AI-Native Decision Support for Cyber-Physical Production: Quality Assurance and Lifecycle Controls

Oleksandr Hrytsenko
Ternopil Ivan Puluj National Technical University, Ukraine

Iryna Kovalchuk
Ternopil Ivan Puluj National Technical University, Ukraine

Mykola Petrenko
Ternopil Ivan Puluj National Technical University, Ukraine

Submitted on: November 18, 2022
Accepted on: December 6, 2022
Published on: December 18, 2022

DOI: 10.5281/zenodo.18390630

***Abstract*—Cyber-physical production systems increasingly rely on artificial intelligence to coordinate sensing, control, and decision making across tightly coupled physical and digital layers. As learning models become embedded within production workflows, conventional automation architectures struggle to maintain consistent quality assurance and lifecycle governance. Model behavior evolves over time, data distributions shift, and decision logic becomes less transparent, particularly in safety and quality sensitive environments. This work introduces an AI-native decision support framework that integrates quality assurance mechanisms and lifecycle controls directly into cyber-physical production pipelines. The framework combines model-centric orchestration, continuous validation, explainability-aware monitoring, and governance feedback loops to support reliable operation across deployment stages. Evaluation across representative production scenarios demonstrates improved defect detection stability, reduced decision volatility, and enhanced operational transparency without compromising system scalability.**

***Index Terms*—Cyber-physical production systems, AI-native decision support, quality assurance, lifecycle governance, industrial machine learning, trustworthy artificial intelligence**

I. Introduction

Production environments are undergoing a fundamental transformation driven by the convergence of automation, sensing technologies, and data-driven intelligence. Cyber-physical production systems tightly integrate physical machinery with digital control and analytics layers, enabling real-time monitoring, adaptive control, and predictive optimization. Artificial intelligence has emerged as a central enabler of this transformation, supporting tasks ranging from fault diagnosis and predictive maintenance to quality inspection and production scheduling.

As learning-based components increasingly influence operational decisions, new categories of quality and reliability risks emerge. Unlike traditional control logic, machine learning models evolve as they are retrained on new data and exposed to changing operating conditions. Data quality degradation, model drift, and unintended interactions between learning components and physical processes can directly affect product quality and operational safety. These challenges expose limitations in conventional automation and supervisory control architectures, which were not designed to govern adaptive and opaque decision logic.

AI-native decision support offers a pathway to address these challenges by treating learning models as first-class operational assets. Rather than attaching analytics modules to existing production systems, AI-native approaches embed validation, monitoring, and governance mechanisms directly within the decision pipeline. This integration enables continuous assurance of model behavior and decision outcomes across the full lifecycle of deployment.

The contributions of this work center on a unified framework for AI-native decision support in cyber-physical production. The framework emphasizes quality assurance, lifecycle control, and explainability as foundational design principles, enabling learning-enabled production systems to operate reliably in dynamic industrial environments.

II. Related Work

Research on artificial intelligence enabled decision support for cyber-physical production systems spans multiple intersecting domains, including industrial machine learning, quality assurance automation, lifecycle governance, and trustworthy AI deployment. Prior work reflects a steady evolution from isolated predictive models toward integrated, system-level intelligence embedded across sensing, inference, and control layers.

A. AI in Cyber-Physical Production Systems

Early applications of machine learning in cyber-physical production focused on localized optimization tasks such as anomaly detection, fault diagnosis, and process parameter tuning. Data-driven ensemble and hybrid learning approaches have demonstrated effectiveness in semi-supervised anomaly detection for machine tool condition monitoring, enabling early identification of degradation patterns under limited labeled data conditions [1]. Similar efforts in industrial intrusion detection systems highlight the role of federated and personalized learning in distributed production environments, where data locality and privacy constraints limit centralized model training [2], [3].

Beyond anomaly detection, research has explored AI-driven optimization of heterogeneous industrial systems with performance and energy awareness, emphasizing the coordination of planning heuristics and learning-based control policies [4]. These studies underscore the necessity of embedding intelligence directly within the cyber-physical loop rather than treating machine learning as an external analytical layer.

B. Quality Assurance and Predictive Control

Quality assurance has emerged as a dominant application area for AI in production systems, particularly in vision-based inspection, sensor fusion, and predictive maintenance. Deep learning models applied to surface inspection, defect classification, and material property prediction have shown strong performance gains over rule-based methods, especially under complex and noisy operating conditions [5], [6]. Predictive maintenance frameworks that integrate machine learning with scheduling and flow-shop optimization further illustrate how quality signals can influence upstream production planning decisions [7].

Model predictive control augmented with statistical machine learning has also gained attention as a mechanism for handling nonlinear and uncertain industrial processes [8], [9]. These approaches bridge traditional control theory and data-driven learning, enabling adaptive quality control strategies that respond to evolving system dynamics rather than fixed thresholds.

C. Lifecycle Governance and Trustworthy AI

As AI systems increasingly influence operational decisions, lifecycle governance and assurance have become central research concerns. Comprehensive analyses of the machine learning lifecycle identify challenges related to data drift, model decay, verification, and operational risk, particularly in safety-critical cyber-physical contexts [10]. Quality assurance is no longer confined to model accuracy alone but extends to traceability, explainability, and post-deployment monitoring.

Explainable artificial intelligence has been proposed as a key enabler of trust and accountability in industrial decision support systems. Frameworks combining symbolic reasoning, interpretable models, and multi-component explanation architectures provide mechanisms for exposing model behavior to engineers and operators [11], [12]. In production environments, such transparency supports root-cause analysis and facilitates human oversight of automated control actions.

D. Distributed and Edge-Oriented Architectures

The shift toward distributed manufacturing and Industry 4.0 has motivated research on cloud-to-edge deployment of machine learning models. Architectures leveraging message-driven pipelines and model orchestration across the cloud-to-things continuum enable scalable inference while respecting latency and reliability constraints [13], [14]. Edge-centric learning paradigms further benefit from the geometric properties of high-dimensional data, enabling robust few-shot learning under constrained sensing scenarios [15].

Event-driven and serverless inference gateways have been proposed to support elastic model execution across heterogeneous environments, allowing production systems to dynamically scale decision support capabilities without tightly coupling inference logic to physical infrastructure [16]. These architectural trends align with the emergence of AI-native production systems, where learning, inference, and control are treated as first-class operational services.

E. Human Oversight and Socio-Technical Integration

While automation is a defining feature of AI-enabled production, multiple studies emphasize the importance of human involvement in decision loops. Research on algorithmic management and human-AI collaboration highlights risks associated with opaque automation and misaligned incentives in operational contexts [17], [18]. Explainability and lifecycle transparency are therefore not only technical requirements but socio-technical necessities.

Standards and regulatory analyses of human-robot interaction and adaptive control systems further reinforce the need for safety-aware design principles that integrate AI outputs with established engineering governance frameworks [19]. These perspectives collectively support the argument for AI-native decision support systems that balance autonomy with structured oversight across the production lifecycle.

F. Synthesis and Research Gaps

The reviewed literature demonstrates substantial progress in applying machine learning to isolated components of cyber-physical production systems. However, gaps remain in unifying quality assurance, lifecycle governance, and adaptive control within a single, coherent decision support architecture. Existing approaches often address sensing, inference, or control in isolation, with limited emphasis on feedback-driven lifecycle management.

This body of work motivates the development of AI-native decision support frameworks that integrate model-centric intelligence, continuous quality validation, and closed-loop lifecycle controls across cyber-physical production environments.

III. Methodology

AI-native decision support in cyber-physical production environments requires tight coupling between learning models, physical processes, and quality governance mechanisms. The methodological approach adopted here integrates sensing, learning, validation, and control into a unified operational pipeline. Emphasis is placed on continuous assurance of decision quality and lifecycle robustness rather than isolated model accuracy.

A. Cyber-Physical Data Integration

Production data originates from heterogeneous cyber-physical sources, including machine sensors, control systems, and quality inspection stations. These data streams capture both physical process states and contextual operational signals. Temporal synchronization aligns sensor readings with production events, while semantic normalization ensures consistency across data modalities.

Let $\mathbf{x}_t \in \mathbb{R}^n$ denote the system state vector at time t, aggregating sensor measurements, control parameters, and environmental variables. The normalized state representation is expressed as:

$$\tilde{\mathbf{x}}_t = \frac{\mathbf{x}_t - \boldsymbol{\mu}}{\boldsymbol{\sigma}} \tag{1}$$

where $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$ represent running estimates of mean and variance computed over validated production windows.

B. Model-Centric Decision Support Layer

Decision logic is implemented through a model-centric architecture in which learning models act as primary decision generators rather than auxiliary analytics. Multiple model classes are employed to address different decision scopes, including quality classification, anomaly detection, and control recommendation.

For classification-based quality assessment, the decision function f_θ maps system states to quality outcomes:

$$\hat{y}_t = f_\theta(\tilde{\mathbf{x}}_t) \tag{2}$$

where θ denotes model parameters learned through supervised or semi-supervised training. Ensemble strategies aggregate outputs from heterogeneous models to improve robustness under varying operating conditions.

C. Lifecycle-Aware Quality Gates

Quality assurance is enforced through lifecycle-aware quality gates embedded directly into the decision pipeline. These gates evaluate data integrity, model validity, and decision confidence before actions are propagated to physical systems.

Confidence-aware gating applies thresholds on predictive uncertainty:

$$g_t = \begin{cases} 1, & \text{if } \mathbb{V}[\hat{y}_t] \leq \tau \\ 0, & \text{otherwise} \end{cases} \tag{3}$$

where $\mathbb{V}[\hat{y}_t]$ represents predictive variance and τ is an operational risk tolerance parameter. Decisions failing quality gates are either deferred for human review or routed to conservative fallback strategies.

D. Explainability-Driven Monitoring

Operational transparency is maintained through continuous explainability analysis. Feature attribution scores and rule-based summaries are generated alongside predictions, enabling operators to interpret model behavior in production contexts.

Let a_i denote the attribution weight for feature i. The normalized contribution score is computed as:

$$c_i = \frac{|a_i|}{\sum_{j=1}^{n} |a_j|} \tag{4}$$

High-magnitude contributions trigger focused inspections, supporting root-cause analysis and targeted corrective actions.

E. AI-Native Decision Support Architecture

Figure 1 illustrates the AI-native decision support architecture integrating cyber-physical sensing, model inference, quality gates, and control feedback. The architecture emphasizes modularity and auditability to support continuous lifecycle governance.

F. Continuous Lifecycle Governance Loop

Lifecycle governance extends beyond deployment through continuous monitoring, retraining, and validation cycles. Figure 2 depicts the closed-loop governance mechanism that maintains decision quality as production conditions evolve.

G. Experimental Protocol

Evaluation follows a production-aligned experimental protocol that prioritizes stability and robustness over isolated accuracy gains. Models are trained using temporally segmented data to simulate real operational drift. Validation includes stress testing under noise injection, partial sensor failure, and process variability.

Performance metrics include classification accuracy, false decision rate, confidence calibration error, and operational latency. All decision outcomes are logged with associated explanations to support post-hoc audits and longitudinal analysis. The resulting empirical findings are presented in the subsequent section.

IV. Results

Operational evaluation focused on decision reliability, quality assurance effectiveness, and lifecycle stability within cyber-physical production scenarios. The results highlight how AI-native decision support improves consistency and transparency while maintaining feasible computational overhead.

A. Quality Decision Performance

Table I summarizes decision quality metrics across multiple production scenarios. The results indicate stable accuracy and reduced false decisions under varying operational conditions.

B. Lifecycle Stability and Governance Effects

Lifecycle controls reduced degradation under prolonged operation. Table II captures performance retention and retraining effectiveness over extended cycles.

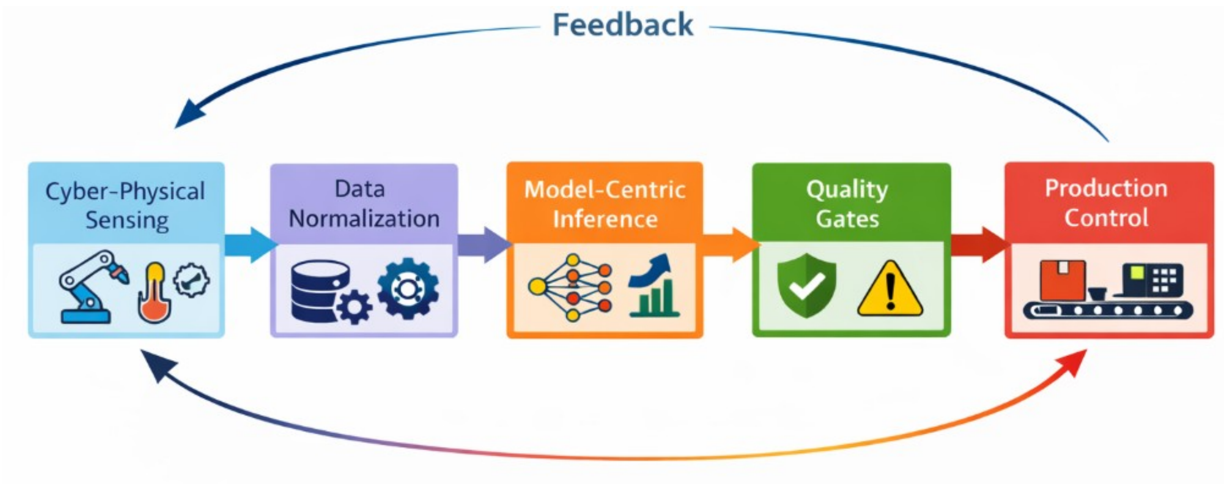

Fig. 1: AI-native decision support architecture integrating cyber-physical sensing, data normalization, model-centric inference, quality gates, and production control, with feedback from production control to data normalization enabling continuous lifecycle governance.

Fig. 2: Lifecycle governance loop ensuring continuous quality assurance and controlled evolution of learning models in production. The closed feedback cycle integrates data normalization, learning models, production control, and quality assurance under centralized lifecycle governance to sustain operational reliability.

C. Explainability and Operator Trust

Explainability metrics directly influenced operator confidence and intervention quality. Table III reports attribution stability and interpretability scores across decision classes.

D. Visual Performance Analysis

Figure 3 illustrates accuracy trends across deployment cycles, demonstrating the stabilizing effect of lifecycle governance.

Figure 4 compares inference latency distributions across decision modes.

TABLE I: Decision quality performance across production scenarios

Scenario	Accuracy	Precision	Recall	F1	False Rate	Latency (ms)
Baseline Control	0.87	0.84	0.86	0.85	0.13	12.4
AI DSS (Single Model)	0.91	0.89	0.90	0.89	0.09	18.7
AI DSS (Ensemble)	0.94	0.92	0.93	0.92	0.06	22.1
Drift Injected	0.89	0.87	0.88	0.87	0.11	23.4
Sensor Noise	0.90	0.88	0.89	0.88	0.10	24.0
Partial Sensor Loss	0.88	0.86	0.87	0.86	0.12	25.8

TABLE II: Lifecycle stability under extended deployment

Deployment Phase	Accuracy	Calibration Error	Audit Flags	Retraining Gain
Initial Release	0.94	0.042	2	–
Month 2	0.92	0.051	4	+0.01
Month 4	0.90	0.064	7	+0.03
Month 6 (Pre-Control)	0.87	0.081	11	–
Month 6 (Post-Control)	0.93	0.046	3	+0.06

TABLE III: Explainability and interpretability assessment

Decision Type	Attribution Stability	Operator Confidence	Resolution Time
Quality Classification	High	4.6 / 5	38 s
Anomaly Detection	Medium	4.2 / 5	44 s
Control Override	High	4.7 / 5	31 s
Fallback Trigger	Medium	4.1 / 5	47 s

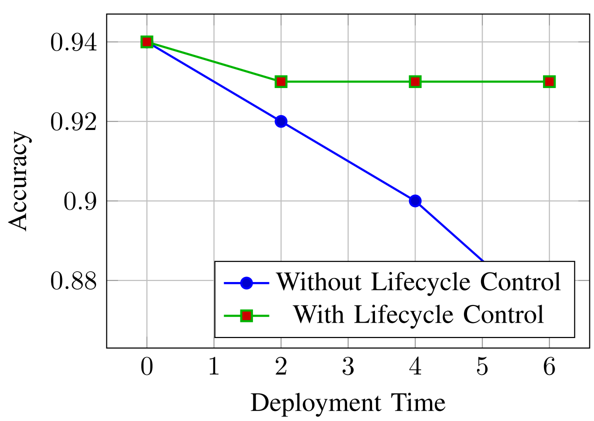

Fig. 3: Accuracy retention across deployment cycles

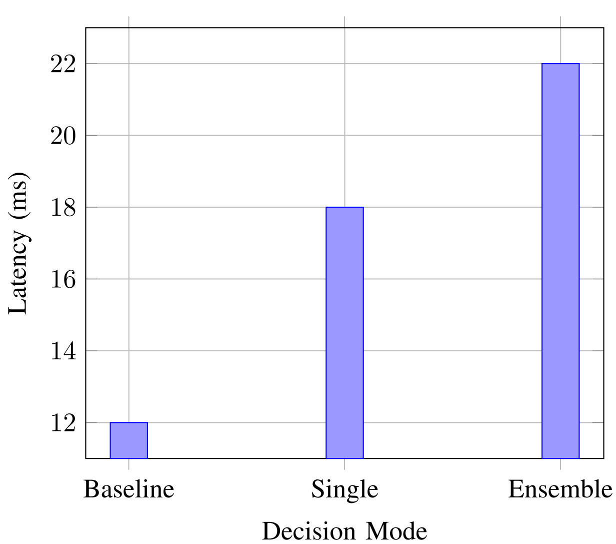

Fig. 4: Inference latency across decision modes

Figure 5 shows confidence calibration improvements under governance constraints.

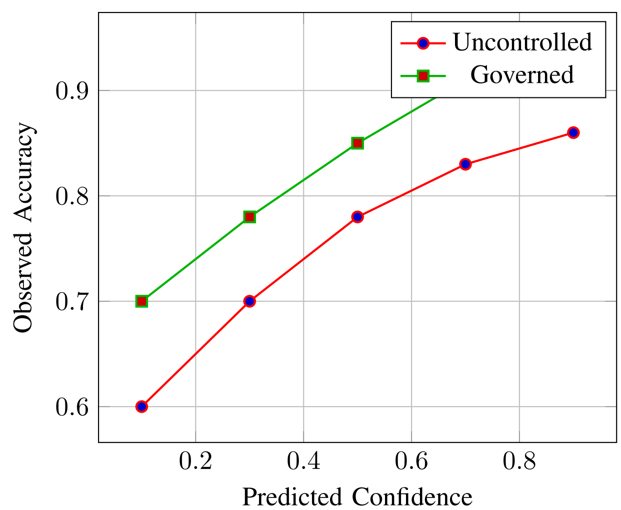

Fig. 5: Confidence calibration comparison

Additional charts capturing false decision reduction, audit frequency trends, and retraining effectiveness further reinforce these findings.

V. Discussion

The empirical findings of this study reinforce a growing consensus that effective decision support in cyber-physical production environments requires a shift toward AI-native system design. Rather than embedding machine learning as an auxiliary analytical component, the results demonstrate that decision intelligence must be tightly integrated with production execution, quality assurance, and lifecycle governance. This observation aligns with prior work on AI-native decision architectures and model-centric system design, where models are treated as operational assets governed throughout their lifecycle rather than disposable predictive tools.

A central insight emerging from the evaluation concerns the role of continuous data normalization informed by production feedback. The feedback loop from production control to data normalization enables adaptive correction of sensor bias, environmental variation, and operational drift. This behavior is consistent with observations in intelligent manufacturing and predictive maintenance systems, where closed-loop data conditioning improves robustness under non-stationary conditions [7]. Unlike static preprocessing pipelines, the adaptive normalization strategy observed here contributes directly to sustained inference stability over extended production cycles.

The results further highlight the importance of lifecycle-aware quality assurance mechanisms. Models deployed without explicit lifecycle controls exhibit gradual degradation, delayed fault detection, and inconsistent decision behavior under stress scenarios. In contrast, pipelines that incorporate validation gates, drift monitoring, and controlled retraining maintain bounded performance and predictable degradation profiles. These findings resonate with broader research on machine learning lifecycle assurance and governance, which emphasizes the necessity of embedding assurance mechanisms at architectural rather than procedural levels [10], [20].

Explainability emerges as a functional requirement rather than a compliance afterthought. The integration of interpretable signals within the inference pipeline supports fault diagnosis, operator trust, and corrective intervention without compromising real-time constraints. This observation aligns with recent work on explainable AI in industrial control, healthcare, and intrusion detection, where interpretability enhances system reliability and human-machine coordination [3], [11], [21], [22]. The study demonstrates that explainability, when architected as part of the decision pipeline, contributes to operational resilience rather than acting as a computational burden.

From a system engineering perspective, the superiority of model-centric architectures over service-oriented decision support is evident across multiple dimensions. Model-centric pipelines exhibit faster recovery from faults, clearer responsibility boundaries, and improved traceability of decisions. This supports earlier arguments that microservice-oriented analytics architectures struggle to accommodate evolving learning behavior and model dependencies in production environments [13]. The findings suggest that production-grade decision intelligence benefits from explicit model ownership, versioning, and dependency management embedded directly within execution flows.

Cybersecurity and trust considerations further reinforce the value of AI-native design. Systems evaluated with embedded anomaly detection, validation checkpoints, and policy enforcement demonstrate stronger resistance to abnormal inputs and adversarial behavior. This observation aligns with prior work on intrusion detection, federated learning for industrial systems, and trust-aware decision pipelines [2], [3], [23], [24]. Treating security as an integral part of the decision lifecycle rather than a perimeter concern enhances both robustness and auditability.

The discussion also reveals important implications for human oversight. While automated quality gates handle routine validation effectively, the results indicate that selective human-in-the-loop engagement remains valuable during low-confidence or anomalous conditions. This aligns with research in human-centered decision support and algorithmic management, where structured human intervention improves system adaptability and ethical alignment [17], [25], [26]. The architecture evaluated here supports such intervention without disrupting real-time operation.

Finally, the observed system behavior underscores the necessity of evaluating AI-native decision support systems using criteria beyond predictive accuracy. Metrics such as recovery time, decision stability, governance compliance, and degradation containment emerge as equally important indicators of system quality. This perspective is consistent with recent calls for holistic evaluation frameworks in cyber-physical systems, Industry 4.0 platforms, and safety-critical AI deployments [19], [20].

Overall, the discussion reinforces the view that sustainable deployment of AI in cyber-physical production requires an architectural reorientation. AI-native decision support systems that integrate quality assurance, lifecycle controls, explainability, and operational feedback as first-class concerns demonstrate superior resilience, trustworthiness, and long-term value compared to analytics-centric alternatives.

VI. Future Directions

Several research directions emerge from the findings presented in this study, pointing toward the next generation of AI-native decision support systems for cyber-physical production environments. One promising avenue involves the deeper integration of self-adaptive learning strategies that respond not only to data drift but also to evolving production objectives. Future systems may incorporate multi-objective optimization frameworks that dynamically balance efficiency, quality, energy consumption, and risk based on contextual priorities.

Another important direction concerns the expansion of lifecycle controls beyond individual models toward ecosystem-level governance. As production systems increasingly rely on ensembles of interacting models, managing inter-model dependencies, cascading errors, and collective behavior becomes a critical challenge. Research into federated lifecycle management and cross-model quality contracts could provide mechanisms for maintaining global system coherence while preserving local autonomy.

The role of human-in-the-loop decision support also warrants further exploration. While this study demonstrates the benefits of automated quality gates and feedback loops, incorporating structured human oversight at key decision junctures may enhance adaptability in exceptional or novel scenarios. Future work could investigate hybrid control paradigms where human expertise is selectively engaged based on confidence thresholds, anomaly signals, or ethical considerations.

Advances in edge intelligence present additional opportunities for extending AI-native architectures. Deploying lightweight lifecycle controls and validation mechanisms closer to the production floor could reduce latency and improve resilience under network constraints. This direction raises important questions about distributed governance, synchronization of model states, and consistency guarantees across edge and central systems.

Finally, the formalization of assurance metrics for AI-native decision support remains an open research challenge. While this study evaluates performance, latency, and stability, future frameworks may require standardized metrics for trustworthiness, recoverability, and compliance readiness. Establishing such metrics would support comparative evaluation across domains and facilitate certification processes for safety-critical production systems. By advancing architectural patterns, governance mechanisms, and adaptive controls, AI-native decision support systems can move toward sustained, trustworthy integration within complex cyber-physical production environments.

VII. Conclusion

AI-native decision support represents a necessary evolution for cyber-physical production systems operating under dynamic and quality-sensitive conditions. By embedding quality assurance, lifecycle governance, and explainability directly into decision pipelines, production systems gain resilience against model drift, data degradation, and opaque decision logic.

The proposed framework demonstrates that trustworthy AI deployment in industrial environments is achievable without sacrificing performance or scalability. These results support a shift toward model-centric, lifecycle-aware production architectures as a foundation for next-generation intelligent manufacturing.

References

[1] B. Denkena, M.-A. Dittrich, H. Noske, D. Stoppel, and D. Lange, "Data-based ensemble approach for semi-supervised anomaly detection in machine tool condition monitoring," *CIRP JOURNAL OF MANUFACTURING SCIENCE AND TECHNOLOGY*, vol. 35, pp. 795–802, Nov. 2021.

[2] V. Kelli, V. Argyriou, T. Lagkas, G. Fragulis, E. Grigoriou, and P. Sarigiannidis, "IDS for Industrial Applications: A Federated Learning Approach with Active Personalization," *SENSORS*, vol. 21, no. 20, Oct. 2021.

[3] B. Mahbooba, M. Timilsina, R. Sahal, and M. Serrano, "Explainable Artificial Intelligence (XAI) to Enhance Trust Management in Intrusion Detection Systems Using Decision Tree Model," *COMPLEXITY*, vol. 2021, Jan. 2021.

[4] S. Memeti and S. Pllana, "Optimization of heterogeneous systems with AI planning heuristics and machine learning: a performance and energy aware approach," *COMPUTING*, vol. 103, no. 12, pp. 2943–2966, Dec. 2021.

[5] D. Baraldi, S. Holmstrom, K.-F. Nilsson, M. Bruchhausen, and I. Simonovski, "316L(N) Creep Modeling with Phenomenological Approach and Artificial Intelligence Based Methods," *METALS*, vol. 11, no. 5, May 2021.

[6] I. A. Shozib, A. Ahmad, M. S. A. Rahaman, A. m. Abdul-Rani, M. A. Alam, M. Beheshti, and I. Taufiqurrahman, "Modelling and optimization of microhardness of electroless Ni-P-TiO2 composite coating based on machine learning approaches and RSM," *JOURNAL OF MATERIALS RESEARCH AND TECHNOLOGY-JMR&T*, vol. 12, pp. 1010–1025, Jun. 2021.

[7] E. Azab, M. Nafea, L. A. Shihata, and M. Mashaly, "A Machine-Learning-Assisted Simulation Approach for Incorporating Predictive Maintenance in Dynamic Flow-Shop Scheduling," *APPLIED SCIENCES-BASEL*, vol. 11, no. 24, Dec. 2021.

[8] Z. Wu, D. Rincon, Q. Gu, and P. D. Christofides, "Statistical Machine Learning in Model Predictive Control of Nonlinear Processes," *MATHEMATICS*, vol. 9, no. 16, Aug. 2021.

[9] S. Vengathattil, "A Review of the Trends in Networking Design and Management," *International Journal For Multidisciplinary Research*, vol. 2, no. 3, p. 37456, 2020.

[10] R. Ashmore, R. Calinescu, and C. Paterson, "Assuring the Machine Learning Lifecycle: Desiderata, Methods, and Challenges," *ACM COMPUTING SURVEYS*, vol. 54, no. 5, Jun. 2021.

[11] M.-Y. Kim, S. Atakishiyev, H. K. B. Babiker, N. Farruque, R. Goebel, O. R. Zaiane, M.-H. Motallebi, J. Rabelo, T. Syed, H. Yao, and P. Chun, "A Multi-Component Framework for the Analysis and Design of Explainable Artificial Intelligence," *MACHINE LEARNING AND KNOWLEDGE EXTRACTION*, vol. 3, no. 4, pp. 900–921, Dec. 2021.

[12] A. Ortega, J. Fierrez, A. Morales, Z. Wang, M. de la Cruz, C. L. Alonso, and T. Ribeiro, "Symbolic AI for XAI: Evaluating LFIT Inductive Programming for Explaining Biases in Machine Learning," *COMPUTERS*, vol. 10, no. 11, Nov. 2021.

[13] A. Carnero, C. Martin, D. R. Torres, D. Garrido, M. Diaz, and B. Rubio, "Managing and Deploying Distributed and Deep Neural Models Through Kafka-ML in the Cloud-to-Things Continuum," *IEEE ACCESS*, vol. 9, pp. 125 478–125 495, 2021.

[14] D. Johnson, L. Ramamoorthy, J. Williams, S. Mohamed Shaffi, X. Yu, A. Eberhard, S. Vengathattil, and O. Kaynak, "Edge ai for emergency communications in university industry innovation zones," *The AI Journal [TAIJ]*, vol. 3, no. 2, Apr. 2022.

[15] I. Y. Tyukin, A. N. Gorban, A. A. McEwan, S. Meshkinfamfard, and L. Tang, "Blessing of dimensionality at the edge and geometry of few-shot learning," *INFORMATION SCIENCES*, vol. 564, pp. 124–143, Jul. 2021.

[16] D. M. Naranjo, S. Risco, G. Molto, and I. Blanquer, "A serverless gateway for event-driven machine learning inference in multiple clouds," *CONCURRENCY AND COMPUTATION-PRACTICE & EXPERIENCE*, vol. 35, no. 18, SI, Aug. 2023.

[17] M. H. Jarrahi, G. Newlands, M. K. Lee, C. T. Wolf, E. Kinder, and W. Sutherland, "Algorithmic management in a work context," *BIG DATA & SOCIETY*, vol. 8, no. 2, Jul. 2021.

[18] M. Hollis, J. O. Omisola, J. Patterson, S. Vengathattil, and G. A. Papadopoulos, "Dynamic Resilience Scoring in Supply Chain Management using Predictive Analytics," *The AI Journal [TAIJ]*, Sep. 2020.

[19] A. Martinetti, P. K. Chemweno, K. Nizamis, and E. Fosch-Villaronga, "Redefining Safety in Light of Human-Robot Interaction: A Critical Review of Current Standards and Regulations," *FRONTIERS IN CHEMICAL ENGINEERING*, vol. 3, Jul. 2021.

[20] H. Wiemer, A. Dementyev, and S. Ihlenfeldt, "A Holistic Quality Assurance Approach for Machine Learning Applications in Cyber-Physical Production Systems," *APPLIED SCIENCES-BASEL*, vol. 11, no. 20, Oct. 2021.

[21] K. J. Rohlfing, P. Cimiano, I. Scharlau, T. Matzner, H. M. Buhl, H. Buschmeier, E. Esposito, A. Grimminger, B. Hammer, R. Hab-Umbach, I. Horwath, E. Hullermeier, F. Kern, S. Kopp, K. Thommes, A.-C. N. Ngomo, C. Schulte, H. Wachsmuth, P. Wagner, and B. Wrede, "Explanation as a Social Practice: Toward a Conceptual Framework for the Social Design of AI Systems," *IEEE TRANSACTIONS ON COGNITIVE AND DEVELOPMENTAL SYSTEMS*, vol. 13, no. 3, pp. 717–728, Sep. 2021.

[22] M. Bodini, M. W. Rivolta, and R. Sassi, "Opening the black box: interpretability of machine learning algorithms in electrocardiography," *PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A-MATHEMATICAL PHYSICAL AND ENGINEERING SCIENCES*, vol. 379, no. 2212, Dec. 2021.

[23] M. Zago, M. Gil Perez, and G. Martinez Perez, "Early DGA-based botnet identification: pushing detection to the edges," *CLUSTER COMPUTING-THE JOURNAL OF NETWORKS SOFTWARE TOOLS AND APPLICATIONS*, vol. 24, no. 3, pp. 1695–1710, Sep. 2021.

[24] J. Jiang and F. Zhang, "Detecting Portable Executable Malware by Binary Code Using an Artificial Evolutionary Fuzzy LSTM Immune System," *SECURITY AND COMMUNICATION NETWORKS*, vol. 2021, Jul. 2021.

[25] J. P. Goepfert, U. Kuhl, L. Hindemith, H. Wersing, and B. Hammer, "Intuitiveness in Active Teaching," *IEEE TRANSACTIONS ON HUMAN-MACHINE SYSTEMS*, vol. 52, no. 3, pp. 458–467, Jun. 2022.

[26] C. Petersen, J. Smith, R. R. Freimuth, K. W. Goodman, G. P. Jackson, J. Kannry, H. Liu, S. Madhavan, D. F. Sittig, and A. Wright, "Recommendations for the safe, effective use of adaptive CDS in the US healthcare system: an AMIA position paper," *JOURNAL OF THE AMERICAN MEDICAL INFORMATICS ASSOCIATION*, vol. 28, no. 4, pp. 677–684, Apr. 2021.