain experts navigate the resulting Pareto fronts. Extending this philosophy to manufacturing process control requires embedding the rule learning process within the operational constraints of industrial software systems, a step that raises additional questions about computational efficiency, memory footprint, and compatibility with programmable logic controllers. The lineage of model predictive control offers instructive parallels, as its long-standing industrial deployment has necessitated careful handling of real-time computation, model maintenance, and operator interaction [16]. Evolutionary algorithms have likewise been surveyed for their utility in control systems engineering, highlighting both their flexibility in handling nonlinear objectives and the practical difficulties of certifying their solutions for safety-critical applications [17].

3. System Architecture for Multi-Objective Rule Optimization in Manufacturing

Designing a viable system architecture for evolutionary optimization of interpretable process control rules requires balancing the exploratory capabilities of population-based search with the deterministic reliability expected in industrial execution environments. At the highest level of abstraction, such an architecture comprises a data ingestion layer that interfaces with plant historians and digital twins, an offline optimization engine responsible for evolving rule populations, an evaluation module that simulates or tests candidate controllers against multiple objective functions, and a deployment pathway that translates selected rules into executable control logic. The data ingestion layer must accommodate the heterogeneity of manufacturing data sources, including time-series variables sampled at disparate frequencies, categorical batch identifiers, and alarm logs. Preprocessing pipelines that perform missing value imputation, outlier removal, and feature engineering must be transparent and auditable to avoid introducing systematic biases that could be magnified by the evolutionary process.

The core optimization engine is predicated on a carefully chosen rule representation. Fuzzy if-then rules with linguistic labels offer a compelling synthesis of interpretability and generality, as they map continuous sensor readings onto humanly meaningful categories such as “high

temperature” or “low vibration” and prescribe control actions that are themselves expressed in linguistic or crisp terms. The evolutionary operators—crossover, mutation, and selection—must be designed to respect the syntactic boundaries of the rule language so that offspring remain interpretable. Research in genetic fuzzy systems has catalogued an array of genetic representations, including Pittsburgh and Michigan approaches, each carrying different implications for population dynamics and convergence behavior [9]. For manufacturing process control, a Pittsburgh-style encoding, where each individual represents a complete rule base, tends to be preferred because it allows direct multi-objective evaluation of the controller as an integrated whole.

The definition of objective functions is arguably the most consequential design decision in the entire architecture. A naive formulation that minimizes root-mean-squared control error and maximizes throughput will frequently produce controllers that are overly aggressive, intolerably brittle, or energetically wasteful. A more mature framework incorporates auxiliary objectives that penalize rule base complexity, measured perhaps by the number of rules, the number of antecedent conditions per rule, or the degree of overlap among fuzzy partitions [10]. Additionally, robustness to input perturbations should be promoted by evaluating candidate controllers on multiple realizations of the digital twin that inject noise and parameter drift in accordance with observed plant variability. Sustainability can be folded into the objective vector through direct energy and material consumption metrics derived from the twin. In this way, the multi-objective optimization process yields a set of non-dominated trade-off solutions from which plant management can select according to prevailing priorities.

The integration of the optimization engine with digital twin infrastructure deserves particular attention. Digital twins enable what-if experimentation without risking physical assets, making them ideal platforms for fitness evaluation during evolutionary search. However, the fidelity and computational cost of the twin impose practical constraints: while high-fidelity computational fluid dynamics or finite element models can accurately predict process behavior, their execution time per evaluation may render population-based optimization infeasible for large search spaces. A layered evaluation strategy is therefore often advisable, in which candidate solutions are first screened using fast surrogate models or reduced-order twins before the most promising individuals are validated against full-fidelity simulations. This staged approach mirrors the hierarchical control structures found in modern manufacturing execution systems and aligns with the broader trend toward edge-cloud orchestration [21]. The rule discovery process itself can benefit from innovations in data mining that efficiently extract diversified rule sets from large databases, such as user-guided embedding techniques that allow operators to steer the search toward operationally meaningful regions of the rule space [18]. The incorporation of such techniques can accelerate convergence and improve the relevance of the discovered rules, while still preserving the multi-objective character of the overall framework.

Once a set of candidate controllers has been evolved, the selection of a single rule base for deployment is inherently a multi-criteria decision-making problem that must involve human stakeholders. Visualization tools that plot the Pareto front and allow interactive exploration of rule base contents have proven effective in fuzzy systems research and can be adapted to the manufacturing context [10]. Operators and process engineers can inspect the linguistic rules, compare them with their own mental models of the process, and assess the operational consequences of the trade-offs displayed. Such collaborative decision processes help build trust and surface tacit knowledge that may not be present in the historical data. After

deployment, a continuous learning loop should be established wherein the performance of the active controller is monitored against its predicted Pareto position, and significant deviations trigger a re-optimization cycle. This lifecycle perspective, which treats the rule base as a living artefact rather than a one-time deliverable, is essential for maintaining alignment between the controller and an evolving production environment.

4. Interpretability, Fairness, and Governance in Rule-Based Control

Interpretability in manufacturing process control is not merely a cosmetic property but a functional requirement that intersects with regulatory compliance, operator training, and incident investigation. When process anomalies occur, the ability to trace the control logic that contributed to an undesirable state can drastically reduce diagnostic time and prevent recurrence. Recent comprehensive surveys of explainable artificial intelligence have delineated a spectrum of interpretability methods, distinguishing post-hoc explanation approaches from inherently transparent model classes. Rule-based systems fall squarely into the latter category, offering a direct path from model internals to human-understandable reasoning chains [20]. However, interpretability is not a monolithic quality; it encompasses multiple dimensions such as simulatability, decomposability, and algorithmic transparency, each of which can be supported or undermined by specific architectural choices. For manufacturing process control, simulatability—the ability of a human to step through the inference process mentally—is particularly prized because it enables operators to anticipate controller behavior in novel situations without executing the actual process.

With interpretability comes the responsibility to scrutinize whether the rules encode unfair or discriminatory patterns. In manufacturing contexts, fairness concerns might arise when a controller assigns different quality thresholds or maintenance intervals to products destined for different markets, or when resource allocation algorithms inadvertently favor certain production lines due to historical biases in sensor placement or data collection practices. The fairness-aware machine learning literature provides conceptual frameworks that can be adapted to this domain, including notions of individual fairness, which demands that similar process states evoke similar control actions, and group fairness, which requires parity of outcomes across distinct product categories or shifts [12, 22]. Embedding fairness metrics as explicit objectives within the evolutionary optimization loop enables the discovery of rule bases that achieve a desirable balance between accuracy and equity. This approach is preferable to post-hoc fairness corrections that can obscure the underlying trade-offs and complicate subsequent maintenance.

Governance of algorithmic process control extends beyond technical performance to encompass the allocation of decision rights between human and automated agents. As manufacturers adopt increasingly sophisticated optimization tools, they must define clear policies regarding the scope of autonomous control actions and the conditions under which human intervention is required. Ethical frameworks developed for artificial intelligence in society provide transferable principles regarding beneficence, non-maleficence, autonomy, and justice that can inform the design of manufacturing control systems [23]. In practice, this may translate to constraints on the feasible region of the optimization that prevent controllers from exceeding safety envelopes or from making abrupt changes that could destabilize downstream operations. The rule-based representation supports governance by making such constraints explicit and auditable; a rule that violates a safety constraint can be identified and eliminated before deployment, in contrast to neural controllers whose internal representations defy straightforward inspection.

Interoperability and standardization are additional governance dimensions that influence the long-term viability of evolutionary rule optimization in manufacturing. The diversity of equipment vendors, communication protocols, and data formats in a typical plant creates a fragmented landscape in which a controller optimized for one machine may be incommunicable with another. National and international standards bodies have responded with reference architectures for smart manufacturing that define common information models and interface specifications [24]. Aligning the rule optimization pipeline with such standards ensures that the resulting controllers can be integrated into supervisory control platforms, shared across facilities, and maintained by personnel with standardized training. Moreover, standardizing the representation of rules and objectives facilitates benchmarking and comparative evaluation across different optimization approaches, accelerating the translation of research advances into industrial practice.

The sustainability lens adds a temporal dimension to governance, as manufacturing systems must be managed not just for immediate efficiency but for their long-term environmental and economic footprint. Circular economy principles advocate for closed-loop material flows and extended product lifecycles, which impose new requirements on process control rules [25]. A controller might, for instance, need to adjust parameters to accommodate recycled feedstock of variable quality or to minimize energy consumption during periods of high carbon intensity on the electric grid. By incorporating lifecycle assessment metrics into the multi-objective formulation, the evolutionary framework can generate rule bases that are responsive to such circular economy mandates. The interpretability of the rules then serves an additional purpose: it enables sustainability officers and external auditors to verify that environmental claims are substantiated by actual control logic rather than by aspirational reporting.

5. Deployment, Infrastructure, and Lifecycle Management

Transitioning a multi-objective evolutionary rule optimizer from a research prototype to a production-grade system demands careful attention to the computational infrastructure on which it will execute. In smart factories, the prevailing architectural paradigm involves a hierarchy of computing resources ranging from low-latency edge devices on the factory floor to cloud-based data lakes and analytics platforms [21]. The offline optimization and model selection stages are naturally suited to cloud execution, where abundant computational resources can be marshalled to explore large populations and evaluate candidates against high-fidelity digital twins. Once a rule base has been selected, the compiled controller—often a compact set of if-then statements—can be deployed to edge gateways or directly to programmable logic controllers with minimal computational overhead. This separation of concerns allows the system to reap the benefits of cloud-scale optimization while preserving the real-time responsiveness required for closed-loop process control.

The reliability of the deployed rule base must be assured through rigorous verification and validation procedures that go beyond aggregate performance statistics. Because the evolutionary search is inherently stochastic and its solutions are not guaranteed to be optimal, a disciplined testing regimen is necessary. This regimen typically includes stress testing against historical edge cases, adversarial perturbation analysis to gauge robustness, and controlled on-site trials conducted under human supervision before full autonomous operation is authorized. The incorporation of robust optimization principles into the objective vector reduces the likelihood of discovering brittle controllers, but it does not eliminate the need for empirical validation [13]. Furthermore, the interpretability of the rule base is a significant

asset during validation, as it allows domain experts to perform manual sanity checks by reading the rules and flagging those that contradict established engineering knowledge.

Lifecycle management of the control system must account for the fact that both the manufacturing process and its objectives evolve over time. Production campaigns may shift in response to market demand, new sensors may be installed that alter the feature space, and equipment aging may drift the process dynamics away from the conditions captured in the original digital twin. A static rule base will inevitably become miscalibrated under such drift, leading to suboptimal or unsafe behavior. An adaptive lifecycle strategy that periodically re-triggers evolutionary search, perhaps on a monthly or per-campaign basis, can maintain alignment between the controller and the real process. The retained interpretability of the rules enables incremental updates: rather than replacing the entire controller at each cycle, operators can choose to accept only those rule modifications that improve performance without sacrificing comprehensibility. Such an approach resonates with the concept of continuous commissioning in building energy management and can be translated to the manufacturing domain with appropriate adjustments.

The human dimension of deployment should not be underestimated. Even the most technically elegant rule-based controller will fail if operators distrust it, circumvent it, or lack the skills to interact with it productively. Organizational change management, targeted training, and participative design processes are therefore essential accompaniments to the technical rollout. When operators are involved in the rule selection process—viewing the Pareto front, inspecting candidate rule bases, and contributing their experiential knowledge—they are more likely to develop a sense of ownership and calibrated trust. This socio-technical integration is in keeping with the vision of smart manufacturing as a human-centered enterprise, in which technology amplifies rather than replaces human expertise [7].

From a broader policy perspective, the proliferation of self-optimizing manufacturing controllers raises questions about liability, certification, and cross-border data governance. If an evolved rule base causes a quality deviation that leads to a product recall, the allocation of responsibility among the technology vendor, the plant operator, and the algorithm developer may be contentious. Regulatory frameworks for industrial artificial intelligence are still nascent, but they will likely require a documented chain of design decisions, a clear specification of the objective trade-offs made, and evidence of fairness and robustness assessments. Interpretable rule-based systems, precisely because their logic can be audited, are better positioned to meet such requirements than opaque neural networks. Standardization efforts that define acceptable formats for rule base exchange and documentation will further facilitate regulatory compliance and insurance underwriting [24]. The cumulative effect of these considerations is that the design of multi-objective evolutionary optimizers for manufacturing process control is as much a governance and infrastructure problem as it is an algorithmic one.

6. Conclusion

The synthesis of multi-objective evolutionary optimization and interpretable rule induction offers a principled pathway toward intelligent, accountable, and sustainable manufacturing process control. This paper has articulated a systems-level perspective that views the optimization engine not as an isolated algorithm but as a component within a broader socio-technical infrastructure encompassing digital twins, edge-cloud architectures, human decision-makers, and regulatory frameworks. The structural trade-offs among accuracy, interpretability, robustness, fairness, and sustainability have been examined, and the

architectural choices that condition the deployability of rule-based controllers have been mapped. It has been argued that fuzzy rule-based representations, when evolved under a multi-objective fitness regime that explicitly penalizes complexity and promotes robustness, can yield controllers that are both performant and transparent, satisfying the dual imperatives of industrial efficiency and human oversight. The governance challenges accompanying autonomous process control—ranging from fairness across product lines to lifecycle management under environmental mandates—are rendered more tractable by the interpretability of the resulting control logic, which transforms the controller from a black box into an auditable artifact. As manufacturing enterprises navigate the transition toward increasingly autonomous operations, the design principles articulated here can guide the development of process control systems that are not only intelligent but also trustworthy, equitable, and aligned with long-term societal goals.

References

1. Deb, K., Pratap, A., Agarwal, S., & Meyarivan, T. (2002). A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2), 182–197.
2. Zitzler, E., Laumanns, M., & Thiele, L. (2001). SPEA2: Improving the strength Pareto evolutionary algorithm. In *Evolutionary Methods for Design, Optimization and Control with Applications to Industrial Problems* (pp. 95–100). CIMNE.
3. Ishibuchi, H., & Nojima, Y. (2013). Analysis of interpretability-accuracy tradeoff of fuzzy systems by multiobjective fuzzy genetics-based machine learning. *IEEE Transactions on Fuzzy Systems*, 21(4), 660–678.
4. Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), 206–215.
5. Tao, F., Zhang, M., Liu, Y., & Nee, A. Y. C. (2018). Digital twin driven prognostics and health management for complex equipment. *CIRP Annals*, 67(1), 169–172.
6. Wuest, T., Weimer, D., Irgens, C., & Thoben, K. D. (2016). Machine learning in manufacturing: advantages, challenges, and applications. *Production & Manufacturing Research*, 4(1), 23–45.
7. Wang, L. (2019). From intelligence science to intelligent manufacturing. *Engineering*, 5(4), 615–618.
8. Kusiak, A. (2018). Smart manufacturing. *International Journal of Production Research*, 56(1-2), 508–517.
9. Herrera, F. (2008). Genetic fuzzy systems: taxonomy, current research trends and prospects. *Evolutionary Intelligence*, 1(1), 27–46.
10. Gacto, M. J., Alcalá, R., & Herrera, F. (2011). Interpretability of linguistic fuzzy rule-based systems: An overview of interpretability measures. *Information Sciences*, 181(20), 4340–4360.
11. Cheng, R., Gen, M., & Tsujimura, Y. (1996). A tutorial survey of job-shop scheduling problems using genetic algorithms—I. Representation. *Computers & Industrial Engineering*, 30(4), 983–997.

12. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), 1–35.
13. Beyer, H. G., & Sendhoff, B. (2007). Robust optimization – A comprehensive survey. *Computer Methods in Applied Mechanics and Engineering*, 196(33-34), 3190–3218.
14. Bi, Z. M., & Wang, L. (2020). Manufacturing system design for sustainability: A review. *Robotics and Computer-Integrated Manufacturing*, 61, 101836.
15. Casillas, J., Cordon, O., Herrera, F., & Magdalena, L. (2003). Accuracy improvements in linguistic fuzzy modeling. In *Accuracy Improvements in Linguistic Fuzzy Modeling* (pp. 3–26). Springer.
16. Qin, S. J., & Badgwell, T. A. (2003). A survey of industrial model predictive control technology. *Control Engineering Practice*, 11(7), 733–764.
17. Fleming, P. J., & Purshouse, R. C. (2002). Evolutionary algorithms in control systems engineering: a survey. *Control Engineering Practice*, 10(11), 1223–1241.
18. Han, Z., Chen, W., Han, Y., Mao, R., & Qin, J. (2026). Fast Diversified Top-k Rule Discovery via User-Guided Embeddings. *IEEE Transactions on Knowledge and Data Engineering*.
19. Kamsu-Foguem, B., Rigal, F., & Mauget, F. (2013). Mining association rules for the quality improvement of the production process. *Expert Systems with Applications*, 40(4), 1034–1045.
20. Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., Garcia, S., Gil-Lopez, S., Molina, D., Benjamins, R., Chatila, R., & Herrera, F. (2020). Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58, 82–115.
21. Mao, Y., You, C., Zhang, J., Huang, K., & Letaief, K. B. (2017). A survey on mobile edge computing: The communication perspective. *IEEE Communications Surveys & Tutorials*, 19(4), 2322–2358.
22. Dwork, C., Hardt, M., Pitassi, T., Reingold, O., & Zemel, R. (2012). Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference* (pp. 214–226). ACM.
23. Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds and Machines*, 28(4), 689–707.
24. Lu, Y., Morris, K. C., & Frechette, S. (2016). Current standards landscape for smart manufacturing systems (NISTIR 8107). National Institute of Standards and Technology.
25. Lieder, M., & Rashid, A. (2016). Towards circular economy implementation: a comprehensive review in context of manufacturing industry. *Journal of Cleaner Production*, 115, 36–51.