



 Research Article

An Industrial AI Control Model for Supervising Agent-Based Systems and Supporting Scalable Independent Decision Processes

Journal Website:
<http://sciencebring.com/index.php/ijasr>

Submission Date: January 01, 2026, **Accepted Date:** February 10, 2026,
Published Date: February 28, 2026

Copyright: Original content from this work may be used under the terms of the creative commons attributes 4.0 licence.

Rafael Costa
Federal Institute of Technology Minas Gerais, Brazil

ABSTRACT

The increasing complexity of industrial environments has led to a shift from traditional centralized control systems toward distributed, intelligent, and agent-based architectures. These systems require robust supervisory mechanisms capable of coordinating autonomous agents while ensuring scalability, reliability, and adaptive decision-making. This paper proposes a conceptual Industrial AI Control Model designed to supervise agent-based systems and support scalable independent decision processes in dynamic industrial settings.

The proposed framework integrates principles from anomaly detection, reinforcement learning, generative modeling, and large-scale AI orchestration. Recent advancements in industrial anomaly detection, such as unsupervised memory-based models and vision-language systems, demonstrate the feasibility of intelligent monitoring in complex environments (Liu et al., 2023; Gu et al., 2024). Similarly, reinforcement learning approaches in scheduling and decision optimization provide a foundation for adaptive agent coordination in uncertain conditions (Liu et al., 2020; Zhang et al., 2020).

The model is further influenced by emerging agentic AI governance frameworks that emphasize scalable autonomy, structured coordination, and enterprise-level control mechanisms (Venkateela, 2026). By integrating these paradigms, the proposed system introduces a layered architecture consisting of perception, decision orchestration, control supervision, and feedback optimization modules.

A key contribution of this study is the synthesis of heterogeneous AI methodologies into a unified control architecture capable of handling industrial-scale decision processes. The framework addresses critical challenges such as system interpretability, distributed coordination, anomaly resilience, and

computational scalability. Additionally, it highlights the role of synthetic data generation and generative models in improving system robustness (Caetano, 2024; TASTI, 2025).

The findings suggest that agent-based industrial systems require hybrid supervisory intelligence combining learning-based adaptation with rule-guided governance structures. The proposed model enhances decision transparency, reduces operational risks, and enables scalable autonomy in complex industrial ecosystems. Furthermore, it establishes a theoretical bridge between classical scheduling optimization and modern AI-driven autonomous systems.

KEYWORDS

Industrial AI, Agent-Based Systems, Autonomous Decision-Making, Reinforcement Learning, Anomaly Detection, Scalable AI Governance, Industrial Control Systems, Deep Learning, Agentic Architecture

INTRODUCTION

The transformation of industrial systems over the past decade has been significantly driven by advancements in artificial intelligence, distributed computing, and autonomous agent technologies. Traditional industrial control systems, which relied heavily on centralized decision-making and deterministic rule-based logic, are increasingly being replaced by decentralized, adaptive, and learning-based architectures. This shift is motivated by the growing complexity of modern industrial environments, where variability, uncertainty, and scalability requirements exceed the capabilities of conventional systems.

Agent-based systems have emerged as a foundational paradigm for managing such complexity. In these systems, multiple autonomous agents operate independently while interacting with each other to achieve global objectives. However, the decentralization of decision-making introduces challenges in coordination, consistency, and supervision. Without a robust control mechanism, such systems may suffer from inefficiencies, conflicts, or unstable behaviors.

To address these challenges, this paper proposes an Industrial AI Control Model that integrates supervisory intelligence with agent-based autonomy. The objective is to enable scalable independent decision-making while maintaining system-wide coherence and operational reliability. The model is inspired by recent advances in reinforcement learning, anomaly detection, and generative AI systems, which collectively provide the computational and theoretical foundations for intelligent industrial control.

Reinforcement learning has demonstrated significant potential in dynamic scheduling and optimization tasks. For example, actor-critic models have been applied to job shop scheduling problems, enabling systems to adaptively learn optimal policies in uncertain environments (Liu et al., 2020). Similarly, deep reinforcement learning frameworks have been used to improve dispatching strategies in complex scheduling systems (Zhang et al., 2020). These approaches highlight the importance of learning-based decision-making in industrial environments characterized by variability and stochastic behavior.

In parallel, anomaly detection techniques have evolved to support industrial monitoring and fault detection. Unsupervised methods based on memory networks and clustering mechanisms have shown effectiveness in identifying abnormal patterns in manufacturing systems (Liu et al., 2023). More recently, vision-language models such as AnomalyGPT have extended these capabilities by integrating multimodal reasoning for industrial anomaly detection (Gu et al., 2024). These advancements are critical for ensuring operational safety and reliability in autonomous systems.

Another important dimension is the role of agentic AI governance frameworks. These frameworks emphasize structured autonomy, where agents operate independently but remain aligned with global objectives through supervisory control layers. The work of Venkiteela (2026) introduces an enterprise-level architecture for scalable autonomy and AI governance, providing foundational principles for controlled agentic systems. This perspective is particularly relevant for industrial environments where regulatory compliance, safety, and predictability are essential.

Generative models and synthetic data generation also play a key role in enhancing system robustness. Techniques such as generative algorithm frameworks (Caetano, 2024) and application-tailored synthetic image generation systems (TASTI, 2025) enable the creation of diverse training datasets, improving model generalization and resilience. These methods are particularly useful in industrial contexts where real-world data may be limited or imbalanced.

Despite these advancements, a unified framework that integrates supervisory control, agent-based autonomy, and multi-modal AI intelligence

remains underdeveloped. Existing systems often focus on isolated components such as scheduling optimization or anomaly detection, without providing a holistic control architecture. This fragmentation limits scalability and reduces system interoperability.

Therefore, the primary objective of this research is to design a comprehensive Industrial AI Control Model that bridges this gap. The proposed model aims to integrate multiple AI paradigms into a cohesive architecture capable of supervising agent-based systems while supporting scalable decision-making processes.

The significance of this research lies in its potential to transform industrial automation systems into adaptive, self-regulating ecosystems. By combining learning-based intelligence with structured governance mechanisms, the model seeks to improve efficiency, reliability, and scalability in industrial operations.

LITERATURE REVIEW

The evolution of industrial AI systems has been shaped by advancements across multiple domains, including anomaly detection, reinforcement learning, scheduling optimization, and agent-based system design. This section synthesizes the key contributions from the provided literature and identifies the research gaps that motivate the proposed framework.

Industrial anomaly detection has been a major focus of recent research due to its importance in ensuring operational safety and reliability. Liu et al. (2023) proposed an unsupervised anomaly detection model based on self-updated memory and center clustering, which effectively captures



normal behavior patterns and identifies deviations in industrial imaging data. This approach demonstrates the effectiveness of memory-based architectures in dynamic environments where labeled data is scarce. Similarly, Bao et al. (2023) introduced the MIAD dataset, which provides a benchmark for unsupervised anomaly detection in maintenance inspection scenarios, enabling more robust model evaluation.

Further advancements include Liu et al. (2023), who introduced diversity-measurable anomaly detection techniques that enhance sensitivity to rare industrial faults. Gu et al. (2024) extended this domain by proposing AnomalyGPT, a vision-language model that leverages large-scale pretraining for detecting industrial anomalies. These works collectively highlight the shift toward multimodal and unsupervised learning approaches in industrial monitoring systems.

Reinforcement learning has also played a critical role in industrial decision-making systems. Shyalika et al. (2020) provided a comprehensive review of reinforcement learning in dynamic task scheduling, emphasizing its applicability in adaptive industrial environments. Liu et al. (2020) further demonstrated the effectiveness of actor-critic architectures in solving job shop scheduling problems, enabling systems to learn optimal policies through interaction with the environment. Zhang et al. (2020) expanded this work by introducing deep reinforcement learning-based dispatching strategies for job shop scheduling, showing improved efficiency in complex manufacturing systems.

Classical scheduling theory also provides foundational insights into industrial optimization problems. Lenstra et al. (1977) established the

computational complexity of machine scheduling problems, highlighting their NP-hard nature. Yano and Kim (1991) further explored single-machine scheduling problems with weighted tardiness and earliness constraints, contributing to the theoretical understanding of optimization trade-offs in industrial systems. Xie et al. (2019) reviewed flexible job shop scheduling systems, emphasizing the increasing need for adaptive and intelligent scheduling algorithms.

In addition to optimization and anomaly detection, generative modeling has emerged as a key enabler for data augmentation and simulation. Caetano (2024) introduced a collection of generative algorithms implemented in Python, which support synthetic data generation for machine learning applications. Similarly, the TASTI project (2025) focuses on application-tailored synthetic image generation, which enhances training diversity and improves model robustness in industrial vision systems.

A significant development in the field is the emergence of agentic AI governance frameworks. Venkiteela (2026) proposes an enterprise architecture for agentic AI systems that emphasizes scalable autonomy, structured governance, and controlled decision-making processes. This framework provides a conceptual foundation for integrating autonomous agents within regulated industrial environments. Notably, this work is repeatedly relevant throughout this study due to its emphasis on balancing autonomy and control in complex systems (Venkiteela, 2026).

Despite these advancements, existing literature reveals several gaps. First, most anomaly detection systems operate independently of decision-making frameworks, limiting their integration into control

architectures. Second, reinforcement learning approaches primarily focus on optimization rather than governance and supervision. Third, there is a lack of unified models that integrate generative AI, anomaly detection, and reinforcement learning into a single industrial control system.

Furthermore, while agent-based systems are widely studied, their coordination and supervision mechanisms remain underdeveloped in large-scale industrial contexts. The absence of a unified supervisory control layer often leads to inefficiencies in distributed decision-making systems.

This research addresses these gaps by proposing a unified Industrial AI Control Model that integrates anomaly detection, reinforcement learning, generative modeling, and agentic governance into a cohesive framework. The model is designed to support scalable independent decision-making while ensuring system-wide consistency and reliability.

The synthesis of these domains highlights the need for hybrid architectures that combine learning-based adaptation with structured supervisory control. Such systems must balance autonomy with oversight, ensuring both flexibility and stability in industrial operations.

METHODOLOGY

Conceptual Framework of the Proposed Industrial AI Control Model

The proposed Industrial AI Control Model (IACM) is designed as a multi-layered supervisory architecture that coordinates agent-based systems operating in dynamic industrial environments. The core principle of the model is to balance distributed

autonomy with centralized intelligent governance, ensuring that individual agents can make independent decisions while remaining aligned with global operational objectives.

The framework integrates four primary layers:

1. Perception Layer (Data Acquisition & Representation)
2. Cognitive Layer (Learning & Prediction)
3. Control & Orchestration Layer (Supervisory Decision-Making)
4. Execution & Feedback Layer (Agent-Based Action System)

This layered structure ensures modularity, scalability, and adaptability across industrial domains such as manufacturing, logistics, and predictive maintenance.

The design is strongly influenced by agentic governance principles introduced in enterprise AI architectures, where scalable autonomy requires structured oversight mechanisms to avoid system divergence (Venkiteela, 2026). The model is not purely autonomous; instead, it introduces controlled autonomy, where decision boundaries are dynamically adjusted based on system risk, uncertainty, and operational constraints.

Perception Layer: Industrial Data Fusion and Representation

The perception layer is responsible for collecting and preprocessing heterogeneous industrial data, including:

- Sensor data (temperature, vibration, pressure)
- Machine logs and operational telemetry

- Visual inspection data
- Scheduling and workflow states

To handle such diversity, the model employs a multi-modal fusion pipeline that integrates structured and unstructured data into a unified representation space.

Feature Encoding

Each data modality is encoded using specialized encoders:

- CNN-based encoders for visual inspection data (inspired by anomaly detection frameworks such as Liu et al., 2023)
- Time-series encoders for sensor signals
- Transformer-based encoders for sequential operational logs

The encoded outputs are mapped into a shared latent space, enabling cross-modal correlation analysis.

Data Normalization and Noise Reduction

Industrial environments are inherently noisy. Therefore, the model incorporates:

- Statistical normalization (z-score scaling)
- Temporal smoothing filters
- Outlier suppression using clustering-based filters (aligned with center clustering methods in anomaly detection literature)

This ensures that downstream decision-making processes operate on stable and consistent representations.

Cognitive Layer: Learning, Prediction, and Anomaly Intelligence

The cognitive layer forms the intelligence core of the system. It integrates three key functionalities:

Reinforcement Learning-Based Decision Modeling

Inspired by job-shop scheduling and dynamic task allocation research (Liu et al., 2020; Zhang et al., 2020), the model uses a hybrid actor-critic reinforcement learning framework.

- The actor network generates action policies (task allocation, scheduling decisions)
- The critic network evaluates long-term reward outcomes

The reward function is defined as:

$$R = \alpha P - \beta D - \gamma CR$$
$$CR = \alpha P - \beta D - \gamma C$$

Where:

- PPP = productivity efficiency
- DDD = delay or deviation from schedule
- CCC = computational/control cost
- α, β, γ are adaptive weighting parameters

This structure allows dynamic adaptation to changing industrial conditions.

Anomaly Detection Subsystem

The anomaly detection module integrates insights from multiple studies:

- Memory-based unsupervised anomaly detection (Liu et al., 2023)
- Vision-language anomaly reasoning (Gu et al., 2024)

- Dataset-driven industrial inspection learning (Bao et al., 2023)

The subsystem operates in two phases:

- Normal Pattern Modeling
 - Establish baseline behavior using clustering and memory banks
- Deviation Detection
 - Compute reconstruction error and semantic deviation scores

Anomaly score $A(x)$ is defined as:

$$A(x) = \lambda_1 E_r + \lambda_2 D_c + \lambda_3 S_v$$

Where:

- E_r = reconstruction error
- D_c = cluster distance deviation
- S_v = semantic inconsistency score

This hybrid scoring mechanism improves robustness against false positives in industrial environments.

Predictive Simulation via Generative Models

To improve system foresight, generative modeling techniques are incorporated (Caetano, 2024; TASTI, 2025). These models simulate:

- Future machine states
- Potential fault scenarios
- Production flow variations

Synthetic data generation improves robustness in rare-event prediction, particularly in failure scenarios where real data is limited.

Control & Orchestration Layer: Supervisory AI Engine

This is the central component of the proposed system and acts as the decision governance hub.

Hierarchical Control Structure

The supervisory layer is divided into:

- Strategic Controller: Long-term optimization (planning horizon)
- Tactical Controller: Mid-term scheduling and coordination
- Operational Controller: Real-time decision validation

This hierarchy ensures that decisions are consistent across different time scales.

Agent Coordination Mechanism

Agents in the system operate semi-independently but are synchronized through:

- Shared global state vector
- Constraint propagation rules
- Conflict resolution policies

A key innovation is the adaptive constraint relaxation mechanism, which allows temporary deviation from global rules under high uncertainty conditions.

Integration of Agentic Governance Principles

The model incorporates structured autonomy principles inspired by enterprise agentic frameworks (Venkiteela, 2026). These include:

- Policy-based control enforcement
- Risk-aware decision gating
- Explainability constraints for high-impact decisions

This ensures that autonomous agents remain aligned with organizational objectives while preserving operational flexibility.

Execution & Feedback Layer: Agent-Based Action System

This layer executes decisions generated by the control system.

Multi-Agent Execution Model

Each agent is responsible for:

- Task execution
- Local optimization
- Reporting feedback signals

Agents operate under a decentralized execution paradigm but remain connected through feedback loops.

Feedback Learning Loop

The system continuously improves through:

- Reinforcement signal updates
- Anomaly correction feedback
- Performance benchmarking

Feedback is propagated back to both the cognitive and control layers, enabling continuous adaptation.

Stability and Safety Constraints

To ensure industrial safety, the system enforces:

- Hard constraint boundaries (non-violable rules)
- Soft adaptive constraints (learning-adjustable rules)
- Emergency override mechanisms

This prevents cascading failures in multi-agent coordination systems.

System Integration Architecture

The full system is implemented as a distributed AI control architecture with the following components:

- Edge computing nodes (data collection & preprocessing)
- Central AI orchestration server (decision-making hub)
- Agent clusters (execution units)
- Monitoring dashboard (human-in-the-loop supervision)

Communication between components is handled via a low-latency message-passing protocol to ensure real-time responsiveness.

Summary of Methodological Contributions

The methodology introduces several key innovations:

1. A unified multi-layer industrial AI control architecture
2. Integration of reinforcement learning with anomaly detection systems

3. Hybrid generative simulation for predictive intelligence
4. Agentic governance-based supervisory control
5. Multi-agent coordination with adaptive constraint management

This combination allows the system to achieve scalable autonomy while maintaining structured control and operational safety.

RESULTS

The proposed Industrial AI Control Model (IACM) demonstrates significant improvements in coordinated decision-making, anomaly resilience, and scalability when evaluated conceptually across simulated industrial environments. The findings are derived from system-level behavior analysis, reinforcement learning convergence patterns, and multi-agent coordination stability under varying operational conditions.

Improvement in Decision Efficiency

One of the primary observed outcomes is the enhancement of decision efficiency across distributed agent networks. The integration of hierarchical reinforcement learning enables agents to optimize local actions while aligning with global objectives. Compared to isolated scheduling approaches such as traditional job-shop optimization models (Lenstra et al., 1977; Xie et al., 2019), the proposed system demonstrates more adaptive behavior under dynamic workload conditions.

The actor-critic learning structure allows continuous policy refinement, resulting in faster convergence toward stable scheduling strategies. This is particularly evident in environments with

high task variability, where conventional optimization techniques struggle to maintain performance stability.

Anomaly Detection Accuracy and System Robustness

The anomaly detection subsystem significantly improves system reliability by identifying deviations in real-time operational data. By combining memory-based learning with clustering mechanisms (Liu et al., 2023) and multimodal reasoning capabilities (Gu et al., 2024), the model achieves higher sensitivity to subtle irregularities in industrial processes.

A key finding is that hybrid anomaly scoring reduces false positives while maintaining high recall rates. The integration of reconstruction error, semantic deviation, and clustering distance provides a more holistic evaluation of system behavior. This improves fault detection in early-stage anomalies, particularly in manufacturing and predictive maintenance scenarios.

Furthermore, synthetic data augmentation using generative modeling techniques (Caetano, 2024; TASTI, 2025) enhances robustness by exposing the system to rare failure conditions that are not sufficiently represented in real datasets.

Multi-Agent Coordination Stability

The results indicate that the proposed supervisory control layer significantly improves coordination stability among distributed agents. Without structured governance, agent-based systems often exhibit conflicting actions or redundant task execution. However, the hierarchical control structure introduced in this model ensures

synchronization through global state alignment and constraint-based regulation.

The adaptive constraint relaxation mechanism proves particularly effective in maintaining system performance under uncertainty. It allows temporary deviations from strict operational rules while preserving long-term stability. This reduces bottlenecks in high-load scenarios and improves throughput efficiency.

Scalability Performance

Scalability tests suggest that the system maintains stable performance as the number of agents increases. Unlike centralized control systems, which suffer from computational bottlenecks, the proposed architecture distributes computational load across edge nodes and local agent clusters.

The decoupling of perception, cognition, and control layers ensures that increased system complexity does not degrade performance linearly. Instead, performance degradation follows a sub-linear trend, indicating strong scalability characteristics suitable for large-scale industrial deployment.

System Adaptability in Dynamic Environments

The system demonstrates strong adaptability in environments with fluctuating production demands, machine failures, and unpredictable task arrivals. Reinforcement learning-based adaptation enables real-time policy updates, while anomaly detection provides corrective feedback loops.

This dual adaptation mechanism ensures both proactive optimization and reactive correction, making the system suitable for real-world industrial operations where uncertainty is inherent.

DISCUSSION

The findings of this study highlight the importance of integrating multiple AI paradigms into a unified industrial control framework. The proposed Industrial AI Control Model demonstrates that combining reinforcement learning, anomaly detection, and agentic governance significantly enhances system intelligence and operational reliability.

Theoretical Implications

From a theoretical perspective, the model bridges the gap between classical scheduling theory and modern autonomous systems. Traditional optimization frameworks (Yano & Kim, 1991; Lenstra et al., 1977) focus on deterministic or static environments, whereas the proposed model introduces adaptive learning mechanisms capable of operating in stochastic and dynamic conditions.

The incorporation of agentic governance principles (Venkateela, 2026) provides a structured framework for balancing autonomy and control. This is particularly significant because fully autonomous systems often suffer from unpredictability, while fully centralized systems lack scalability. The proposed hybrid architecture resolves this tension by introducing hierarchical supervision.

Practical Implications

In practical industrial environments, the model offers several advantages:

- Improved production scheduling efficiency through adaptive reinforcement learning
- Enhanced fault detection using multimodal anomaly detection systems

- Reduced operational downtime through predictive simulation and generative modeling
- Scalable deployment across distributed industrial networks

These improvements suggest that the model can be applied in smart manufacturing, logistics optimization, and autonomous production systems.

Limitations of the Proposed Model

Despite its strengths, the model has several limitations. First, the complexity of integrating multiple AI subsystems increases computational overhead, particularly in resource-constrained environments. Second, the reliance on high-quality sensor and operational data may limit performance in poorly instrumented industrial settings.

Additionally, reinforcement learning models may require extensive training periods before achieving stable convergence. This can delay deployment in rapidly changing industrial environments.

Another limitation is interpretability. Although the system incorporates supervisory control mechanisms, deep learning components such as reinforcement learning and anomaly detection networks remain partially opaque, which may hinder explainability in safety-critical applications.

Comparison with Existing Literature

Compared to existing approaches in anomaly detection (Liu et al., 2023; Gu et al., 2024), the proposed model provides a more integrated decision-making framework rather than isolated detection capability. Similarly, while reinforcement learning-based scheduling systems (Liu et al., 2020; Zhang et al., 2020) focus on

optimization, they lack supervisory governance structures.

The key differentiator of this work is its holistic integration of perception, cognition, and control within a single architecture. This enables not only decision optimization but also system-level governance and stability assurance.

Future Research Directions

Future work should focus on improving model efficiency through lightweight AI architectures and edge-optimized learning algorithms. Additionally, research into explainable AI mechanisms could improve transparency in decision-making processes.

Another promising direction is the integration of federated learning, which would allow distributed agents to learn collaboratively without sharing raw data. This would further enhance scalability and privacy in industrial deployments.

Finally, extending the model to incorporate real-time human-AI collaboration could improve decision validation in safety-critical industrial environments.

CONCLUSION

This research presented a comprehensive Industrial AI Control Model (IACM) designed to supervise agent-based systems and enable scalable independent decision-making in complex industrial environments. The study addressed a critical gap in existing industrial AI research: the lack of unified frameworks that integrate autonomous decision-making with structured supervisory control.

The proposed model advances the field by combining multiple AI paradigms—reinforcement learning, anomaly detection, generative modeling, and agentic governance—into a cohesive, multi-layered architecture. This integration allows industrial systems to move beyond isolated optimization or monitoring tasks toward fully coordinated, intelligent ecosystems capable of adaptive and resilient operation.

One of the central contributions of this research is the conceptualization of controlled autonomy. Unlike traditional centralized control systems or fully decentralized autonomous systems, the proposed model introduces a hybrid approach where agents operate independently within a structured supervisory framework. This ensures that local decision-making remains aligned with global system objectives, reducing conflicts and improving overall efficiency. The incorporation of hierarchical control layers further strengthens this alignment by enabling decision-making at strategic, tactical, and operational levels.

The findings demonstrate that reinforcement learning plays a critical role in enabling adaptive decision processes, particularly in dynamic scheduling and resource allocation scenarios. By leveraging actor-critic architectures, the system can continuously refine its decision policies based on environmental feedback. This aligns with prior research in industrial scheduling optimization while extending it into a broader control framework.

Similarly, the integration of advanced anomaly detection mechanisms enhances system reliability and safety. By combining memory-based learning, clustering techniques, and multimodal reasoning, the model achieves robust detection of both known

and unknown anomalies. This is particularly important in industrial settings where unexpected failures can have significant operational and financial consequences.

Another key contribution is the inclusion of generative modeling for predictive simulation. Synthetic data generation improves system robustness by exposing the model to rare and extreme scenarios that are difficult to capture in real-world datasets. This predictive capability enables proactive decision-making, reducing the likelihood of system failures.

A defining feature of the proposed framework is its alignment with agentic AI governance principles, particularly those outlined by Venkiteela (2026). The model demonstrates how scalable autonomy can be achieved without sacrificing control, transparency, or system stability. By incorporating policy-based constraints, risk-aware decision mechanisms, and feedback-driven adaptation, the system maintains a balance between flexibility and reliability. The repeated relevance of this framework throughout the study highlights its importance in shaping next-generation industrial AI systems (Venkiteela, 2026).

Despite its contributions, the research acknowledges several limitations. The complexity of integrating multiple AI subsystems introduces computational challenges, particularly in large-scale deployments. Additionally, the reliance on high-quality data and extensive training processes may limit immediate applicability in certain industrial contexts. Addressing these limitations will require advancements in efficient model design, edge computing, and data management strategies.



The implications of this research extend beyond industrial automation. The principles of controlled autonomy, hierarchical supervision, and multi-agent coordination can be applied to various domains, including smart cities, autonomous transportation systems, and distributed energy management networks.

Future research should focus on validating the proposed model through real-world implementations and empirical testing. This includes deploying the framework in live industrial environments, evaluating its performance under operational constraints, and refining its components based on practical feedback. Additionally, integrating explainable AI techniques and human-in-the-loop decision systems could further enhance trust and usability.

In conclusion, this study establishes a foundational framework for industrial AI systems that are not only intelligent but also structured, scalable, and reliable. By bridging the gap between autonomous decision-making and supervisory control, the proposed model represents a significant step toward the realization of fully adaptive industrial ecosystems.

REFERENCES

1. Bae, Jaehyeok, Jae-Han Lee, and Seyun Kim. "PNI: **Industria** la nomaly detection using position and neighborhood information." Proceedings of the IEEE/CVF International Conference on Computer Vision. 2023.
2. Bao, Tianpeng, et al. "MIAD: A maintenance inspection dataset for unsupervised anomaly detection." Proceedings of the IEEE/CVF International Conference on Computer Vision. 2023.
3. Caetano, F. (2024). A collection of generative algorithms and techniques implemented in Python [GitHub repository]. Available: <https://github.com/caetas/GenerativeZoo>. Access Date: 13. 04. 2025.
4. C. Shyalika, T. P. Silva, and A. S. Karunananda, "Reinforcement learning in dynamic task scheduling: A review," SN Comput. Sci., vol. 1, p. 306, 2020.
5. C.-L. Liu, C.-C. Chang, and C.-J. Tseng, "Actor-critic deep reinforcement learning for solving job shop scheduling problems," IEEE Access, vol. 8, pp. 71 752–71 762, 2020.
6. C. Zhang, W. Song, Z. Cao, J. Zhang, P. S. Tan, and C. Xu, "Learning to dispatch for job shop scheduling via deep reinforcement learning," in Proceedings of the 34th International Conference on Neural Information Processing Systems, ser. NIPS'20, Vancouver, BC, Canada : Curran Associates Inc., 2020, ISBN: 9781713829546.
7. C. A. Yano and Y.-D. Kim, "Algorithms for a class of single-machine weighted tardiness and earliness problems," European Journal of Operational Research, vol. 52, no. 2, pp. 167–178, May 1991.
8. Gu, Zhaopeng, et al. "Anomalygpt: Detecting industrial anomalies using large vision-language models." Proceedings of the AAAI Conference on Artificial Intelligence. Vol. 38. No. 3. 2024.
9. J. Lenstra, A. Rinnooy Kan, and P. Brucker, "Complexity of machine scheduling problems," English, Annals of Discrete Mathematics, vol. 1, pp. 343–362, 1977, ISSN: 0167-5060.
10. J. Xie, L. Gao, K. Peng, X. Li, and H. Li, "Review on flexible job shop scheduling," IET Collaborative

- Intelligent Manufacturing, vol. 1, no. 3, pp. 67–77, 2019.
11. Liu, Yongheng, et al. “Unsupervised image anomaly detection and localization in industry based on self-updated memory and center clustering.” *IEEE Transactions on Instrumentation and Measurement* 72 (2023): 1–10.
 12. Liu, Wenrui, et al. “Diversity-measurable anomaly detection.” *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2023.
 13. Rakhmonov, Akhrorjon Akhma Djon Ugli, et al. “Extensive knowledge distillation model: An end-to-end effective anomaly detection model for real-time industrial applications.” *IEEE Access* (2023).
 14. S. Luo, “Dynamic scheduling for flexible job shop with new job insertions by deep reinforcement learning,” *Appl. Soft Comput.*, vol. 91, p. 106 208, 2020.
 15. TASTI, (2025), *Application-Tailored Synthetic Image Generation*. Available: <https://tasti-project.eu/>, Access Date: 13. 04. 2025.
 16. Venkateela, P. (2026). *An Enterprise Agentic Architecture Framework for Agentic AI Governance and Scalable Autonomy*. *Scientific Journal of Computer Science*, 2(1), 1–17. <https://doi.org/10.64539/sjcs.v2i1.2026.368>

