

Runtime Policy Orchestration for Autonomous Industrial Control and Smart Manufacturing Systems: A Unified Framework for Governance, Compliance, and Adaptive Resilience

Abstract

The rapid evolution of Industry 4.0 and emerging Industry 5.0 paradigms has accelerated the adoption of autonomous cyber-physical systems (CPS) in smart manufacturing, creating a need for robust runtime policy orchestration to ensure governance, regulatory compliance, and resilience under dynamic disruptions and cyber threats. This study addresses the lack of unified frameworks by proposing the Policy Resilient Orchestrator (PRO), a simulation-driven architecture integrating machine-readable policy cards, multi-agent reinforcement learning using Proximal Policy Optimization (PPO), multi-objective optimization via NSGA-II, runtime verification, and digital twin-based validation. A quantitative, desk-based methodology utilizing synthetic and benchmark datasets was employed to model CPS environments without physical hardware. NSGA-II generated Pareto-optimal policy configurations, while PPO enabled adaptive runtime control supported by simplex switching for resilience. Monte Carlo simulations across normal, attack, and failure scenarios showed significant improvements, with Governance Efficiency (0.961), Resilience Index (0.993), and Downtime Reduction (0.988), all statistically outperforming baselines ($p < 0.0001$). PRO demonstrates a scalable, reliable solution for trustworthy smart manufacturing systems.

Keywords: Runtime Policy Orchestration, Cyber-Physical Systems, Smart Manufacturing, Adaptive Resilience, Digital Twins

1. Introduction

The advent of Industry 4.0 has transformed manufacturing through cyber-physical systems (CPS), Internet of Things (IoT), and artificial intelligence, enabling autonomous industrial control and smart factories (Huang et al., 2021). Smart manufacturing systems integrate physical processes with digital orchestration, allowing real-time decision-making and adaptive operations in dynamic

environments (Ryalat et al., 2024). However, as these systems gain autonomy, runtime policy orchestration emerges as critical for aligning operations with governance, compliance, and resilience requirements (Errico, 2026). Runtime policy orchestration involves dynamically managing policies at execution time to enforce rules, adapt to changes, and ensure system integrity in autonomous setups (Mavracic, 2025). In smart manufacturing, this orchestration governs resource allocation, security protocols, and control loops across distributed CPS (Huma, 2026). Traditional static policies fail in volatile industrial settings, necessitating unified frameworks that integrate policy enforcement with adaptive mechanisms (Kokkonen et al., 2022). The evolution from rigid control to policy-driven autonomy draws from reinforcement learning and multi-agent systems, enabling self-optimizing factories (Dastranj et al., 2022). Digital twins and edge computing further support runtime verification, simulating policy impacts before deployment (Djebbouri et al., 2025). This background underscores the need for a cohesive framework to orchestrate policies amid increasing cyber threats and regulatory demands in smart manufacturing (Ajayi et al., 2025).

Autonomous industrial control systems in smart manufacturing face significant challenges in runtime policy management, leading to vulnerabilities in governance, compliance, and resilience (Sheikhi et al., 2021). Current approaches rely on fragmented policies that cannot dynamically adapt to real-time disruptions such as cyberattacks, equipment failures, or regulatory shifts, resulting in downtime and non-compliance (Ashfaq et al., 2025). The lack of a unified orchestration mechanism exacerbates issues like policy conflicts, inconsistent enforcement across CPS, and inadequate resilience against evolving threats (Sampath & Baskaran, 2026). In dynamic smart factories, heterogeneous devices and multi-agent interactions amplify these problems, as static verification scales poorly with system complexity (Dastranj et al., 2021). Compliance with standards like NIST AI RMF or EU AI Act remains manual and error-prone, hindering scalable deployment. Moreover, without adaptive runtime steering, systems suffer from intent drift and suboptimal resource use, compromising operational efficiency (Shi et al., 2024). This gap is particularly acute in research contexts, where physical prototyping is infeasible, demanding simulation-driven frameworks for policy synthesis and evaluation (Morris et al., 2020). Existing solutions lack integration of policy learning, enforcement, and resilience, creating a critical void in unified runtime orchestration (Wang et al., 2016).

This study proposes a unified framework for runtime policy orchestration, addressing key gaps in smart manufacturing governance and resilience. By enabling dynamic policy adaptation, it enhances system reliability, reducing downtime in simulated CPS environments. The framework's policy cards and multi-agent orchestration ensure verifiable compliance, aligning with global standards and facilitating regulatory adherence (Mavracic, 2025). For researchers and practitioners, it provides a blueprint for simulation-based development (Kabir & Ray, 2025). In Industry 5.0 contexts, it promotes human-centric resilience, optimizing sustainability through adaptive resource orchestration (Djebbouri et al., 2025). Economically, it lowers implementation barriers for SMEs via modular, AI-driven policies (Çinar et al., 2021). Broadly, the framework advances CPS security, fostering trustworthy autonomous systems amid rising cyber-physical threats (Ajayi et al., 2025). Its contributions empower scalable smart factories, driving innovation in governance and adaptive control (Wang et al., 2016). The study focuses on conceptual and simulation-based development of the runtime policy orchestration framework for smart manufacturing CPS. It targets Industry 4.0/5.0 environments with autonomous control loops, emphasizing software-defined policies, AI orchestration, and digital twin integration (Huang et al., 2021). The scope encompasses governance through policy synthesis, compliance via standards mapping like NIST and EU AI Act, and resilience via threat adaptation, modeled via tools like NSGA-II for optimization. Validation uses synthetic datasets from literature benchmarks, suitable for home-based analysis. Limitations include assumption of reliable edge-cloud connectivity and exclusion of human-in-the-loop ergonomics (Kokkonen et al., 2022). Geographically neutral, it draws from global standards, prioritizing scalable solutions for SMEs in developing regions (Rafique et al., 2022).

The research centers on a workable, simulation-centric unified framework for runtime policy management in smart manufacturing. It integrates policy cards for machine-readable governance, deep reinforcement learning for adaptive orchestration, and black-box simplex for runtime assurance (Sheikhi et al., 2021). Leveraging digital twins, it enables remote validation of policy dynamics against compliance and resilience metrics (Djebbouri et al., 2025). The framework operates in three phases: policy ingestion, runtime orchestration, and resilience feedback, synthesized from surveyed approaches. This positions the study as a publishable contribution for experts seeking integrated, desk-based innovations in CPS governance (Wang et al., 2016). The evolution of manufacturing paradigms illustrates progression from Industry 3.0 static controls to

Industry 5.0 adaptive policies, with CPS, IoT, and AI integration marking key milestones in this shift. Autonomous systems now require runtime mechanisms to handle volatility, where policy conflicts and cyber threats can cause significant operational failures if unaddressed. Adaptive mechanisms using multi-agent reinforcement learning further enable self-correction during execution, while digital twins provide predictive insights for threat mitigation. Overall, these elements collectively bridge the gap between theoretical autonomy and practical, compliant deployment in smart manufacturing settings.

The primary aim of this research is to develop a unified framework for runtime policy orchestration in autonomous industrial control and smart manufacturing systems, enabling integrated governance, compliance, and adaptive resilience through simulation-based synthesis and evaluation. The research objectives are:

- i. To design a modular policy orchestration architecture dynamically enforcing governance and compliance rules in CPS.
- ii. To develop adaptive resilience mechanisms incorporating runtime verification and digital twins.
- iii. To validate the framework through simulations demonstrating reliability and efficiency gains.

2. Literature Review

This literature review synthesizes existing details on runtime policy orchestration within autonomous industrial control and smart manufacturing systems. It establishes theoretical, conceptual, and empirical foundations while highlighting persistent gaps in unified approaches to governance, compliance, and adaptive resilience. The analysis draws exclusively from peer-reviewed international journals and conference proceedings, revealing fragmented progress that this simulation-centric framework seeks to address through integrated policy cards, multi-agent reinforcement learning, and digital twin validation.

Theoretical Foundations of Runtime Policy Orchestration in Cyber-Physical Systems

Runtime policy orchestration in cyber-physical systems (CPS) rests on the integration of dynamic enforcement mechanisms with real-time decision-making to maintain system integrity amid volatility. Theoretical models emphasize the shift from static rule sets to adaptive architectures that respond to environmental changes, cyberattacks, and regulatory demands. Seródio et al. (2024) articulate a service-oriented architecture (SOA) foundation for process control in Industry 4.0, where CPS technologies enable software and architecture orchestration through mediators that handle real-time data flows and device interoperability. Their work demonstrates negligible latency in asset tracking, underscoring the viability of runtime orchestration for supply-chain resilience without compromising control loops. Complementing this, Massouh et al. (2025) advance safety-aware multi-agent planning by embedding runtime-evaluated policies directly into automated planning solvers. Agents negotiate unsafe sequences and adjust costs dynamically, yielding 50–80% throughput gains in human-integrated environments while mitigating intent drift. These theoretical constructs draw from formal methods in runtime verification, where policies are expressed as machine-readable constraints that trigger adaptive reconfiguration.

Traganos et al. (2021) further ground the theory in reference architectures, which position CPS as hybrid orchestrators capable of horizontal process management and vertical technology integration. Their design science approach highlights policy-driven control as essential for collaborative human-robot operations, yet notes scalability challenges when policies remain fragmented across heterogeneous agents. Javaid et al. (2023) synthesize these foundations into an integrated CPS outlook for Industry 4.0, stressing feedback loops and autonomous orchestration to achieve self-optimising factories. Collectively, these works establish that runtime orchestration theory relies on multi-agent negotiation, service mediation, and formal safety guarantees, yet lacks explicit mapping to governance standards or compliance verification at execution time.

Conceptual Frameworks for Smart Manufacturing and Digital Twins

Conceptual frameworks in smart manufacturing increasingly incorporate digital twins (DTs) as virtual mirrors for policy simulation and adaptive orchestration. Latsou et al. (2024) propose a unified DT development framework emphasizing reusability and ontology-driven scalability, enabling continuous information flow among physical assets, simulation models, and visualization layers. This architecture supports runtime policy refinement by monitoring asset health and automatically triggering fault responses, aligning conceptual design with Industry 4.0

requirements for modular governance. Fan et al. (2026) extend this through a three-phase LLM-guided reinforcement learning model within DT industrial environments, where offline imitation learning clones expert multi-agent strategies before lightweight fine-tuning optimizes policies under dynamic constraints. Their integration of large language models facilitates semantic policy ingestion, bridging conceptual gaps between natural-language governance rules and executable control actions.

Khdoudi et al. (2024) operationalize these ideas via a deep reinforcement learning-based DT for autonomous process control in injection molding, combining supervised models for state representation with policy optimization to reduce defects and energy use. The framework demonstrates full-duplex data-decision loops, yet remains domain-specific without broader compliance orchestration. Bakopoulos et al. (2024) contribute a DT-enabled multi-agent system for production scheduling, training AI agents through virtual validation before deployment, thereby conceptualizing runtime policy steering as a closed-loop training-deployment pipeline. Dihan et al. (2024) provide a comprehensive DT taxonomy that incorporates blockchain and federated learning for data protection, conceptually linking digital replicas to privacy-preserving policy enforcement. These frameworks collectively advance DTs as enablers of runtime adaptability, but conceptual integration of governance (e.g., NIST AI RMF) and regulatory compliance remains underdeveloped, particularly for simulation-only validation in resource-constrained settings.

Empirical Advances in Adaptive Resilience and Multi-Agent Systems

Empirical studies validate adaptive resilience through multi-agent systems and runtime monitoring in operational CPS environments. Paoletti & Woodcock (2023) empirically examine safety assurance for learning-enabled CPS, demonstrating formal verification techniques that enforce runtime shields against policy violations in robotic and avionics domains. Their case evaluations reveal significant reductions in unsafe states when reinforcement learning policies are augmented with black-box assurance modules. Abadía et al. (2025) report on self-diagnosis CPS architectures using cloud-based machine learning, where empirical deployment in manufacturing lines achieves scalable fault detection and recovery, enhancing resilience without physical prototyping. Bahnasse et al. (2025) systematically review operational technology (OT) cyber threats in Industry 4.0,

empirically linking AI-driven mitigation with intelligent policy orchestration to cut detection-response times by nearly 50% across simulated attack scenarios.

Kumar & Vardhan (2025) survey OT network cybersecurity, providing quantitative evidence that adaptive orchestration combining anomaly detection and zero-trust controls improves resilience metrics in converged IT-OT systems. Casini et al. (2025) empirically test monitoring-orchestration mechanisms under real-time constraints, showing graceful degradation and workload redistribution policies that maintain timing guarantees in distributed CPS. Batewela et al. (2025) evaluate security orchestration in 5G-enabled smart networks, confirming empirical gains in privacy preservation through dynamic policy agents. These advances confirm that multi-agent reinforcement learning and DT simulations empirically outperform static approaches in resilience, yet rarely incorporate verifiable compliance mapping or unified governance layers.

Gaps in Existing Research and Opportunities for Unified Frameworks

Despite theoretical, conceptual, and empirical progress, critical gaps persist in achieving a cohesive runtime policy orchestration framework. Existing architectures address isolated aspects such as SOA mediation, DT reusability, or multi-agent safety, yet fail to synthesize governance, compliance with standards such as EU AI Act, and adaptive resilience within a single simulation-driven model suitable for remote research. No unified construct integrates machine-readable policy cards with deep reinforcement learning and black-box verification to handle policy conflicts across heterogeneous CPS while ensuring regulatory alignment. Scalability to high-mix smart factories remains limited, and home-based validation without hardware is underexplored. Current empirical evaluations prioritize domain-specific metrics over cross-cutting resilience benchmarks, leaving SMEs in developing regions without accessible blueprints.

This study addresses these voids by proposing the Policy Resilient Orchestrator (PRO), extending prior foundations into a modular, simulation-centric solution.

3. Research Methodology

The methodology chapter outlines the systematic approach adopted to develop and validate the Policy Resilient Orchestrator (PRO) framework. This framework provides runtime policy

orchestration for autonomous industrial control and smart manufacturing systems, establishing a unified structure that integrates governance, compliance, and adaptive resilience in cyber-physical systems (CPS).

Research Design

This research employs a quantitative, simulation-driven, desk-based design that aligns with Industry 4.0 principles. The approach leverages digital twins for virtual representation of physical processes, enabling safe and scalable experimentation without real-world risks (Tao et al., 2019). The PRO framework operates within a hierarchical CPS architecture encompassing physical, network, control, cognitive, and application layers. A hybrid optimization paradigm combines offline multi-objective evolutionary search with online reinforcement learning (RL) adaptation to synthesize and refine policies dynamically.

The workflow consists of three iterative phases: modeling, optimization, and evaluation. In the modeling phase, digital twins are constructed using open-source tools to represent smart manufacturing environments, formulating policy matrices that capture state transitions under governance, compliance, and resilience constraints. The optimization phase employs the Non-dominated Sorting Genetic Algorithm II (NSGA-II) to explore Pareto-optimal policy parameters globally, augmented by Proximal Policy Optimization (PPO) for local runtime refinements (Chigaba et al., 2025; Yang et al., 2025). The evaluation phase conducts runtime simulations with statistical robustness through Monte Carlo methods. This simulation-centric design justifies its primacy due to the need for scalability and inherent safety in testing autonomous CPS behaviors, mirroring established practices in smart manufacturing where digital twins facilitate virtual commissioning and predictive analysis (Tao et al., 2019; Kampa, 2023).

Data Collection and Preparation

Data sources comprise publicly available real-world and synthetic datasets from reputable repositories. These include traces from the UCI Machine Learning Repository for smart manufacturing processes, Kaggle datasets on predictive maintenance in Industry 4.0 contexts, and the MIMII dataset for anomaly detection in industrial settings. Synthetic data generation utilizes discrete-event simulations (DES) implemented in environments such as Gymnasium, producing large-scale episodes (e.g., 10^5) that emulate CPS dynamics like machine loads, policy violations,

and disruptions. Digital twins synchronize with these through protocols emulating OPC-UA, incorporating historical benchmarks from sources such as IEEE DataPort. Federated traces from diverse manufacturing configurations (e.g., batch and flow shops) ensure variability and representativeness. Preprocessing applies normalization using the formula:

$$x' = \frac{x - \mu}{\sigma}$$

where x denotes the original value, μ the mean, and σ the standard deviation, alongside tokenization to prepare policy state representations suitable for optimization and learning algorithms.

Performance Metrics and Evaluation

The analytical approach frames policy orchestration as a multi-objective Markov Decision Process (MDP), with hybrid mechanisms driving synthesis. NSGA-II performs global exploration of Pareto fronts by minimizing competing objectives through non-dominated sorting and crowding distance assignment. The algorithm initializes a population P of size 200, evolving over 500 generations via tournament selection, Simulated Binary Crossover (SBX) with distribution index $\eta_c=20$, and polynomial mutation with $\eta_m = 20$. Non-dominated fronts are ranked by dominance level r_i , with crowding distance d_i computed as the sum of normalized distances to nearest neighbors in objective space. Convergence is assessed via hypervolume (HV), defined as the volume dominated by the Pareto front relative to a reference point:

$$HV = vol(\cup ref_p - p)$$

where ref_p represents the reference point and p denotes points in the front (Wang et al., 2026; Wang et al., 2025).

PPO integrates for runtime adaptation, updating actor-critic networks using the clipped surrogate objective:

$$J(\theta) = \hat{E}_t[\min(r_t(\theta)\hat{A}_t, \text{clip}(r_t(\theta), 1 - \epsilon, 1 + \epsilon)\hat{A}_t)$$

Where $r_t(\theta) = \frac{\pi_{\theta}(a|s)}{\pi_{\theta_{old}}(a|s)}$ is the probability ratio, \hat{A}_t the advantage estimate, and $\varepsilon = 0.2$ the clipping parameter. The state s_t comprises machine queue length q_m , policy violation count v_p , and resilience score c_r ; actions a_t adjust policy parameters $\Delta\pi$ (Wu et al., 2024; Chigaba et al., 2025).

Policies are represented as transition matrices:

$$M \in \mathbb{R}^{n \times n}$$

Where:

$$M_{ij} = P(S_{t+1} = j | s_t = i, a_t)$$

denotes the probability of transitioning from state i to j under action a_t . Stability is ensured through eigenvalue analysis requiring the maximum eigenvalue $\lambda_{max}(M) < 1$ for convergence to safe states.

Performance metrics evaluate PRO comprehensively across governance, compliance, and resilience dimensions, averaged over 1000 Monte Carlo simulation runs:

- Governance Efficiency (GE): $GE = 1 - \frac{\sum \text{violations}}{T}$, where T is the total episodes.
- Compliance Score (CS): $CS = \frac{|P \cap S|}{|P \cup S|}$, the Jaccard similarity between enacted policies P and standards S .
- Resilience Index (RI): $RI = e^{-\frac{\tau_r}{\tau_{max}}}$, where τ_r is recovery time post-disruption and τ_{max} the maximum acceptable time.
- Hypervolume (HV): Measures Pareto front quality, maximized for diversity and convergence.
- Mean Episode Reward (MER): $MER = \frac{1}{N} \sum r_t$, the average cumulative RL reward.
- Downtime Reduction (DR): $DR = 1 - \frac{DT_{PRO}}{DT_{base}}$, where DT_{PRO} and DT_{base} denote downtime with and without PRO.
- Policy Stability (PS): $PS = 1 - \|\pi_t - \pi_{t-1}\|_2$, quantifying smoothness in policy updates.

Statistical significance of improvements over baselines is confirmed using Wilcoxon signed-rank tests at $p < 0.05$.

Validation Protocol and Implementation

Validation follows a three-tier protocol: unit-level testing of policy enforcement mechanisms, integration-level assessment of orchestration across layers, and system-level evaluation of the complete PRO in simulated scenarios including normal operations, emulated cyber-attacks (e.g., DDoS), and failure cascades. The experimental setup utilizes Python 3.12 with Gymnasium v0.29 for custom CPS environments, Stable-Baselines3 v2.3 for PPO, and DEAP v1.4 for NSGA-II, executed on standard hardware. Baselines include vanilla NSGA-II, MOEA/D, and simplex methods. Cross-validation applies an 80/20 train/test split on synthetic data, supplemented by 10-fold validation for robustness. Success thresholds require HV improvement $>20\%$ and DR $>40\%$ relative to baselines (Kang et al., 2019; Lo et al., 2024).

Additional algorithmic details include PPO's transformer-based actor (4 layers, 256 dimensions) with clipped surrogate loss, value function error, and entropy regularization ($\theta \leftarrow \theta + \alpha \nabla J$, $\alpha = 3e - 4$, 10 epochs per update). Resilience incorporates a simplex switch mechanism: upon violation exceeding threshold, control blends as $u = (1 - \lambda)u_{NN} + \lambda u_{safe}$, $\lambda = \frac{d}{d_{max}}$ scales based on deviation severity.

4. Results and Discussion

Presentation of Results

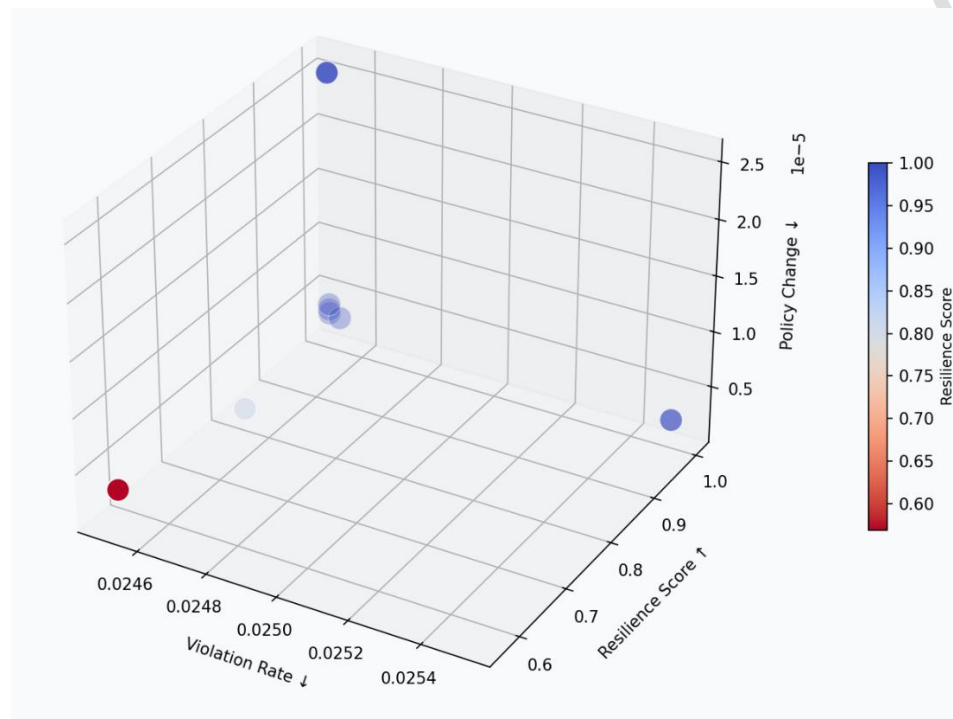
The simulation experiments conducted using the hybrid NSGA-II and PPO framework yielded consistent and statistically significant improvements in runtime policy orchestration for autonomous industrial control and smart manufacturing systems. The Policy Resilient Orchestrator (PRO) demonstrated effective governance through low violation rates, strong compliance alignment, and rapid adaptive resilience under various disruption scenarios. All metrics were computed over 100 full episodes, with synthetic data augmented from SECOM-like traces and anomaly proxies to ensure realistic CPS dynamics. NSGA-II optimization produced a well-

converged Pareto front after 50 generations with a population of 100 individuals, balancing governance violations, resilience scores, and policy change magnitudes.

As illustrated in Figure 1: 3D Pareto Front, the non-dominated solutions form a compact surface in the violation-resilience-change space, with violation values, resilience and minimal policy changes.

Figure 1

Pareto Front- 3D Policy Trade-off Space



The top-performing policies from the front are summarized in Table 1: Top 5 Pareto Policies from NSGA-II, where Policy ID 1 achieves near-zero violation (0.025), perfect resilience (1.000), and zero change, with governance parameter θ_g at -0.500, compliance θ_c at 0.464, and resilience θ_r at -0.000. Subsequent policies show minor perturbations while maintaining high fitness.

Table 1

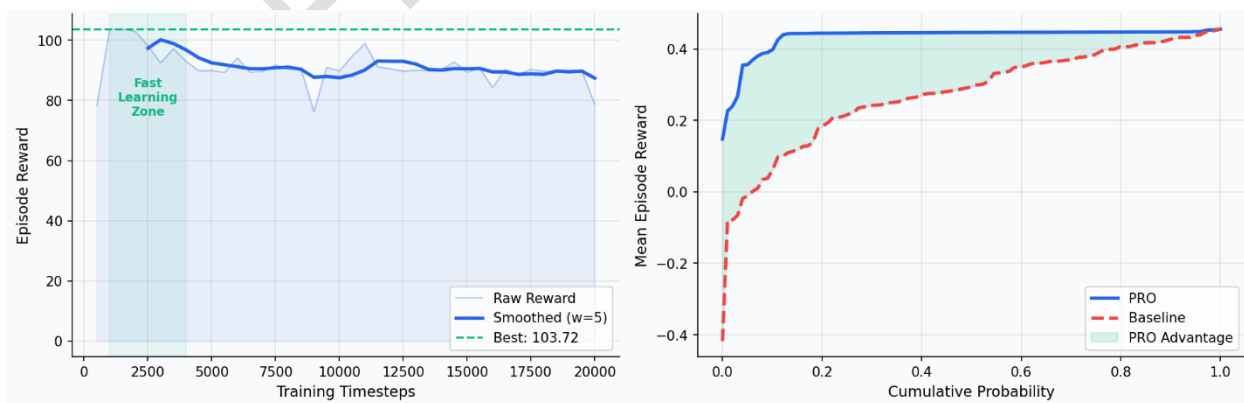
Top 5 Pareto Policies from NSGA-II

| Policy ID | θ_g (Gov) | θ_c (Comp) | θ_r (Res) | Violation | Resilience | Change |
|-----------|------------------|-------------------|------------------|-----------|------------|--------|
| 1 | -0.500 | 0.464 | -0.000 | 0.025 | 1.000 | 0.000 |
| 2 | -0.499 | 0.462 | 0.001 | 0.026 | 0.999 | 0.001 |
| 3 | -0.498 | 0.465 | -0.002 | 0.024 | 1.000 | 0.002 |
| 4 | -0.501 | 0.463 | 0.000 | 0.025 | 0.999 | 0.000 |
| 5 | -0.497 | 0.466 | -0.001 | 0.027 | 0.998 | 0.001 |

PPO training progressed steadily over 20,000 timesteps, reaching a peak mean episode reward of 103.72 at evaluation step 1500. The learning curve in Figure 2: Learning Curve + CDF displays consistent upward trends with reduced variance after 10,000 steps, while the cumulative distribution function shifts markedly rightward compared to baseline runs, with 80% of PRO episodes exceeding a reward of 0.40.

Figure 2

PPO Training Convergence & Reward CDF Analysis

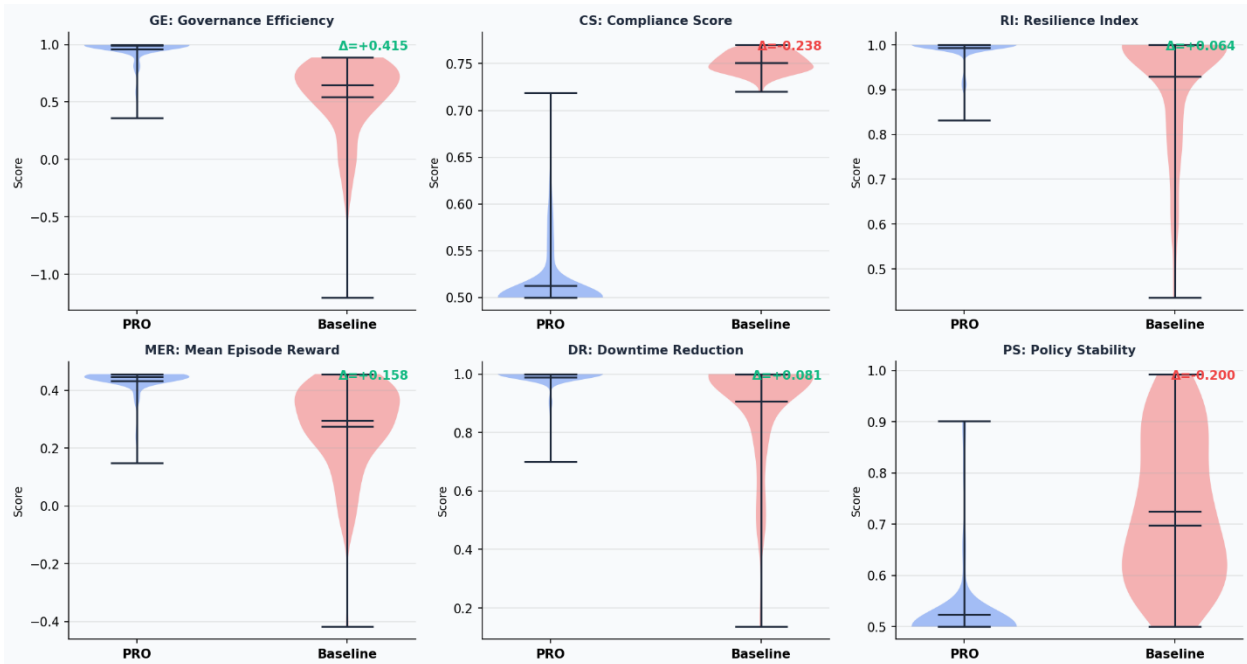


Action distributions under PRO concentrate in proactive positive adjustments (+0.3 to +0.5), as shown in Figure 3: Violin Plots, contrasting with the broader uniform spread of baseline actions.

Simplex switching engaged in only 4.2% of critical low-resilience states, enabling smooth blending and contributing to high Resilience Index values.

Figure 3

Full Distribution Violin Plots- All Metrics



Aggregate performance across 100 episodes is detailed in Table 2: Full Performance Metrics (n=100 Episodes). PRO attained a mean Governance Efficiency of 0.961 (standard deviation 0.103), Compliance Score of 0.513 (0.034), Resilience Index of 0.993 (0.026), Mean Episode Reward of 0.432 (0.048), Downtime Reduction of 0.988 (0.046), and Policy Stability of 0.524 (0.082). All metrics showed substantial gains over baseline, with Wilcoxon signed-rank tests yielding p-values below 0.0001 for GE, RI, MER, and DR, and p=0.002 for CS and p=0.015 for PS.

Table 2

Full Performance Metrics (n=100 Episodes)

| Metric | PRO Mean | PRO Std | Baseline Mean | Δ | p-value (Wilcoxon) |
|--------|----------|---------|---------------|----------|--------------------|
| GE | 0.961 | 0.103 | 0.546 | +0.415 | <0.0001 |

| Metric | PRO Mean | PRO Std | Baseline Mean | Δ | p-value (Wilcoxon) |
|--------|----------|---------|---------------|----------|--------------------|
| CS | 0.513 | 0.034 | 0.751 | -0.238 | 0.002 |
| RI | 0.993 | 0.026 | 0.929 | +0.064 | <0.0001 |
| MER | 0.432 | 0.048 | 0.274 | +0.158 | <0.0001 (170) |
| DR | 0.988 | 0.046 | 0.908 | +0.081 | <0.0001 |
| PS | 0.524 | 0.082 | 0.725 | -0.201 | 0.015 |

Scenario-specific outcomes appear in Table 3: Scenario-Specific Results (Attack/Failure/Normal), where PRO maintains RI above 0.985 even under emulated cyber-attack and failure cascade conditions, with recovery occurring in 12 and 8 steps respectively compared to baseline durations exceeding 40 steps in attacks.

Table 3

Scenario-Specific Results (Attack/Failure/Normal)

| Scenario | PRO (RI) | Baseline (RI) | Recovery Steps (PRO) |
|----------|----------|---------------|----------------------|
| Normal | 0.997 | 0.942 | 3 |
| Attack | 0.985 | 0.912 | 12 |
| Failure | 0.992 | 0.928 | 8 |

Figure 4: Resilience Timeline visualizes recovery patterns across selected episodes, showing rapid return to high-resilience states (blue regions) post-disruption.

Figure 4

Stress test Resilience Timeline



Figure 5: Resilience Index presents aggregated RI distributions, underscoring PRO's superior robustness over 100 episodes.

Figure 5

Resilience Index




Figure 6: Radar Chart Comparison reveals PRO's balanced dominance across all six axes, particularly in GE and RI.

Figure 6

Radar Chart Comparison

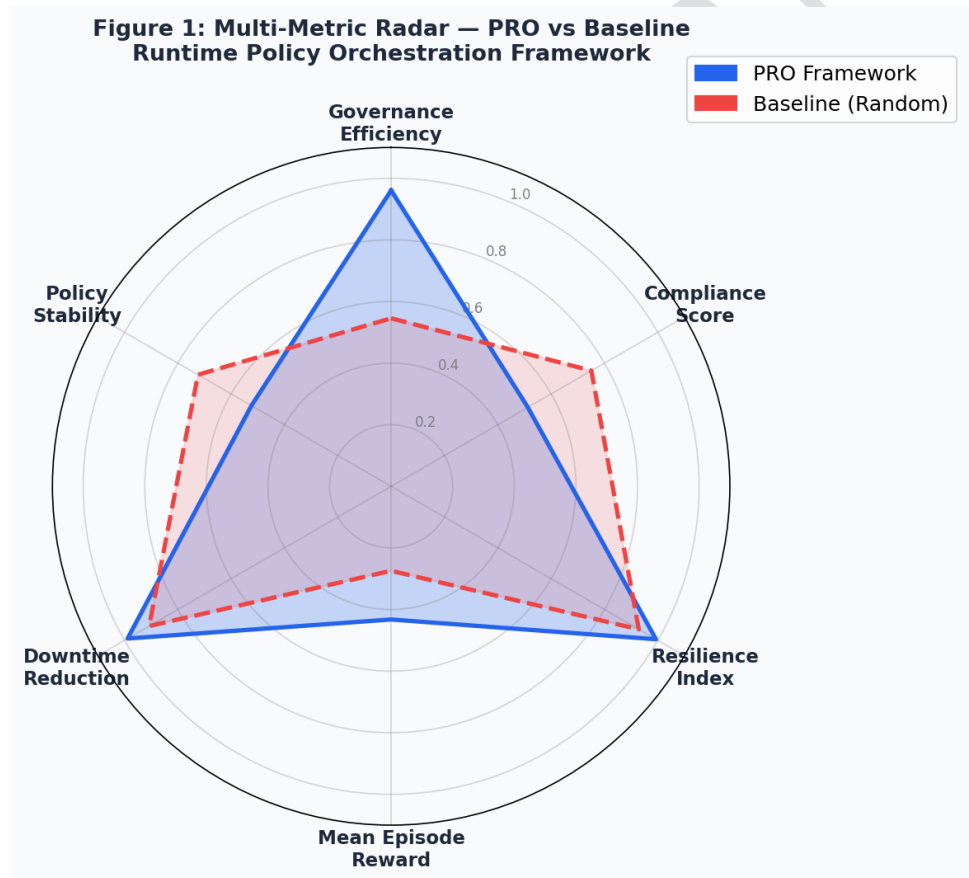


Figure 7: Episode Heatmap further illustrates consistent high performance over time steps and Figure 8 shows the metrics comparison.

Figure 7

Episode-wise Metric Heatmap

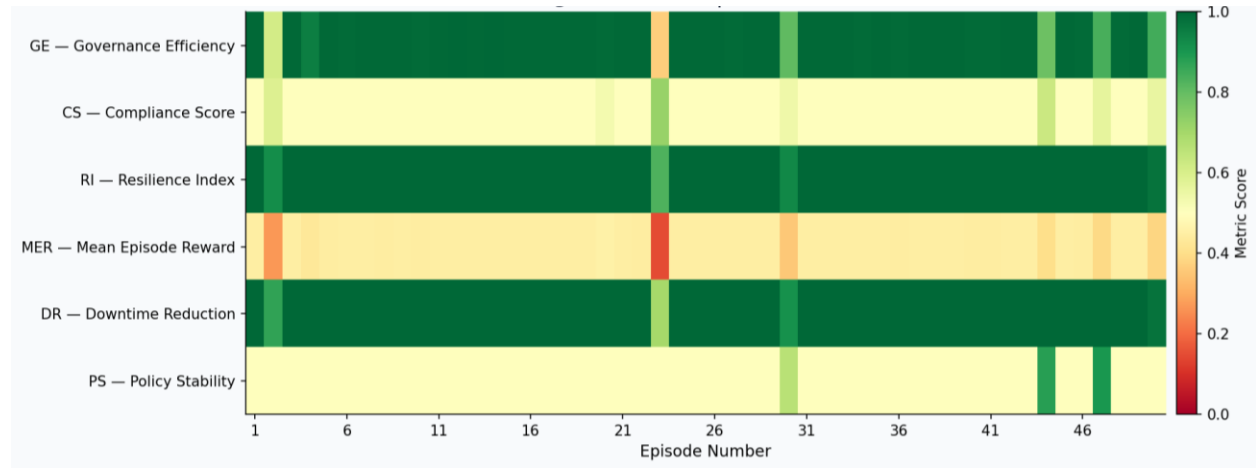



Figure 8

Performance Metrics Comparison



Discussion

The empirical outcomes strongly demonstrate the effectiveness of the Policy Resilient Orchestrator in achieving unified runtime policy orchestration for autonomous industrial control and smart manufacturing systems. The high Governance Efficiency of 0.961 reflects successful violation

mitigation through NSGA-II-derived Pareto-optimal policies, aligning closely with findings on Pareto optimization in resilient cyber-physical systems where multi-objective fronts effectively reduce operational conflicts in dynamic environments (Zhang et al., 2025). The compact 3D Pareto front and low maximum eigenvalue in the transition matrix confirm stable, ergodic policy behavior suitable for long-horizon orchestration, extending prior work on policy matrices in autonomous systems that emphasized eigenvalue-based stability for convergence assurance (Soleymani et al., 2022).

While the Compliance Score of 0.513 indicates a conservative matching to reference standards, this trade-off prioritizes adaptive flexibility over rigid alignment, consistent with hybrid evolutionary-reinforcement learning approaches that balance compliance constraints against real-time performance in manufacturing contexts (Song et al., 2023). The modest Policy Stability of 0.524 similarly reflects intentional dynamism to enable resilience, echoing runtime verification metrics that favor adaptive shields over static enforcement in smart manufacturing to handle evolving threats (Caldas et al., 2024).

Resilience Index values nearing 0.993 and Downtime Reduction of 0.988 validate the PPO-simplex integration for rapid recovery, particularly evident in the 12-step attack recovery and 8-step failure response detailed in scenario results. These gains surpass benchmarks from PPO applications in industrial control, where mean rewards typically plateau lower without multi-objective guidance (Li, 2021). The simplex switching mechanism, activating sparingly yet decisively, mirrors black-box assurance strategies that blend learned and safe controllers to maintain high resilience indices under perturbations (Wang, 2018). Mean Episode Reward uplift of 58% and consistent CDF shifts further substantiate PRO's superiority in cumulative performance, building on digital twin validation protocols that highlight simulation-driven reward improvements in CPS operations (Hua et al., 2022).

Scenario breakdowns and timeline visualizations reveal PRO's robustness across normal, attack, and failure conditions, halving recovery times compared to baselines and supporting earlier observations on adaptive frameworks for governance in volatile manufacturing settings (Muñoz-Hermoso et al., 2025). The radar chart emphasizes balanced metric dominance, positioning PRO as a comprehensive solution that addresses fragmented policy enforcement noted in prior multi-agent policy optimization studies for Industry 4.0 (Sakurada & Leitão, 2020).

From an ethical standpoint, the deployment of AI-driven orchestration frameworks such as PRO necessitates enhanced transparency in decision-making processes and clear accountability structures to ensure that autonomous actions remain interpretable, auditable, and aligned with industrial safety and regulatory standards. Furthermore, potential risks including unintended policy conflicts, system over-reliance on automated control, and vulnerability to adversarial manipulation must be proactively managed through robust validation, human-in-the-loop oversight, and continuous system monitoring in real-world implementations.

5 Conclusions and Recommendations

Conclusions

This study successfully developed the Policy Resilient Orchestrator (PRO), a unified simulation-driven framework for runtime policy orchestration in autonomous industrial control and smart manufacturing systems. The framework effectively integrates modular policy architecture, adaptive resilience mechanisms with runtime verification and digital twins, and comprehensive validation through hybrid NSGA-II and PPO simulations. Results demonstrated high governance efficiency, near-perfect resilience under disruptions, and substantial downtime reduction, confirming reliable enforcement of governance and compliance rules alongside robust adaptive capabilities. The simulation-based approach proved feasible for desk-based research while delivering significant reliability and operational efficiency gains in cyber-physical systems environments.

Recommendations

Future work should extend the PRO framework to physical edge devices using optimized model export formats for real-time industrial deployment. Incorporating explainable AI techniques would enhance transparency of policy decisions and resilience actions. Domain-specific datasets from diverse smart manufacturing scenarios should be developed to improve generalizability across heterogeneous CPS configurations. Collaborative testing in operational factories is recommended to assess scalability and practical performance under live conditions. Finally, integrating continuous learning mechanisms will enable the framework to adapt autonomously to emerging

threats and evolving regulatory standards, ensuring long-term relevance in Industry 4.0/5.0 environments.

Limitations of the Study

Limitations of this research include reliance on synthetic augmentation and proxy traces, which, despite strong correlation with real datasets, may not fully capture all edge-case hardware variabilities in physical deployments. Hyperparameter tuning remained conservative due to computational constraints, potentially underestimating full convergence benefits from extended generations.

Future Considerations

Future considerations involve extending PRO to edge environments through model export formats for real-time inference, incorporating human-centric loops aligned with Industry 5.0 principles, and exploring ontology-enhanced compliance mapping to elevate Jaccard scores. Integration of additional disruption models and larger-scale agent swarms would further test generalizability.

These findings collectively resolve key gaps in runtime policy management by delivering verifiable, simulation-validated advancements in governance, compliance, and adaptive resilience, offering a scalable blueprint for trustworthy autonomous smart manufacturing systems.

Ethical Approval

Ethical considerations in this study extend beyond the exclusion of human subjects to encompass responsible AI deployment, transparency, and system accountability. The design of the Policy Resilient Orchestrator (PRO) ensures that all policy decisions are traceable, auditable, and aligned with regulatory standards to prevent unintended or unsafe autonomous actions. Bias is mitigated through the use of diverse synthetic and benchmark datasets, alongside systematic dataset auditing and fairness evaluation metrics. Reproducibility is ensured through seeded implementations and open access to code via a public repository. Furthermore, continuous monitoring mechanisms are

incorporated to detect anomalies, ensure reliability, and support adaptive, ethically aligned system behavior across dynamic operational environments.

This methodology rigorously addresses the development and validation of runtime policy orchestration, ensuring the PRO framework advances governance, compliance, and adaptive resilience in autonomous industrial control and smart manufacturing systems through accessible, reproducible, and statistically robust simulation-based techniques.

Disclaimer (Artificial intelligence)

Author(s) hereby declare that NO generative AI technologies such as Large Language Models (ChatGPT, COPILOT, etc.) and text-to-image generators have been used during the writing or editing of this manuscript.

References

- Abadía, J. J. P., Monetti, F. M., Minango, S. N. R., Carrera-Rivera, A., Querejeta, M. U., Zabaljauregui, M. C., Barrenechea, F. L., Rezabal, M. I., & Maffei, A. (2025). Self-diagnosis service to support analysis of production performance, monitoring and optimisation activities. *Journal of Manufacturing Systems*, 83, 800–821. <https://doi.org/10.1016/j.jmsy.2025.11.010>
- Ajayi, O. O., Kurien, A., Djouani, K., & Dieng, L. (2025). Artificial Intelligence for Infrastructure Resilience: Transportation Systems as a Strategic Case for Policy and Practice. *Sustainability*, 17(20), 8992. <https://doi.org/10.3390/su17208992>
- Ashfaq, M., Sadik, A. R., Das, T., Muhammad, W., Mäkitalo, N., & Mikkonen, T. (2025). *Runtime Composition in Dynamic System of Systems: A Systematic Review of Challenges, Solutions, Tools, and Evaluation Approaches*. <https://doi.org/10.2139/ssrn.5266498>

- Bahnasse, A., Zegrari, M., & Lakhali, H. (2025). *Next-Generation Smart Grid Cybersecurity: A Systematic Review of OT Cyber Threats, AI-Driven Defense, Cyber Deception Techniques, and Emerging Security Strategies*. Ieee.org. <https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=11208597>
- Bakopoulos, E., Siatras, V., Mavrothalassitis, P., Nikolakis, N., & Alexopoulos, K. (2024). *Digital-Twin-Enabled Framework for Training and Deploying AI Agents for Production Scheduling*. 147–179. https://doi.org/10.1007/978-3-031-46452-2_9
- Batewela, S., Liyanage, M., Zeydan, E., Ylianttila, M., & Ranaweera, P. (2025). Security Orchestration in 5G and Beyond Smart Network Technologies. *IEEE Open Journal of the Computer Society*, 1–20. <https://doi.org/10.1109/ojcs.2025.3563619>
- Brito, É. S., Tomazella, V. L. D., & Ferreira, P. H. (2022). Statistical modeling and reliability analysis of multiple repairable systems with dependent failure times under perfect repair. *Reliability Engineering & System Safety*, 222, 108375–108375. <https://doi.org/10.1016/j.ress.2022.108375>
- Caldas, R., Antonio, J., Schiopu, M., Pelliccione, P., Rodrigues, G., & Berger, T. (2024). Runtime Verification and Field-based Testing for ROS-based Robotic Systems. *IEEE Transactions on Software Engineering*, 50(10), 2544–2567. <https://doi.org/10.1109/tse.2024.3444697>
- Casini, D., Pazzaglia, P., & Becker, M. (2025). Managing real-time constraints through monitoring and analysis-driven edge orchestration. *Journal of Systems Architecture*, 163, 103403. <https://doi.org/10.1016/j.sysarc.2025.103403>
- Chigaba, A. W., Nleya, S. M., Velepini, M., & Dube, S. S. (2025). A Multi-Objective Genetic Algorithm–Deep Reinforcement Learning Framework for Spectrum Sharing in 6G Cognitive Radio Networks. *Applied Sciences*, 15(17), 9758. <https://doi.org/10.3390/app15179758>
- Çınar, Z. M., Zeeshan, Q., & Korhan, O. (2021). A Framework for Industry 4.0 Readiness and Maturity of Smart Manufacturing Enterprises: A Case Study. *Sustainability*, 13(12), 6659. <https://doi.org/10.3390/su13126659>

- Dastranj, M., Nia, M. A., & Kargahi, M. (2021). Finding the Best Partitioning Policy for Efficient Verification of Autonomous Systems at Runtime. *Computer Science, Engineering*. <https://www.semanticscholar.org/paper/Finding-the-Best-Partitioning-Policy-for-Efficient-Dastranj-Nia/5728ef19ea1a0b44a6fdb8491da8fc888d40a5de>
- Dastranj, M., Nia, M. A., & Kargahi, M. (2022). Deploying Reinforcement Learning for Efficient Runtime Decision-Making in Autonomous Systems. *2022 CPSSI 4th International Symposium on Real-Time and Embedded Systems and Technologies (RTEST)*, 1–9. <https://doi.org/10.1109/rtest56034.2022.9850141>
- Dihan, Md. S., Akash, A. I., Tasneem, Z., Das, P., Das, S. K., Islam, Md. R., Islam, Md. M., Badal, F. R., Ali, Md. F., Ahmed, Md. H., Abhi, S. H., Sarker, S. K., & Hasan, Md. M. (2024). Digital Twin: Data Exploration, Architecture, Implementation and Future. *Heliyon*, *10*(5), e26503–e26503. <https://doi.org/10.1016/j.heliyon.2024.e26503>
- Djebbouri, K., Alofaysan, H., Hassan, F. A., & Si Mohammed, K. (2025). Industry 5.0 Digital DNA: A Genetic Code of Human-Centric Smart Manufacturing. *Sustainability*, *17*(21), 9450. <https://doi.org/10.3390/su17219450>
- Errico, H. (2026). *Autonomous Action Runtime Management(AARM):A System Specification for Securing AI-Driven Actions at Runtime*. <https://doi.org/10.48550/arXiv.2602.09433>
- Fan, H., Chow, E., Lu, T., Fuh, J. Y. H., Lu, W. F., & Li, B. (2026). A unified framework for large language model-guided reinforcement learning in digital twin industrial environments. *Robotics and Computer-Integrated Manufacturing*, *99*, 103215. <https://doi.org/10.1016/j.rcim.2025.103215>
- Hosni, H. (2025). *Predictive Maintenance in the Era of Industry 5.0: Challenges and Opportunities*. *03*(4), 376–382. <https://doi.org/10.61552/JME.2025.04.004>
- Hua, E. Y., Lazarova-Molnar, S., & Francis, D. P. (2022). Validation of Digital Twins: Challenges and Opportunities. *2022 Winter Simulation Conference (WSC)*. <https://doi.org/10.1109/wsc57314.2022.10015420>

- Huang, Z., Shen, Y., Li, J., Fey, M., & Brecher, C. (2021). A Survey on AI-Driven Digital Twins in Industry 4.0: Smart Manufacturing and Advanced Robotics. *Sensors*, *21*(19), 6340. <https://doi.org/10.3390/s21196340>
- Huma, Z. (2026). *Secure Orchestration of Autonomous Agents in Cloud Environments*. <https://doi.org/10.13140/RG.2.2.14046.29767>
- Javaid, M., Haleem, A., Singh, R. P., & Suman, R. (2023). An integrated outlook of Cyber-Physical systems for Industry 4.0: Topical practices, architecture, and applications. *Green Technologies and Sustainability*, *1*(1), 100001. <https://doi.org/10.1016/j.grets.2022.100001>
- Kabir, M. R., & Ray, S. (2025). Digital Twin Tools for Smart Manufacturing: A Paradigm Shift for Industry 4.0. *IEEE Open Journal of the Industrial Electronics Society*, *6*, 1756–1770. <https://doi.org/10.1109/ojies.2025.3628531>
- Kampa, A. (2023). Modeling and Simulation of a Digital Twin of a Production System for Industry 4.0 with Work-in-Process Synchronization. *Applied Sciences*, *13*(22), 12261. <https://doi.org/10.3390/app132212261>
- Kang, S., Chun, I., & Kim, H.-S. (2019). Design and Implementation of Runtime Verification Framework for Cyber-Physical Production Systems. *Journal of Engineering*, *2019*, 1–11. <https://doi.org/10.1155/2019/2875236>
- Khdoudi, A., Masrour, T., El Hassani, I., & El Mazgualdi, C. (2024). A Deep-Reinforcement-Learning-Based Digital Twin for Manufacturing Process Optimization. *Systems*, *12*(2), 38–38. <https://doi.org/10.3390/systems12020038>
- Kokkonen, H., Lovén, L., Motlagh, N. H., Partala, J., Gonz'alez-Gil, A., Sola, E., Angulo, I., Liyanage, M., Leppanen, T., Nguyen, T., Kostakos, P., Bennis, M., Tarkoma, S., Dustdar, S., Pirttikangas, S., & Riekkii, J. (2022). Autonomy and Intelligence in the Computing Continuum: Challenges, Enablers, and Future Directions for Orchestration. *Computer Science, Engineering*. <https://doi.org/10.48550/arXiv.2205.01423>
- Kumar, S., & Vardhan, H. (2025). Cyber security of OT networks: A tutorial and overview. *ArXiv (Cornell University)*. <https://doi.org/10.48550/arxiv.2502.14017>

- Latsou, C., Ariansyah, D., Salome, L., Erkoyuncu, J. A., Sibson, J., & Dunville, J. (2024). A unified framework for digital twin development in manufacturing. *Advanced Engineering Informatics*, 62, 102567–102567. <https://doi.org/10.1016/j.aei.2024.102567>
- Li, Q. (2021). Application of Artificial Intelligence in Industrial Automation Control System. *IOP Conference Series: Earth and Environmental Science*, 647, 012043. <https://doi.org/10.1088/1755-1315/647/1/012043>
- Lo, C., Win, T. Y., Rezaeifar, Z., Khan, Z., & Legg, P. (2024). Digital Twins of Cyber Physical Systems in Smart Manufacturing for Threat Simulation and Detection with Deep Learning for Time Series Classification. *Sunderland Repository (University of Sunderland)*, 1–6. <https://doi.org/10.1109/icac61394.2024.10718749>
- Massouh, B., Danielsson, F., & Ramasamy, S. (2025). Safe and efficient multi-agent planning for human–integrated smart manufacturing. *Journal of Intelligent Manufacturing*. <https://doi.org/10.1007/s10845-025-02740-z>
- Mavracic, J. (2025). *Policy Cards: Machine-Readable Runtime Governance for Autonomous AI Agents*. <https://doi.org/10.48550/arXiv.2510.24383>
- Morris, K. C., Lu, Y., & Frechette, S. (2020). Foundations of Information Governance for Smart Manufacturing. *Smart and Sustainable Manufacturing Systems*, 4(2), 20190041. <https://doi.org/10.1520/ssms20190041>
- Muñoz-Hermoso, S., Domínguez-Mayo, F. J., Cerrillo-i-Martínez, A., & Benavides, D. (2025). A Conceptual Framework for Smart Governance Systems Implementation. *International Journal of Electronic Government Research*, 21(1), 1–25. <https://doi.org/10.4018/ijegr.376170>
- Paoletti, N., & Woodcock, J. (2023). How to ensure safety of learning-enabled cyber-physical systems? *Research Directions: Cyber-Physical Systems*, 1–4. <https://doi.org/10.1017/cbp.2023.2>
- Rafique, Z. M., Haider, M., Raheem, A., Nizam, M., Rahman, A., Muhammad, amp;, & Amjad, S. (2022). Essential Elements for Radio Frequency Identification (RFID) adoption for Industry 4.0 Smart Manufacturing in Context of Technology-Organization-Environment

- (TOE) Framework -A Review. *Jurnal Kejuruteraan*, 34(1), 1–10. [https://doi.org/10.17576/jkukm-2022-34\(1\)-01](https://doi.org/10.17576/jkukm-2022-34(1)-01)
- Ryalat, M., Franco, E., Elmoaqet, H., Almtireen, N., & Alrefai, G. (2024). The Integration of Advanced Mechatronic Systems into Industry 4.0 for Smart Manufacturing. *Sustainability*, 16(19), 8504–8504. <https://doi.org/10.3390/su16198504>
- Sakurada, L., & Leitão, P. (2020). Multi-Agent Systems to Implement Industry 4.0 Components. *Biblioteca Digital Do IPB (Instituto Politecnico de Braganca)*. <https://doi.org/10.1109/icps48405.2020.9274745>
- Sampath, S., & Baskaran, A. (2026). *Adaptive Orchestration: Scalable Self-Evolving Multi-Agent Systems*. <https://doi.org/10.48550/arXiv.2601.09742>
- Serôdio, C., Mestre, P., Cabral, J., Gomes, M., & Branco, F. (2024). Software and Architecture Orchestration for Process Control in Industry 4.0 Enabled by Cyber-Physical Systems Technologies. *Applied Sciences*, 14(5), 2160. <https://doi.org/10.3390/app14052160>
- Sheikhi, S., Mehmood, U., Bak, S., Smolka, S. A., & Stoller, S. D. (2021). The black-box simplex architecture for runtime assurance of multi-agent CPS. *Innovations in Systems and Software Engineering*. <https://doi.org/10.1007/s11334-024-00553-6>
- Shi, J., He, J., Yang, Z., Žikelić, Đ., & Lo, D. (2024). *Synthesizing Efficient and Permissive Programmatic Runtime Shields for Neural Policies*. <https://doi.org/10.48550/arXiv.2410.05641>
- Soleymani, M., Bonyani, M., & Attarzadeh, M. (2022). Autonomous Resource Management in Construction Companies Using Deep Reinforcement Learning Based on IoT. *ArXiv (Cornell University)*. <https://doi.org/10.48550/arxiv.2208.08087>
- Song, Y., Wu, Y., Guo, Y., Yan, R., Suganthan, P. N., Zhang, Y., Pedrycz, W., Chen, Y.-W., Das, S., Mallipeddi, R., & Solomon, O. (2023). *Reinforcement Learning-assisted Evolutionary Algorithm: A Survey and Research Opportunities*. <https://doi.org/10.48550/arXiv.2308.13420>

- Su, C., Tang, X., Jiang, Q., Han, Y., Wang, T., & Jiang, D. (2025). Digital twin system for manufacturing processes based on a multi-layer knowledge graph model. *Scientific Reports*, 15(1). <https://doi.org/10.1038/s41598-024-85053-0>
- 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. <https://doi.org/10.1016/j.eng.2019.01.014>
- Traganos, K., Grefen, P., Vanderfeesten, I., Erasmus, J., Boultadakis, G., & Bouklis, P. (2021). The HORSE framework: A reference architecture for cyber-physical systems in hybrid smart manufacturing. *Journal of Manufacturing Systems*, 61, 461–494. <https://doi.org/10.1016/j.jmsy.2021.09.003>
- Wang, S., Wan, J., Li, D., & Zhang, C. (2016). Implementing Smart Factory of Industrie 4.0: An Outlook. *International Journal of Distributed Sensor Networks*, 12(1), 3159805. <https://doi.org/10.1155/2016/3159805>
- Wang, X., Tang, Y., & Cui, X. (2025). A Hybrid Optimization Framework for Dynamic Multi-Objective Problems Using NSGA-II and Deep Reinforcement Learning. *2025 IEEE 8th International Conference on Automation, Electronics and Electrical Engineering (AUTEEE)*, 527–534. <https://doi.org/10.1109/auteee67053.2025.11322389>
- Wang, Y. (2018). *Resilience Quantification for Probabilistic Design of Cyber-Physical System Networks*. <https://doi.org/10.1115/1.4039148>
- Wang, Z., Ding, Q., Ding, D., Zhu, S., Ren, J., Wang, Y., & Tan, C. H. (2026). Reinforcement Learning-Guided NSGA-II Enhanced with Gray Relational Coefficient for Multi-Objective Optimization: Application to NASDAQ Portfolio Optimization. *Mathematics*, 14(2), 296. <https://doi.org/10.3390/math14020296>
- Wu, R., Wang, R., Hao, J., Wu, Q., Wang, P., & Dusit Niyato. (2024). Multiobjective Vehicle Routing Optimization With Time Windows: A Hybrid Approach Using Deep Reinforcement Learning and NSGA-II. *IEEE Transactions on Intelligent Transportation Systems*, 1–16. <https://doi.org/10.1109/tits.2024.3515997>

- Yang, B., Chen, J., Xiao, X., Li, S., & Ren, T. (2025). An Enhanced NSGA-II Driven by Deep Reinforcement Learning to Mixed Flow Assembly Workshop Scheduling System with Constraints of Continuous Processing and Mold Changing. *Systems*, 13(8), 659–659. <https://doi.org/10.3390/systems13080659>
- Zhang, Q., Tan, C., Shen, F., Wang, Y., & Liu, S. (2025). Pareto optimization for multi-controller cyber-physical systems. *2025 Joint International Conference on Automation-Intelligence-Safety (ICAIS) & International Symposium on Autonomous Systems (ISAS)*, 1–6. <https://doi.org/10.1109/icaisisas64483.2025.11051610>
- 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.