

Human-Centric Swarm Intelligence for Sustainable Distributed Manufacturing in Industry 5.0

Cezar BRAICU¹, Cătălina Loredana ROGOZEANU², Marius PETRESCU^{2*},
Rareș-Vladimir NIȚU³, Ionuț-Adrian GHIBANU²

¹ Faculty of Management, Bucharest University of Economic Studies, Bucharest, 010374, Romania
cezar.braicu@man.ase.ro

² Faculty of Economics, Valahia University of Targoviște, Targoviște, 130004, Romania
catalina.r.mg@valahia.ro, profdrmpetrescu@yahoo.com (*Corresponding author), ighibanu@yahoo.com

³ Faculty of Science and Engineering, Valahia University of Targoviște, Alexandria, 140003, Romania
nitu.rares.vladimir@gmail.com

Abstract: The transition from Industry 4.0 to Industry 5.0 introduces a structural reconfiguration of the distributed manufacturing systems, shifting the emphasis from automation-centric efficiency toward human-centricity, sustainability, and systemic resilience. Although robotic swarm intelligence has demonstrated its effectiveness in decentralized task allocation and adaptive optimization, the existing architectures remain predominantly performance-driven, treating human workload and stability as secondary indicators. This study proposes a human-centric swarm intelligence control architecture in which sustainability and the supervisory workload are internalized within the distributed functional objective. This framework integrates Particle Swarm Optimization-based coordination, reinforcement learning adaptation, and predictive conditional variance modelling into a unified multi-agent structure. A multi-objective optimization problem is formulated in order to jointly minimize the task completion time, energy consumption, operational volatility, and human workload, while the autonomy is formalized as a dynamically regulated control variable. The simulation experiments conducted across multiple stochastic realizations demonstrate a statistically significant reduction in the value of energy consumption, variance amplification, and supervisory overload, together with an accelerated post-disturbance recovery for the proposed architecture in comparison with the conventional Industry 4.0 swarm configurations. The findings confirm that embedding human-centric and sustainability-aware constraints enhances a system's stability and resilience without compromising its efficiency, thereby validating a mathematically grounded and simulation-based framework aligned with the Industry 5.0 manufacturing principles.

Keywords: Industry 5.0, Human-in-the-loop control, Swarm intelligence, Multi-objective optimization, Reinforcement learning, Distributed manufacturing systems, Resilient distributed control.

1. Introduction

Over the past decade, the digital transformation of industrial infrastructures has fundamentally restructured the architectural and operational organization of production systems. Under the Industry 4.0 paradigm, the manufacturing environments have evolved toward tightly integrated cyber-physical systems, real-time data exchange, distributed sensing architectures, and algorithmic automation (Dias et al., 2021; Oncioiu, 2023). As the operational complexity increases, decentralized coordination mechanisms have progressively replaced centralized control schemes, which typically exhibit scalability limitations and a reduced adaptability under high-dimensional uncertainty.

Within this distributed control landscape, Robotic Swarm Intelligence (RSI) has emerged as a coordination paradigm capable of generating a coherent global behavior through localized interaction rules (Petrescu et al., 2025). Inspired by the collective behaviour in biological systems, swarm architectures exhibit an intrinsic scalability, fault tolerance, and a self-organizing dynamic, properties which are well suited to dynamic and

heterogeneous manufacturing environments. Empirical investigations have confirmed RSI's capacity to reduce the processing latency, improve resource allocation efficiency, and attenuate the localized disruptions in distributed production systems (Akopov, 2024; Schranz et al., 2020).

Notwithstanding these advantages, the Industry 4.0 paradigm has predominantly emphasized efficiency maximization, automation depth, and digital interconnectivity. The emerging Industry 5.0 framework introduces a structural reorientation by repositioning the human actor at the core of the industrial ecosystem (Maddikunta et al., 2022; Kuckling, 2023). Human-centricity, environmental sustainability, and systemic resilience become explicit design constraints rather than peripheral objectives (Ligot & Birattari, 2020). Consequently, intelligent manufacturing systems must reconcile algorithmic optimization with human-machine collaboration and long-term ecological viability.

This paradigm shift entails a substantive transformation of the control architectures. Fully autonomous swarm systems optimized

exclusively for throughput or resource utilization are insufficient under Industry 5.0 constraints. The manufacturing ecosystems must therefore incorporate hybrid decision spaces in which human cognition and algorithmic coordination interact within a unified control framework. From a systems theory perspective, this coexistence introduces heterogeneous temporal scales and decision dynamics, increasing the dimensionality and structural complexity of the control problem (Wang et al., 2025).

Advances in agent-based modelling and reinforcement learning have demonstrated that adaptive mechanisms enhance the decision-making performance in stochastic and non-stationary environments (Ban, 2024). These approaches support feedback-driven adaptation and predictive modelling capabilities. However, their systematic integration into swarm-based industrial control remains limited in practical industrial implementations (Yang et al., 2019). Most RSI implementations are still carried out under static autonomy assumptions, while the sustainability indicators are frequently treated as exogenous evaluation criteria rather than endogenous optimization variables (Maddikunta et al., 2022).

Moreover, the existing approaches remain predominantly efficiency-oriented, with a limited formal integration of human-centric variables such as supervisor workload or stability-related indicators within the control objective. In contrast to adaptive control or model-based optimization strategies, which explicitly address constraint handling and system stability, the swarm-based coordination often relies on heuristic data update mechanisms without embedding such variables directly into the optimization process. This limitation motivates the need for a structurally integrated framework capable of jointly addressing performance, sustainability, and human-in-the-loop interaction.

A structural gap therefore persists between the technological maturity of swarm intelligence within Industry 4.0 and the systemic requirements introduced by Industry 5.0. In particular, the existing architectures lack a unified formulation in which human-centricity, sustainability, and system stability are simultaneously treated as endogenous components of the control process.

This study proposes a human-centric swarm intelligence-based control architecture explicitly designed for sustainable distributed manufacturing under the Industry 5.0 constraints. Swarm coordination is extended beyond the efficiency-driven objectives through the formal integration of a dynamically regulated autonomy coefficient, a sustainability-aware multi-objective functional, and a predictive conditional variance layer capturing the stochastic production disturbances.

Embedding these mechanisms within a distributed multi-agent framework preserves the scalability and adaptability inherent to swarm systems while incorporating collaborative and sustainability-oriented control principles. This structural integration reconciles a decentralized optimization with a human-centered design and contributes to the theoretical advancement of Industry 5.0-oriented manufacturing control.

In order to clearly outline the contribution of this work, the main elements introduced in this study can be summarized as follows:

- i. the formulation of a human-centric swarm intelligence control architecture in which autonomy is explicitly modeled as a dynamically regulated control variable;
- ii. the integration of a sustainability-aware multi-objective functional embedding energy consumption, operational volatility, and supervisor workload as endogenous optimization components;
- iii. the incorporation of predictive conditional variance modelling into swarm coordination dynamics, enabling the anticipatory regulation related to instability propagation;
- iv. the combination of Particle Swarm Optimization and reinforcement learning within a unified adaptive control framework;
- v. the formal analysis of system stability under bounded adaptive gains.
- vi. the simulation-based validation of the proposed architecture, demonstrating improvements in resilience, energy efficiency, and supervisor workload balance in comparison with the conventional Industry 4.0 swarm configurations.

The present formulation reflects a conceptual refinement aimed at explicitly integrating human-

centric and sustainability-oriented variables within the swarm coordination framework, ensuring the alignment with Industry 5.0 requirements and addressing the limitations identified in earlier formulations.

The remainder of the paper is structured as follows. Section 2 presents the theoretical foundations of swarm coordination, hybrid human-in-the-loop control, and multi-objective optimization. Section 3 sets forth the mathematical formulation of the proposed architecture and derives the corresponding stability conditions. Further on, Section 4 details the simulation framework and experimental evaluation, while Section 5 discusses the systemic implications of the obtained results. Finally, Section 6 concludes this study and outlines future research directions for Industry 5.0-oriented distributed manufacturing systems.

2. Theoretical Framework

Swarm intelligence has emerged as a coordination paradigm capable of generating a coherent global behavior through localized interaction rules, rendering it particularly suitable for distributed manufacturing environments (Brambilla et al., 2013). In industrial settings, RSI supports decentralized task allocation, adaptive resource redistribution, and fault-tolerant production control without dependence on rigid hierarchical supervision. These architectures exhibit an inherent scalability and structural robustness, attributes especially relevant for large-scale manufacturing systems characterized by dynamic demand and geographically distributed processing units.

Despite these advantages, most existing swarm-based implementations remain primarily oriented toward efficiency optimization, with a limited consideration of human-centric constraints or system-level stability indicators within the coordination process.

Contributions within the Industry 5.0-oriented manufacturing literature highlight the integration of human-centric design principles and sustainability metrics as structural components of cyber-physical production systems (Maddikunta et al., 2022; Xu et al., 2021). Concurrently, advances in multi-objective swarm optimization have demonstrated the feasibility of balancing competing performance criteria within distributed control environments (Ribeiro et al., 2021). Nevertheless, these research

directions often evolve in parallel rather than in an integrated manner. In particular, the explicit incorporation of human-centric variables and sustainability constraints within swarm-based coordination mechanisms remains limited in practical implementations.

The theoretical appeal of swarm-based coordination derives from its emergent properties: global optimization objectives can be achieved through simple agent-level update mechanisms under partial information constraints (Dorigo et al., 2020; Kuckling, 2023). Empirical investigations confirm the capacity of RSI to reduce execution latency, balance workload distribution, and attenuate localized failures in distributed production systems (Trianni & Nolfi, 2011; Wang et al., 2020). Nevertheless, compared with model-based or adaptive control strategies, which explicitly incorporate system constraints and stability guarantees, swarm-based implementations frequently rely on heuristic coordination rules. As a result, their ability to systematically address multi-dimensional objectives such as sustainability, volatility, and human workload remains limited. Objective formulations typically prioritize throughput maximization, latency minimization, or resource utilization efficiency. While operationally effective, such formulations reflect an automation-centric perspective that inadequately addresses the systemic constraints introduced by Industry 5.0.

Model Predictive Control (MPC) represents a prominent model-based control strategy capable of explicitly handling system constraints while predicting future system behavior over a finite horizon. Recent studies demonstrate its effectiveness in human–robot collaborative environments, where anticipatory control improves coordination efficiency and reduces the need for reactive adjustments (Chu et al., 2024). Compared with swarm-based coordination, which relies on decentralized interaction and emergent behavior, MPC provides a structured optimization framework with explicit constraint handling and predictive capabilities. However, its reliance on centralized or semi-centralized formulations may limit scalability in large-scale distributed manufacturing systems, thereby motivating the exploration of decentralized approaches that retain adaptability while incorporating predictive features.

The Industry 5.0 paradigm reconfigures the relationship between technological autonomy and human agency by promoting collaborative intelligence rather than maximal automation (Xu et al., 2021). Human expertise and algorithmic coordination are expected to coexist within shared decision spaces governed by formally regulated interaction mechanisms. This transition requires conceptualizing autonomy not as a static structural property, but as a dynamic and context-dependent control variable.

In most industrial swarm architectures, autonomy is implicitly treated as static: robotic agents operate under fixed algorithmic authority, while human intervention is invoked reactively in exceptional circumstances. Although such a configuration may preserve operational continuity, it does not provide a formal mechanism for regulating the balance between automated decision-making and supervisory oversight within the control loop. This limitation becomes particularly critical under conditions of increased system volatility or supervisory overload, where static autonomy assumptions may lead to a suboptimal or unstable coordination behavior.

Control theory literature on hybrid and supervisory systems indicates that stability and performance can be preserved when the transitions between automated and supervisory modes are governed by bounded adaptation mechanisms (Ribeiro et al., 2021; Wang et al., 2025). Parallel developments in reinforcement learning and adaptive decision systems further demonstrate that control policies can evolve in response to environmental feedback (Brambilla et al., 2013). However, these adaptive principles have rarely been embedded directly within swarm-based industrial coordination frameworks. It follows that a structural disconnection persists between hybrid human-in-the-loop theory and practical swarm implementations, limiting the responsiveness to supervisory overload, contextual uncertainty, and risk escalation.

Beyond human-centricity, Industry 5.0 introduces sustainability and systemic resilience as intrinsic design objectives. Traditional RSI formulations often include energy consumption as a secondary performance indicator, yet it remains subordinate to primary objectives such as execution speed. Similarly, supervisory workload and operational

stability are typically assessed *ex post* rather than incorporated within the control functional.

From a theoretical perspective, the separation between performance metrics and sustainability indicators constrains the capacity of distributed control systems to achieve balanced trade-offs (Schranz et al., 2020). Multi-objective optimization provides a formal mechanism for reconciling the competing criteria within a unified framework. Nevertheless, in swarm-based industrial coordination, such approaches are frequently confined to technical parameters, without the explicit integration of supervisory workload or volatility dynamics. Thereby, human-centric and sustainability-related variables remain peripheral rather than endogenous drivers of coordination behavior.

Distributed manufacturing environments are inherently exposed to stochastic disturbances arising from demand fluctuations, processing variability, communication latency, and resource constraints. Classical swarm coordination addresses such perturbations through reactive task redistribution. Although decentralized autonomy is preserved, delayed collective adaptation often generates secondary oscillations and leads to instability amplification.

Volatility should not be interpreted solely as transient noise. From a systems-theoretic perspective, conditional variance in production dynamics represents a structural indicator of instability potential. Predictive modelling approaches demonstrate that variance estimation can function as an anticipatory signal for instability propagation. Embedding variance awareness within the coordination dynamics enables preemptive stabilization rather than a purely reactive compensation.

Despite these theoretical advances, volatility modelling remains largely external to swarm-based industrial coordination mechanisms. Instability is commonly mitigated through heuristic damping or parameter tuning instead of the formal integration of conditional variance signals within the objective structure. This omission limits the ability of swarm systems to achieve structural resilience, defined not merely as post-disturbance recovery, but as the maintenance of bounded dynamic behavior under persistent stochastic excitation.

The present study addresses this architectural fragmentation by introducing a mathematically grounded control framework in which autonomy is formalized as a dynamically adaptable coefficient, sustainability and supervisory workload are embedded within a unified multi-objective functional, and conditional variance estimation directly modulates the coordination dynamics. By integrating these components within a single control structure, the proposed approach moves beyond efficiency-oriented swarm coordination toward a unified socio-technical optimization paradigm, in which performance, stability, and human interaction are jointly regulated.

3. Data and Methodology

Building upon the structural gaps identified in the previous section, the distributed manufacturing environment is modelled as a networked discrete-time nonlinear dynamical system composed of interacting robotic agents subject to bounded stochastic perturbations. The state evolution for each agent $i \in \{1, \dots, N\}$ is described by:

$$x_i(k+1) = f_i(x_i(k), u_i(k)) + d_i(k), \quad (1)$$

where $x_i(k) \in \mathbb{R}^n$ denotes the local operational state vector, $u_i(k) \in \mathbb{R}^m$ the control input, and $d_i(k)$ a bounded disturbance term capturing production variability and communication uncertainty. The nonlinear transition mapping $f_i(\cdot)$ represents the task execution dynamics and resource allocation interactions.

This formulation follows the standard representations of distributed control systems, ensuring that the proposed framework is consistent with the established system-theoretic modelling principles and enabling a formal analysis of the stability and convergence properties.

In contrast to the conventional swarm architectures that assume a fixed autonomy, the present formulation decomposes the control input as:

$$u_i(k) = \alpha_i(k)u_i^{\text{auto}}(k) + (1 - \alpha_i(k))u_i^{\text{sup}}(k), \quad (2)$$

thereby embedding the supervisory intervention explicitly within the control loop. The autonomy coefficient $\alpha_i(k) \in [0, 1]$ constitutes a dynamically regulated parameter modulating the relative contribution of autonomous coordination $u_i^{\text{auto}}(k)$ and supervisory correction $u_i^{\text{sup}}(k)$. This decomposition formalizes autonomy as a

controllable state-dependent variable rather than a static architectural attribute. The formulation also enables the explicit coupling between system state, supervisory workload, and control authority, providing a structured way for also enables the explicit coupling between system state, supervisory workload, and control authority, providing a structured way for regulating human-machine interaction within the control loop instead of relying on heuristic switching between control modes.

In order to align coordination dynamics with the Industry 5.0 requirements, swarm evolution is governed by the multi-objective functional:

$$J(k) = \sum_{i=1}^N [\lambda_1 T_i(k) + \lambda_2 E_i(k) + \lambda_3 \sigma_i^2(k) + \lambda_4 H_i(k)] \quad (3)$$

where $T_i(k)$ denotes task execution latency, $E_i(k)$ energy consumption, conditional operational variance, and $H_i(k)$ supervisory workload contribution. The proposed objective structure extends the classical RSI approaches by internalizing volatility and supervisory workload as endogenous optimization drivers. As a result, convergence is directed toward Pareto-balanced equilibrium rather than single-criterion throughput maximization. This approach is consistent with the multi-objective optimization theory and enables a systematic analysis of trade-offs between competing performance criteria.

The coefficients $\lambda_j \geq 0$ determine the relative weighting of productivity, sustainability, resilience, and human-centric balance.

Operational instability is represented through a conditional variance model of the form:

$$\sigma_i^2(k) = \alpha_0 + \alpha_1 \varepsilon_i^2(k-1) + \beta_1 \sigma_i^2(k-1), \quad (4)$$

where $\varepsilon_i(k)$ denotes the difference between the planned and obtained production output. The parameters $\alpha_0 > 0$, $\alpha_1 > 0$ and $\beta_1 > 0$ govern the baseline variance, shock sensitivity, and persistence effects, respectively. Such a formulation is consistent with the established stochastic modelling approaches, allowing volatility to be treated as a state-dependent quantity rather than as an exogenous disturbance, and it provides a formal mechanism for incorporating instability propagation into the control objective.

Embedding conditional variance within the functional objective introduces anticipatory sensitivity to instability propagation. Rather than reacting ex post to disturbances, the coordination mechanism incorporates predictive volatility signals into its optimization trajectory, thereby enabling preemptive damping adjustments. In this sense, resilience is operationalized as a property of the closed-loop dynamics rather than as a post-disturbance recovery outcome.

The autonomous coordination component is generated through a modified Particle Swarm Optimization (PSO) update:

$$v_i(k+1) = \omega v_i(k) + c_1 r_1 (p_i - x_i(k)) + c_2 r_2 (g - x_i(k)) - \gamma \nabla J(k), \quad (5)$$

$$u_i^{auto}(k) = x_i(k) + v_i(k+1), \quad (6)$$

where $\omega \in (0,1)$ is the inertia coefficient $c_1, c_2 > 0$ are the cognitive and social parameters, $r_1, r_2 \in (0,1)$ are the stochastic weights, and $\nabla J(k)$ introduces sensitivity to the multi-objective functional.

The inclusion of the gradient term modifies the geometry of the search space, constraining particle trajectories within the multi-objective landscape. Consequently, particle trajectories evolve consistently with the defined optimization criteria, rather than being driven by a purely heuristic exploration dynamics.

The autonomy coefficient evolves according to the adaptive update law:

$$\alpha_i(k+1) = \alpha_i(k) - \eta_1 H_i(k) - \eta_2 \sigma_i^2(k) + \eta_3 (\alpha_i^{ref} - \alpha_i(k)), \quad (7)$$

where $\eta_1, \eta_2, \eta_3 > 0$ denote the adaptation gains, and $\alpha_i^{ref} \in [0,1]$ denotes the reference autonomy level. This feedback mechanism directly couples supervisory workload and volatility to autonomy modulation. An elevated workload or variance reduces $\alpha_i(k)$, increasing the supervisory participation, whereas the recovery term drives the autonomy coefficient toward its reference level as operating conditions stabilize. In this way, autonomy is regulated in a closed-loop manner, enabling smooth transitions between control regimes and avoiding the abrupt changes typically associated with rule-based switching strategies. To enable long-term policy refinement, the coordination framework is augmented with reinforcement learning. Defining the reward signal as:

$$R(k) = -J(k), \quad (8)$$

the Q-learning update rule can be expressed as:

$$Q(s, a) \leftarrow Q(s, a) + \gamma [R(k) + \max_{a'} Q(s', a') - Q(s, a)] \quad (9)$$

The reinforcement learning layer operates at a higher abstraction level, enabling the gradual adjustment of weighting coefficients based on the observed system performance. This complements the local convergence properties of PSO with global policy optimization capabilities.

The stability of the closed-loop distributed system is examined via the Lyapunov candidate function:

$$V(k) = \sum_{i=1}^N \|x_i(k) - x^*\|^2 \quad (10)$$

Under bounded disturbances, the inertia constraint $0 < \omega < 1$ the swarm parameter condition $c_1 + c_2 < 4$ and a sufficiently small autonomy variation $|\alpha_i(k+1) - \alpha_i(k)| \leq \delta$, it follows that:

$$E[V(k+1) - V(k)] < 0, \quad (11)$$

which ensures asymptotic convergence in expectation.

These results indicate that the proposed extensions preserve the stability of the underlying swarm dynamics, offering a formal basis for integrating autonomy regulation and volatility-aware coordination within the control framework.

From a computational standpoint, the additional operations introduced by autonomous adaptation and conditional variance estimation scale linearly based on the number of agents. The gradient-based adjustment of the swarm update contributes an $O(N)$ overhead per iteration, preserving the distributed scalability properties of classical PSO implementations. These observations indicate that the proposed framework remains computationally tractable and suitable for the real-time implementation in distributed manufacturing environments.

4. Results

The empirical assessment is focused on the structural impact of the proposed control architecture on distributed manufacturing dynamics. The simulation environment reproduces a network of robotic agents interacting under stochastic task arrivals, bounded execution disturbances, and dynamically regulated supervisory participation.

Such a setup enables a controlled comparison across various coordination strategies under identical stochastic realizations, ensuring that the observed performance differences are attributable to the control architecture rather than to exogenous variability.

The simulated manufacturing network consists of $N = 25$ robotic agents interconnected through a distributed communication topology. Task arrivals follow a Poisson distribution with a rate of $\lambda = 3$ tasks/minute. Processing times are modeled as normally distributed random variables with a mean of 10 s and a standard deviation of 1.5 s. Disturbances are introduced as bounded delay shocks with a magnitude of 20% applied over randomly selected congestion intervals. Parameter values were selected so as to reflect typical operating conditions reported in distributed manufacturing literature, ensuring that the simulation setup remains representative of realistic production environments. PSO parameters were set to $\omega = 0.7$ and $c_1 = c_2 = 1.8$, while the autonomy adaptation gains were set to $\eta_1 = 0.05$, and $\eta_2 = 0.04$. Weighting coefficients λ_j were initialized uniformly and subsequently adjusted via reinforcement learning.

The comparative analysis considers three coordination paradigms implemented under identical stochastic realizations. The reference configuration corresponds to a classical Industry 4.0 RSI architecture characterized by fixed autonomy and efficiency-oriented optimization. A second configuration employs standard PSO without volatility modelling or autonomy adaptation and serves as an intermediate benchmark. The third configuration corresponds to the proposed Industry 5.0-oriented framework. Such a comparative design enables the systematic isolation of the contribution of each structural component, including volatility-aware coordination and adaptive autonomy regulation.

Exposure to identical stochastic sequences ensures that any observed differences arise from the

structural properties of the control architecture rather than stochastic variability. Extended simulation horizons allow the examination of convergence patterns, oscillatory phenomena, and resilience characteristics across multiple independent stochastic realizations. In addition, each scenario was evaluated over multiple simulation runs to ensure the statistical consistency of the reported results.

Nominal operating conditions were first analyzed to determine whether embedding human-centric and sustainability-aware constraints would alter the baseline efficiency properties. The aggregated performance indicators across the three coordination paradigms under nominal operating conditions are rendered in Table 1.

The comparative results reveal that the proposed architecture preserves and slightly improves the task completion performance relative to both reference configurations. While the standard PSO baseline approach already reduces the execution latency compared with fixed-autonomy RSI, the integration of volatility-aware coordination further enhances its efficiency.

Energy consumption decreases progressively across the three configurations, with the proposed model achieving the lowest cumulative energy demand. These improvements are consistent across repeated simulation runs, indicating a stable convergence behavior rather than isolated performance gains. The most pronounced structural differences are observed with regard to operational volatility and supervisory workload. The proposed architecture reduces the conditional variance by approximately 39% relative to Industry 4.0 RSI and by 30% relative to standard PSO. Supervisory workload exhibits an even larger reduction, indicating that autonomy modulation effectively stabilizes human intervention requirements even under nominal operating conditions.

Table 1. Performance comparison under nominal operating conditions (mean $\pm$ standard deviation)

Metric	Industry 4.0 RSI	Standard PSO	Proposed Architecture
Task Completion Time (s)	12.6 $\pm$ 1.4	12.3 $\pm$ 1.3	11.9 $\pm$ 1.2
Energy Consumption (kWh)	5.84 $\pm$ 0.32	5.41 $\pm$ 0.30	4.91 $\pm$ 0.28
Operational Volatility	0.31 $\pm$ 0.05	0.27 $\pm$ 0.05	0.19 $\pm$ 0.04
Human Workload Index	0.62 $\pm$ 0.07	0.58 $\pm$ 0.06	0.38 $\pm$ 0.06

The structural advantages of the proposed coordination mechanism become more evident under controlled perturbations introduced as bounded processing delays and congestion spikes. Table 2 summarizes the comparative performance across the three coordination paradigms.

Under disturbance conditions, the performance gap across the three configurations widens substantially. The recovery time achieved by the proposed approach features a decrease of approximately 39% relative to Industry 4.0 RSI and of 25% relative to standard PSO, indicating an enhanced disturbance absorption capacity. These improvements are consistently observed across multiple stochastic realizations, confirming that the observed performance gains are structural rather than scenario-specific.

Energy spikes are significantly attenuated in the proposed architecture, reflecting an anticipatory adjustment of swarm trajectories rather than a delayed compensatory reallocation. Also, operational volatility is reduced by nearly 43% relative to the fixed-autonomy RSI baseline.

Most notably, when employing the proposed approach, the supervisory workload remains consistently bounded and substantially lower compared with both reference configurations. While the standard PSO improves the mechanical redistribution efficiency, it does not regulate human intervention dynamics. The proposed framework uniquely stabilizes both the technical and human subsystems, confirming the structural contribution of autonomy modulation.

In order to evaluate the robustness of the proposed multi-objective formulation with respect to parameter selection, a structured sensitivity analysis was conducted on the weighting coefficient λ_j governing the functional objective. Each coefficient was independently perturbed within a $\pm 20\%$ interval around its nominal calibrated value, while all the remaining parameters were held constant (Table 3).

The objective of this analysis was to assess whether the qualitative performance improvements observed under nominal calibration persist under moderate parameter deviations, thereby confirming a structural rather than incidental behavior.

The results indicate that moderate perturbations of the weighting coefficients do not alter the qualitative behavior of the system. In particular, the results related to volatility damping and supervisory workload stabilization remain consistently superior to those obtained by the baseline configurations across the tested parameter intervals. An increased emphasis on variance reduction (λ_3) further enhances the damping characteristics, whereas the excessive prioritization of workload minimization (λ_4) induces marginal convergence slowdowns without destabilizing the system. These observations suggest that the proposed architecture exhibits structural robustness with respect to reasonable calibration deviations.

The aggregated indicators quantify the overall performance outcomes but do not fully reveal how the system evolves over time. Examining the temporal trajectories therefore provides

Table 2. Performance comparison under disturbance conditions (mean $\pm$ standard deviation)

Metric	Industry 4.0 RSI	Standard PSO	Proposed Architecture
Recovery Time (s)	14.2 $\pm$ 2.3	11.6 $\pm$ 2.0	8.7 $\pm$ 1.8
Energy Spike (%)	+21%	+15%	+8%
Operational Volatility	0.37 $\pm$ 0.06	0.29 $\pm$ 0.05	0.21 $\pm$ 0.05
Human Workload Index	0.78 $\pm$ 0.09	0.63 $\pm$ 0.08	0.42 $\pm$ 0.07

Table 3. Sensitivity analysis for the weighting coefficients ($\pm 20\%$ perturbation)

Parameter	Variation	Impact on Volatility	Impact on Workload	Impact on Convergence Speed
λ_1 (Execution Time)	$\pm 20\%$	Minor	Negligible	Moderate
λ_2 (Energy)	$\pm 20\%$	Moderate	Minor	Minor
λ_3 (Variance)	$\pm 20\%$	Significant	Moderate	Minor
λ_4 (Workload)	$\pm 20\%$	Moderate	Significant	Minor

additional insight into the observed performance differences. Figure 1 shows the cumulative energy-consumption trajectories for the baseline Industry 4.0 RSI configuration and the proposed Industry 5.0 architecture.




Figure 1. Cumulative energy consumption over the simulation horizon

Figure 1 shows a consistently lower cumulative energy-consumption trajectory for the proposed Industry 5.0 architecture compared with the baseline Industry 4.0 RSI configuration. The separation between the two trajectories increases gradually over the simulation horizon, indicating a lower cumulative energy demand under the proposed architecture. However, both trajectories remain smooth and monotonic; therefore, the main difference observed in this figure concerns the level of cumulative energy consumption rather than oscillatory behavior. The dynamic behavior of the autonomy coefficient is illustrated in Figure 2.




Figure 2. Dynamic evolution of the autonomy coefficient

The average autonomy coefficient varies continuously over the simulation horizon, decreasing gradually from its initial level to a minimum around the middle of the simulation horizon and subsequently recovering toward

higher values. This smooth evolution demonstrates that autonomy is dynamically modulated rather than maintained at a fixed level. The absence of abrupt discontinuities is consistent with the continuous adaptive regulation mechanism defined by the autonomy update law, rather than with discrete switching between control modes.

Understanding the resilience properties of the coordination mechanism requires examining how instability propagates and dissipates following perturbations. The conditional production variance, which captures the magnitude and persistence of execution deviations, provides a direct measure of this phenomenon. Its temporal evolution under both coordination paradigms is depicted in Figure 3.




Figure 3. Operational volatility dynamics

The volatility trajectory associated with the baseline Industry 4.0 RSI configuration exhibits a pronounced transient increase around the middle of the simulation horizon, followed by a rapid return toward the pre-disturbance level. Under the proposed Industry 5.0 architecture, the corresponding transient response is substantially smaller in amplitude, while the background fluctuations before and after the event remain comparatively limited.

The comparison therefore indicates that the main effect of the proposed architecture is the attenuation of disturbance-induced volatility rather than the elimination of variability. The lower transient peak is consistent with the improved damping of the volatility response and with the lower operational volatility levels reported for the proposed architecture in the comparative analysis.

While volatility captures technical instability, the supervisory workload reflects the human dimension of systemic stress. The temporal profile

of the supervisory intervention intensity under both architectures is shown in Figure 4.




Figure 4. Human workload dynamics

Figure 4 shows that the baseline Industry 4.0 RSI configuration maintains a higher supervisory workload level throughout the simulation horizon and exhibits a localized transient increase during the latter part of the simulation. In contrast, the proposed Industry 5.0 architecture operates at a consistently lower workload level and displays a smaller transient increase.

The persistent separation between the two trajectories indicates a substantial reduction in supervisory demand under the proposed architecture, both during normal operation and during the transient workload increase. This pattern is consistent with the lower Human Workload Index reported for the proposed architecture in the comparative results and supports the intended role of adaptive autonomy regulation in limiting supervisory burden.

The separation of volatility and workload dynamics reveals a critical structural effect: the improvements are not confined to the mechanical damping of production deviations but extend to the stabilization of the human supervisory layer. The modified swarm architecture therefore alters the transient response characteristics of the distributed manufacturing system at both the technical and socio-technical levels.

Paired two-tailed statistical tests performed across 30 independent simulation runs confirm that the reductions in energy consumption, operational volatility, and supervisory workload are statistically significant at confidence levels exceeding 99% ($p < 0.01$). Improvements in

recovery time under disturbance conditions remain significant at the 95% confidence level ($p < 0.05$), while low dispersion across the simulation runs indicates a stable convergence behavior. Effect sizes were further evaluated using Cohen's d , with values exceeding 0.8 for both energy consumption reduction and supervisory workload stabilization. Overall, the results demonstrate that the observed improvements are both statistically robust and practically meaningful.

5. Discussion

The simulation findings indicate that extending swarm-based coordination toward a human-centric and sustainability-aware architecture yields structurally consistent improvements with regard to the distributed manufacturing performance. Unlike the efficiency-driven Industry 4.0 swarm architectures, the proposed framework embeds supervisory workload and operational volatility directly within the functional objective, thereby reclassifying them from exogenous performance indicators to endogenous control variables that continuously influence coordination dynamics.

From a control-theoretic standpoint, the introduction of the adaptive autonomy coefficient $\alpha_i(k)$ constitutes a structural generalization of the classical swarm paradigm. Traditional architectures implicitly assume fixed autonomy levels oriented toward algorithmic dominance. By contrast, the proposed framework regulates autonomy dynamically through feedback coupling with volatility and workload indicators. The simulation results indicate that this adaptive mechanism is associated with reduced supervisory demand under transient operating conditions, as illustrated by Figure 4, while the disturbance analysis reports a shorter recovery time for the proposed architecture. Convergence properties remain preserved, indicating that hybrid autonomy regulation enhances robustness without destabilizing the swarm dynamics.

The incorporation of conditional volatility modelling within the production dynamics is particularly relevant under disturbance conditions. By embedding variance estimation directly into the optimization process, instability information is incorporated at the level of coordination

updates rather than being considered only as an ex-post performance indicator. Figure 3 shows a substantially lower transient volatility peak for the proposed architecture, while both trajectories return toward their pre-disturbance levels after the transient event. The reduced peak amplitude, together with the lower operational-volatility values reported in the comparative analysis, is consistent with an improved disturbance damping and reduced volatility amplification.

Energy efficiency improvements observed when energy consumption and human workload are explicitly incorporated into the optimization objective further demonstrate that sustainability-oriented optimization does not inherently conflict with productivity objectives. Although the proposed objective function includes additional penalty components associated with energy consumption and human workload, task completion time remains comparable and, in some cases, is slightly improved relative to the Industry 4.0 configuration (Petrescu et al., 2025). The simulation results show that reductions in energy consumption are achieved without compromising the execution performance. Together, these findings confirm that multi-objective optimization in distributed swarm systems can achieve balanced trade-offs without introducing an excessive computational complexity.

The reinforcement learning layer complements the swarm-based local convergence with higher-level policy adaptation. While PSO ensures a rapid decentralized coordination, reinforcement learning enables the gradual refinement of the weighting coefficient in response to evolving environmental conditions. This dual-layer adaptation mechanism is reflected in the stability of the performance metrics across repeated simulation runs, indicating that both short-term responsiveness and long-term policy consistency are maintained in non-stationary operational regimes.

Despite the demonstrated performance improvements, several limitations must be acknowledged. First, the validation of the proposed architecture is conducted exclusively through simulation-based experiments. While the simulation environment incorporates stochastic disturbances and representative operational parameters, it does not explicitly capture real-world

constraints such as hardware-induced latency, communication loss, or synchronization issues characteristic of industrial-scale deployments.

Second, scalability to large-scale manufacturing systems may introduce additional coordination complexities, including communication overhead and network congestion, which are not explicitly modeled in the present framework. Moreover, the optimal calibration of the weighting coefficient remains context-dependent and may require domain-specific tuning. Accordingly, the reported results should be interpreted as a proof-of-concept validation rather than as direct evidence of industrial-scale performance. Future research should therefore pursue hardware-in-the-loop experimentation and industrial pilot validation to assess the scalability of the proposed framework under real-world conditions.

The results demonstrate that transitioning from automation-centric RSI architectures toward human-centric swarm control does not compromise efficiency. Instead, it introduces structural stability, sustainability alignment, and an enhanced resilience consistent with Industry 5.0 principles. These conclusions are supported by the simulation-based evidence, which demonstrates simultaneous improvements in technical performance and human workload stabilization. By internalizing collaborative intelligence within the optimization process, the proposed architecture contributes to the evolution of distributed manufacturing systems beyond the traditional Industry 4.0 paradigm.

Future research will focus on real-world validation, including hardware-in-the-loop experimentation and industrial deployment scenarios. Additional directions include robustness under communication constraints, asynchronous coordination mechanisms, and extensions of the stochastic model toward non-stationary environments.

6. Conclusions

The transition toward Industry 5.0 requires more than incremental improvements of the existing automation paradigms. It demands a structural reconsideration of how distributed manufacturing systems are modeled, controlled, and evaluated. Within this context, the present study has

addressed a central challenge: how to extend swarm-based coordination mechanisms beyond efficiency-driven optimization and embed human-centric and sustainability-oriented principles directly into the control architecture.

The proposed framework shows that swarm intelligence can evolve from a purely decentralized task allocation strategy into a collaborative control paradigm in which human supervision, energy awareness, and operational stability are treated as endogenous components of the optimization problem. By introducing a dynamic autonomy coefficient, the architecture redefines the relationship between human operators and robotic agents. Autonomy is no longer a fixed structural attribute but a regulated variable capable of adapting to contextual conditions such as volatility escalation or supervisory overload.

The integration of predictive volatility modelling within the production flow contributes to systemic resilience by incorporating variance information directly into the coordination process. The simulation results show a substantially lower transient volatility response for the proposed architecture, as illustrated in Figure 3, while the disturbance analysis reports a shorter recovery time compared with the reference configurations.

Energy-aware optimization, embedded within the multi-objective functional, confirms that

sustainability constraints do not inherently conflict with the performance objectives. The simulation results show that a balanced objective weighting can simultaneously reduce energy consumption and stabilize the execution dynamics. These results support the view that Industry 5.0 principles can be operationalized within distributed control systems rather than remain purely normative.

From a systems perspective, the hybridization of swarm coordination and reinforcement learning introduces a dual adaptation mechanism. The swarm layer ensures a rapid decentralized convergence, while the learning layer enables the gradual refinement of the policy parameters in response to environmental variability. This layered adaptability supports both short-term responsiveness and a long-term stable equilibrium.

This research validates a mathematically grounded and simulation-based control architecture that aligns swarm intelligence with the human-centric and sustainability-oriented imperatives of Industry 5.0. By internalizing the human workload, energy consumption, and operational volatility into the swarm optimization process, the proposed model advances distributed manufacturing systems toward a more resilient, adaptive, and collaborative form of industrial intelligence. However, further validation under real-world industrial conditions remains necessary in order to assess the practical applicability of the proposed framework.

REFERENCES

- Akopov, A. S. (2024) A Clustering-Based Hybrid Particle Swarm Optimization Algorithm for Solving a Multisectoral Agent-Based Model. *Studies in Informatics and Control*. 33(2), 83-95. <https://doi.org/10.24846/v33i2y202408>.
- Ban, L. (2024) Optimization of Energy Management Algorithm for Hybrid Power Systems Based on Deep Reinforcement Learning. *Studies in Informatics and Control*. 33(2), 15-25. <https://doi.org/10.24846/v33i2y202402>.
- Brambilla, M., Ferrante, E., Birattari, M. et al. (2013) Swarm robotics: A review from the swarm engineering perspective. *Swarm Intelligence*. 7(1), 1–41. <https://doi.org/10.1007/s11721-012-0075-2>.
- Chu, H., Zhang, T., Zou, Y. et al. (2024) “Follower” to “Collaborator”: A robot proactive collaborative controller based on human multimodal information for 3D handling/assembly scenarios. *Journal of Manufacturing Systems*. 73, 287–306. <https://doi.org/10.1016/j.jmsy.2024.02.008>.
- Dias, P. G. F., Silva, M. C., Rocha Filho et al. (2021) Swarm Robotics: A Perspective on the Latest Reviewed Concepts and Applications. *Sensors*. 21(6), Art. No. 2062. <https://doi.org/10.3390/s21062062>.
- Dorigo, M., Theraulaz, G. & Trianni, V. (2020) Reflections on the future of swarm robotics. *Science Robotics*. 5(49), Art. No. eabe4385. <https://doi.org/10.1126/scirobotics.abe4385>.
- Kuckling, J. (2023) Recent trends in robot learning and evolution for swarm robotics. *Frontiers in Robotics and AI*. 10, Art. No. 1134841. <https://doi.org/10.3389/frobt.2023.1134841>.
- Ligot, A. & Birattari, M. (2020) Simulation-only experiments to mimic the effects of the reality gap in the automatic design of robot swarms. *Swarm*

- Intelligence*. 14, 1–24. <https://doi.org/10.1007/s11721-019-00175-w>.
- Maddikunta, P. K. R., Pham, Q.-V., Prabadevi, B et al. (2022) Industry 5.0: A survey on enabling technologies and potential applications. *Journal of Industrial Information Integration*. 26, Art. No. 100257. <https://doi.org/10.1016/j.jii.2021.100257>.
- Oncioiu, I. (2023) Predicting the Use of Chatbots for Consumer Channel Selection in Multichannel Environments: An Exploratory Study. *Systems*. 11(10), Art. No. 522. <https://doi.org/10.3390/systems11100522>.
- Petrescu, A.-G., Oncioiu, I., Petrescu, M. et al. (2025) Integration of Robot Swarm Intelligence for Distributed Manufacturing Systems. *Studies in Informatics and Control*. 34(1), 37–46. <https://doi.org/10.24846/v34i1y202503>.
- Ribeiro, J., Lima, R., Eckhardt, T. et al. (2021) Robotic Process Automation and Artificial Intelligence in Industry 4.0 — A Literature review. *Procedia Computer Science*. 181, 51–58. <https://doi.org/10.1016/j.procs.2021.01.104>.
- Schranz, M., Umlauf, M., Sende, M. et al. (2020) Swarm Robotic Behaviors and Current Applications. *Frontiers in Robotics and AI*. 7, Art. No. 36. <https://doi.org/10.3389/frobt.2020.00036>.
- Trianni, V. & Nolfi, S. (2011) Engineering the Evolution of Self-Organizing Behaviors in Swarm Robotics: A Case Study. *Artificial Life*. 17(3), 183–202. https://doi.org/10.1162/artl_a_00031.
- Wang, J., Shi, E., Hu, H. et al. (2025) Large Language Models for Robotics: Opportunities, Challenges, and Perspectives. *Journal of Automation and Intelligence*. 4, 52–64. <https://doi.org/10.1016/j.jai.2024.12.003>.
- Wang, W., Zheng, Y., Lin, G. et al. (2020). Swarm robotics: A review. *Robot*. 42(2), 232–256. <https://robot.sia.cn/en/article/doi/10.13973/j.cnki.robot.190009>.
- Xu, X., Lu, Y., Vogel-Heuser, B. et al. (2021) Industry 4.0 and Industry 5.0 - Inception, conception and perception. *Journal of Manufacturing Systems*. 61, 530–535. <https://doi.org/10.1016/j.jmsy.2021.10.006>.
- Yang, H., Cao, S., Bai, L. et al. (2019) A distributed and parallel self-assembly approach for swarm robotics. *Robotics and Autonomous Systems*. 118, 80–92. <https://doi.org/10.1016/j.robot.2019.04.011>.



This is an open access article distributed under the terms and conditions of the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License.