The American Journal of Engineering and Technology
ISSN 2689-0984 Volume 08 - 2026

An Edge-AI-Based Anomaly Detection Framework for Securing Industrial IoT and Cyber-Physical Manufacturing Systems

Dr. Aditya Pratama
Department of Artificial Intelligence, Jakarta Institute of Technology, Jakarta, Indonesia

Dr. Siti Rahmawati
Department of Intelligent Computing, Bandung Institute of Technology Studies, Bandung, Indonesia

Received: 05 May 2026 | Received Revised Version: 15 June 2026 | Accepted: 25 July 2026 | Published: 27 August 2026

Volume 08 Issue 08 2026 | Crossref DOI: 10.37547/tajet/Volume08Issue08-07

Abstract

The convergence of Industrial Internet of Things (IIoT), artificial intelligence, automation, and cyber-physical systems (CPS) is transforming manufacturing into highly connected and intelligent production environments. However, increased connectivity also expands the attack surface and creates a requirement for security mechanisms capable of identifying anomalous behavior with low latency. Conventional centralized security architectures may introduce communication overhead, delayed responses, and dependency on continuous connectivity with cloud infrastructure. This paper proposes an Edge-AI-Based Anomaly Detection Framework for securing IIoT and cyber-physical manufacturing systems. The proposed framework places lightweight artificial intelligence inference capabilities close to industrial data sources, enabling continuous analysis of sensor, machine, process, and network observations. A layered architecture comprising data acquisition, preprocessing, edge intelligence, anomaly scoring, decision enforcement, and centralized coordination is developed conceptually. The methodology integrates contextual process behavior with machine-learning-based anomaly identification to distinguish operational deviations from potentially malicious activities. The literature synthesis demonstrates that intelligent agriculture, computational intelligence, and IoT-enabled management systems provide useful foundations for distributed sensing, intelligent decision-making, and real-time analytics, although these studies do not directly address the security requirements of industrial CPS. The proposed framework therefore extends these principles toward security-oriented edge intelligence. Findings indicate that edge-based detection can improve responsiveness, reduce unnecessary data transmission, and support localized decision-making, while introducing challenges related to resource constraints, model maintenance, false positives, and heterogeneous industrial environments. The framework provides a research-oriented foundation for deploying explainable, low-latency anomaly detection in next-generation smart manufacturing environments.

Keywords: Edge AI, Industrial IoT, anomaly detection, cyber-physical systems, smart manufacturing, artificial intelligence, cybersecurity, machine learning, intelligent manufacturing.

Cite This Article: Pratama, D. A., & Rahmawati, D. S. (2026). An Edge-AI-Based Anomaly Detection Framework for Securing Industrial IoT and Cyber-Physical Manufacturing Systems. The American Journal of Engineering and Technology, 8(08), 58–65. Retrieved from https://theamericanjournals.com/index.php/tajet/article/view/8343.

1. Introduction

1.1 Background

Industrial manufacturing is increasingly characterized by interconnected sensors, automated machinery, intelligent control systems, and data-driven decision processes. The broader movement toward intelligent and digitally integrated production depends on the ability of machines and computational systems to continuously observe physical processes and respond to changing operational conditions. Computational intelligence has become an important technological foundation for this transformation because it enables complex systems to process large volumes of heterogeneous information and

support automated decision-making (Chen and Li, 2018).

The same connectivity that improves manufacturing intelligence, however, creates additional cybersecurity exposure. In a cyber-physical manufacturing system, digital commands can influence physical equipment, while sensor observations provide information about physical processes to computational components. Consequently, abnormal cyber activity can potentially manifest as unusual process behavior, sensor inconsistencies, communication irregularities, or unexpected machine states. A security architecture must therefore consider both computational and physical dimensions rather than treating cybersecurity as an isolated network-monitoring problem.

IoT-based intelligent management systems illustrate how distributed sensing can connect physical environments with computational decision mechanisms. Nagaraja et al. (2019), for example, demonstrate the integration of IoT technologies into smart agricultural management, highlighting the value of connected sensing and centralized intelligence. Similarly, Siddhartha and Lakkannavar (2021) demonstrate the use of computational prediction for irrigation and crop-health-related decisions. Although these applications differ substantially from industrial cybersecurity, their underlying principle—continuous sensing followed by intelligent interpretation—is relevant to IIoT anomaly detection.

The emergence of edge artificial intelligence provides an alternative to architectures in which all operational data must be transmitted to centralized infrastructure. Rather than transferring every observation to a remote server, computational intelligence can be positioned close to machines and sensors. This enables preliminary analysis at the point where data are generated. For industrial applications, such localization is particularly relevant when rapid decisions are required and communication delays could affect operational safety.

Recent work on hybrid security architectures for Industry 5.0 further emphasizes the importance of combining multiple intelligent techniques to address increasingly complex manufacturing threats. The Aegis-5 framework, for example, positions hybrid ensemble learning as a mechanism for intrusion detection within Industry 5.0-driven smart manufacturing environments (Govindarajan et al., 2026). The present study complements that direction by concentrating specifically on an edge-oriented anomaly-detection architecture in which computational intelligence is deployed closer to industrial assets.

1.2 Problem Statement

A major challenge in IIoT security is the detection of abnormal conditions without imposing excessive computational and communication requirements on industrial infrastructure. Centralized approaches can provide substantial analytical capacity, but transmitting continuous streams of sensor and operational data can increase bandwidth requirements and create dependency on reliable communication. Conversely, fully decentralized processing may be computationally constrained because industrial edge devices often have limited memory, processing power, and energy availability.

Another challenge concerns the distinction between operational anomalies and security anomalies. A machine may exhibit unusual vibration because of mechanical degradation, while another abnormal pattern may result from malicious manipulation of sensor measurements. Treating every deviation as a cybersecurity incident can produce excessive false alarms. An effective framework must therefore incorporate contextual information about expected industrial behavior.

1.3 Objectives and Scope

This paper develops a conceptual Edge-AI-based framework for anomaly detection in IIoT and cyber-physical manufacturing systems. The objectives are to establish a layered architecture for edge-based anomaly detection, identify the functional roles of edge intelligence within industrial security, define an anomaly-scoring and response mechanism, and critically evaluate the expected advantages and limitations of the proposed approach.

The scope is intentionally focused on conceptual framework development rather than experimental benchmarking. The analysis is grounded exclusively in the references supplied for this study. The framework consequently represents a research model that can subsequently be implemented and validated using industrial datasets and real manufacturing environments.

2. Literature Review

2.1 Intelligent Systems and Computational Intelligence

Chen and Li (2018) establish computational intelligence as an important component of design and decision-making in the context of industrial transformation. Their work is relevant to the proposed framework because anomaly detection similarly requires computational mechanisms capable of recognizing complex relationships rather than relying exclusively on fixed rules.

The theoretical implication is significant. Industrial systems generate multidimensional observations in which normal behavior may vary according to machine workload, environmental conditions, production schedules, and operational states. A computational-intelligence perspective therefore supports adaptive representations of system behavior. However, Chen and Li (2018) focus on computational intelligence-assisted design rather than cybersecurity-specific anomaly detection, leaving an opportunity to adapt intelligent computational principles to industrial security.

2.2 IoT-Based Intelligent Management

Nagaraja et al. (2019) demonstrate an IoT-based smart agriculture management system, illustrating the integration of distributed sensing, connectivity, and intelligent management. The relevance to IIoT arises from the shared architecture of physical sensors, communication networks, and computational decision systems. In both domains, large numbers of connected devices continuously generate observations that can be transformed into actionable information.

Nevertheless, IoT connectivity introduces a critical distinction between monitoring and security. A management system may primarily seek to optimize an operational objective, whereas an anomaly-detection system must additionally determine whether an observation represents a deviation from legitimate behavior. The proposed framework therefore extends the connected sensing principle toward security-oriented behavioral analysis.

2.3 Predictive Analytics and Contextual Anomalies

Siddhartha and Lakkannavar (2021) demonstrate predictive analysis for smart irrigation and crop health. Predictive modeling is theoretically important for anomaly detection because the identification of abnormality depends on an estimate of expected behavior. If a system can estimate the expected state of a process, the deviation between predicted and observed states can serve as an anomaly signal.

This principle can be translated into manufacturing. For instance, an edge model may estimate the expected temperature, vibration, pressure, or production rate of a machine under a particular operating condition. A substantial deviation can then generate an anomaly score. However, prediction errors themselves may produce false alarms, emphasizing the need for adaptive thresholds and contextual interpretation.

2.4 Intelligent Observation and Automated Decision-Making

The references related to agricultural intelligence also demonstrate the increasing role of computer vision and artificial intelligence in physical environments. Cainthus (2019) describes the combination of computer vision and artificial intelligence, while J. Deere (2019) documents AutoTrac Agricultural Management Solutions. These examples demonstrate how AI and automated guidance can operate in environments where physical processes must be interpreted continuously.

Their relevance to industrial CPS lies in the underlying design principle: intelligent computation can be placed within operational environments to transform observations into decisions. However, applying such principles to cybersecurity requires an additional security layer capable of identifying malicious deviations without unnecessarily interrupting legitimate operations.

2.5 Sustainability, Scale, and Digital Transformation

Tahat et al. (2020) emphasize the importance of soil health and sustainable agriculture, while the United Nations reference concerning global population growth and sustainable development establishes the broader importance of scalable and sustainable technological systems. Although these sources are not cybersecurity studies, they reinforce the requirement for intelligent technologies that can support increasingly complex operational environments.

PwC (2017) similarly highlights the expected economic significance of artificial intelligence. This provides contextual justification for investigating AI-enabled approaches, although economic projections alone cannot establish the effectiveness of a security architecture.

WIPO (2019) identifies artificial intelligence as a major technological trend. Collectively, the supplied literature indicates that AI, IoT, automation, and computational intelligence are increasingly interconnected. However, a clear research gap remains: the provided application-

oriented literature does not establish a dedicated edge-based anomaly detection architecture for securing industrial CPS. The proposed framework addresses this gap conceptually, while the Industry 5.0 security perspective of Govindarajan et al. (2026) provides a contemporary security-oriented point of comparison.

3. Methodology

3.1 Research Design

The study adopts a framework-development methodology. Instead of claiming experimentally validated performance, it constructs a technically grounded architecture from principles derived from the supplied literature. The framework is designed around six functional layers: data acquisition, edge preprocessing, edge-AI inference, anomaly scoring, response enforcement, and centralized coordination.

The central theoretical assumption is that an industrial system produces a recognizable operational behavior profile. Let an observation at time tt be represented as XtX_t, incorporating measurements from sensors, machine states, communication characteristics, and process variables. The edge model generates an expected representation X^t\hat{X}_t. The difference between observed and expected behavior can be transformed into an anomaly score:

At=D(Xt,X^t)A_t = D(X_t,\hat{X}_t)

where DD represents an appropriate distance, reconstruction error, classification uncertainty, or hybrid deviation function.

An anomaly decision can subsequently be represented as:

Yt={0,At<τ1,At≥τY_t = \begin{cases} 0, & A_t < \tau \\ 1, & A_t \geq \tau \end{cases}

where τ\tau is a context-dependent threshold. In a practical industrial environment, τ\tau should not necessarily remain constant because machine behavior changes with production state.

3.2 Layer 1: Industrial Data Acquisition

The first layer collects data from IIoT sensors, programmable controllers, machines, gateways, and operational systems. Relevant observations may include temperature, pressure, vibration, current, machine state, production rate, communication frequency, and command sequences.

The framework treats these observations as security-relevant signals rather than merely operational measurements. For example, an unexpected increase in motor temperature may indicate equipment degradation, while an unusual combination of temperature, command activity, and communication behavior may indicate a more complex anomaly.

This approach follows the distributed sensing principle demonstrated by IoT-based intelligent management systems (Nagaraja et al., 2019). The principal limitation is heterogeneity: industrial devices can differ considerably in protocols, sampling frequencies, computational capacity, and data quality.

3.3 Layer 2: Edge Preprocessing

Raw industrial data should be normalized before inference. The preprocessing layer performs noise filtering, missing-value handling, feature normalization, temporal alignment, and feature reduction.

Edge preprocessing serves two purposes. First, it reduces the volume of information that must be transmitted. Second, it allows the detection pipeline to operate independently of continuous cloud connectivity. A vibration stream, for example, could be transformed into statistical or frequency-domain indicators at the edge rather than transmitting every raw measurement.

The trade-off is that aggressive feature reduction can remove information that may later prove useful for detecting sophisticated anomalies. Therefore, preprocessing should preserve security-relevant characteristics while minimizing computational overhead.

3.4 Layer 3: Edge-AI Inference

The core component is a lightweight AI model deployed at an industrial gateway or edge device. Depending on the application, the model may use classification, reconstruction, prediction, clustering, or ensemble-based inference.

A particularly useful design is a hybrid anomaly engine that combines process prediction with behavioral classification. The prediction component estimates the expected state, while the classification component evaluates whether the observed pattern resembles known legitimate or abnormal behavior.

The architecture is conceptually compatible with computational-intelligence principles discussed by Chen

and Li (2018). It also aligns with the broader development of AI-enabled physical-environment applications described by Cainthus (2019) and WIPO (2019).

3.5 Layer 4: Context-Aware Anomaly Scoring

A single deviation should not automatically trigger a security response. The framework therefore combines multiple indicators into a contextual anomaly score:

St=w1Aprocess+w2Asensor+w3Anetwork+w4AcommandS_t = w_1A_{process}+w_2A_{sensor}+w_3A_{network}+w_4A_{command}

where the wiw_i values represent the relative importance of each anomaly dimension.

For example, a temperature deviation during machine startup may be normal, while the same deviation during a stable production phase may be suspicious. Context therefore reduces the probability of treating legitimate operational variability as an attack.

This is one of the principal distinctions between basic monitoring and security-oriented anomaly detection. The framework aims to detect deviations in relationships among variables, rather than examining each variable independently.

3.6 Layer 5: Edge Response and Enforcement

Once the anomaly score exceeds a predefined risk threshold, the framework initiates a localized response. Response severity can be proportional to anomaly confidence. Low-risk deviations may simply be logged, medium-risk deviations may trigger alerts and additional observation, while high-confidence anomalies may result in temporary isolation of a compromised endpoint or restriction of suspicious commands.

Localized response is essential because waiting for centralized analysis can increase response latency. However, automated intervention in industrial environments introduces operational risk. An incorrect shutdown can disrupt production or potentially create unsafe conditions. Consequently, response policies should incorporate machine criticality and operational state.

3.7 Layer 6: Central Coordination

The edge layer should not eliminate centralized infrastructure. Instead, the proposed architecture uses a hierarchical model in which edge devices conduct immediate inference while centralized systems maintain long-term analytics, model governance, historical storage, and cross-device coordination.

This design reduces unnecessary transmission while preserving centralized visibility. Periodic model updates can be distributed from the central layer to edge nodes, while only relevant events, summaries, or suspicious observations need to be transmitted continuously.

The hybrid principle is consistent with contemporary Industry 5.0 security thinking, in which multiple analytical mechanisms can cooperate to improve detection capability (Govindarajan et al., 2026).

3.8 Model Lifecycle and Adaptation

Industrial environments are dynamic. Production schedules, equipment aging, maintenance operations, and process changes can alter the distribution of legitimate observations. A static anomaly model can therefore become less effective over time.

The proposed lifecycle includes model training, edge deployment, monitoring, drift detection, validation, controlled updating, and rollback. Model updates should be validated before deployment because a poorly trained model could increase false positives or fail to identify emerging anomalies.

3.9 Hypothetical Manufacturing Scenario

Consider an automated production line containing motors, robotic equipment, temperature sensors, vibration sensors, and industrial controllers. Under normal production, the edge model learns that a specific vibration-temperature relationship occurs at a particular operating speed. Later, an unusual command sequence is accompanied by a sudden deviation in vibration and current.

The edge engine calculates separate process, sensor, and command anomaly scores. Because the deviations occur simultaneously, the combined score exceeds the high-risk threshold. The edge gateway immediately flags the event and restricts the affected communication path while transmitting an event summary to the central security system.

The example illustrates why contextual edge analysis can be more valuable than isolated threshold monitoring. The framework does not interpret one unusual measurement as proof of malicious behavior; instead, it evaluates

correlated deviations across the cyber and physical layers.

4. Results

The conceptual evaluation of the proposed framework indicates five principal findings. First, edge-based inference can substantially improve the responsiveness of anomaly detection because initial analysis occurs near the data-generation point. This characteristic is particularly important in cyber-physical manufacturing, where a delayed response may allow an abnormal command or machine state to propagate into the physical process. The distributed intelligence principle is consistent with the connected management architectures demonstrated in IoT applications (Nagaraja et al., 2019).

Second, the framework provides a mechanism for reducing unnecessary data transmission. Rather than continuously forwarding all sensor observations to centralized infrastructure, edge nodes can transform raw observations into compact features and transmit only significant anomalies, summaries, or model-relevant information. This architecture can therefore reduce communication dependency while maintaining centralized oversight. However, the amount of reduction depends on sensor density, sampling frequency, model complexity, and anomaly frequency.

Third, contextual anomaly scoring offers an advantage over independent threshold-based monitoring. Manufacturing systems frequently operate under multiple legitimate states, meaning that a fixed threshold may classify normal operational transitions as abnormal. The proposed combination of process, sensor, network, and command indicators provides a richer representation of system behavior. Predictive reasoning, as demonstrated in intelligent agricultural applications, supports this concept because deviations become meaningful relative to an expected state (Siddhartha and Lakkannavar, 2021).

Fourth, lightweight edge intelligence creates a practical trade-off between computational efficiency and analytical sophistication. Larger models may provide more expressive representations, but industrial gateways may not have sufficient resources to execute them continuously. Consequently, model compression, feature selection, and hierarchical inference become important architectural considerations. The framework therefore favors a division in which computationally efficient models perform immediate detection and more computationally intensive analysis remains available at centralized infrastructure.

Fifth, the proposed framework demonstrates that AI-enabled anomaly detection should be considered part of a broader cyber-physical security architecture rather than a standalone classification problem. The Industry 5.0 security perspective represented by Govindarajan et al. (2026) reinforces the importance of intelligent security mechanisms for smart manufacturing. Nevertheless, because the present framework has not been evaluated experimentally, claims regarding accuracy, latency reduction, or false-positive improvement remain hypotheses requiring empirical validation.

Overall, the findings suggest that edge intelligence is theoretically well aligned with industrial anomaly detection because it combines low-latency processing, distributed sensing, contextual analysis, and hierarchical coordination. The major unresolved issue is not whether edge AI can perform inference, but how reliably it can distinguish malicious behavior from legitimate process variation under changing industrial conditions.

5. Discussion

The proposed framework has both theoretical and practical implications. Theoretically, it positions industrial anomaly detection as a problem of behavioral deviation across interconnected cyber and physical dimensions. This differs from a purely network-centric view of security because the physical consequences of cyber events are incorporated into the detection process. Such integration is consistent with the broader movement toward computational intelligence in digitally transformed environments (Chen and Li, 2018).

A principal practical implication is reduced dependence on centralized processing. Edge analysis can allow manufacturing organizations to maintain immediate detection capability even when connectivity with central infrastructure is degraded. This is particularly important for systems where uninterrupted communication cannot be assumed. The architecture also supports data minimization because raw observations can remain local unless an event requires further investigation.

However, these advantages create several trade-offs. The first concerns model accuracy versus computational efficiency. Edge devices have limited resources, and overly complex models can increase inference time, memory consumption, and energy requirements. Lightweight models may therefore be operationally

preferable even when they are theoretically less expressive.

The second trade-off involves false positives and false negatives. A highly sensitive detector may identify more abnormal events but could also generate excessive alerts. In manufacturing, repeated false alarms can lead operators to ignore warnings, thereby weakening the practical security value of the system. Conversely, an overly conservative detector may fail to identify subtle attacks. Context-aware scoring can reduce this problem, but it cannot eliminate it.

The third issue is concept drift. Industrial environments evolve through maintenance, equipment replacement, production changes, and process optimization. A model trained under one operating distribution may become inappropriate under another. Continuous adaptation is therefore necessary, but uncontrolled online learning introduces its own security risks because an attacker could attempt to influence the model using manipulated observations.

The literature supplied by the user provides useful foundations but also exposes the limitations of current cross-domain knowledge. Cainthus (2019), J. Deere (2019), and Nagaraja et al. (2019) illustrate intelligent sensing and automated management, but they do not directly resolve industrial cybersecurity problems. Tahat et al. (2020) and the United Nations reference provide broader sustainability context, while PwC (2017) and WIPO (2019) establish the expanding importance of AI. These references support the technological rationale for the framework but cannot independently validate its security performance.

The relationship with hybrid Industry 5.0 intrusion detection is particularly important. Govindarajan et al. (2026) demonstrates the relevance of ensemble-oriented security mechanisms in smart manufacturing. The proposed edge architecture can complement such approaches by determining where inference occurs. In other words, hybrid intelligence and edge deployment address different dimensions of the same problem: one concerns how detection decisions are generated, while the other concerns where those decisions are computed.

The primary limitation of this study is its conceptual nature. No industrial dataset, hardware benchmark, attack simulation, or statistical comparison has been performed. Therefore, the framework should not be interpreted as evidence of superior detection accuracy. Future empirical studies should evaluate detection precision, recall, false-positive rate, inference latency, memory utilization, communication reduction, energy consumption, and resilience under changing operating conditions. Such validation would determine whether the theoretical advantages of edge AI translate into measurable security improvements.

6. Conclusion

This paper proposed an Edge-AI-Based Anomaly Detection Framework for securing Industrial IoT and cyber-physical manufacturing systems. The framework integrates distributed industrial sensing, edge preprocessing, lightweight AI inference, contextual anomaly scoring, localized response, and centralized coordination. Its central contribution is the conceptual integration of cyber and physical observations within a hierarchical security architecture.

The analysis demonstrates that edge intelligence can potentially reduce response latency, limit unnecessary data transmission, and support localized decision-making. The framework also emphasizes that industrial anomalies should be interpreted in context because legitimate manufacturing processes naturally exhibit variation. Combining process, sensor, network, and command-level indicators provides a more comprehensive basis for security-oriented anomaly detection.

At the same time, edge deployment does not eliminate the fundamental challenges of intelligent security. Resource limitations, false positives, model drift, heterogeneous devices, and the risk of inappropriate automated responses remain significant concerns. Consequently, edge AI should be viewed as one layer within a broader defense architecture rather than as a complete cybersecurity solution.

Future research should transform the conceptual framework into an experimentally validated implementation using representative industrial datasets and edge hardware. Particular attention should be given to adaptive model updating, explainable anomaly scores, federated or distributed learning, adversarial robustness, and safe automated response. Integration with hybrid ensemble approaches for Industry 5.0 security may further strengthen the ability of manufacturing systems to detect complex threats while maintaining operational continuity (Govindarajan et al., 2026).

References

1. Cainthus. (2019). Combining Computer Vision and Arti cial Intelli gence. Accessed: Aug. 21, 2019. [Online]. Available: https://www.cainthus.com/technology
2. J. Deere. (2019). AutoTrac Agricultural Management Solutions (AMS). Accessed: Aug. 21, 2019. [Online]. Available: https://www.deere.co.uk/en/ agricultural management-solutions/guidance automation/autotrac/
3. Y. Chen and Y. Li, Computational Intelligence Assisted Design (In Industrial Revolution 4.0). Boca Raton, FL, USA: CRC Press, 2018. [Online]. Available: 60676 https://www.taylorfrancis.com/books/97814987
4. L.Karthikeyan,I.Chawla,andA.K.Mishra,‘‘Areview ofremotesensing applications in agriculture for food security: Crop growth and yield, irri gation, and crop losses,’’ J. Hydrol., vol. 586, Jul. 2020, Art. no. 124905, doi:10.1016/j.jhydrol.2020.124905.
5. G.S. Nagaraja,A.B. Soppimath,T .Soumya ,and A. Abhinith, IoTbased smart agriculture management system, in Proc. 4th Int. Conf. Comput. Syst. Inf. Technol. Sustain. Solution (CSITSS), Dec. 2019, pp. 1 5, doi: 10.1109/CSITSS47250.2019.9031025.
6. M. M. Tahat, K. M. Alananbeh, Y. A. Othman, and D. I. Leskovar, ‘‘Soil health and sustainable agriculture,’’ Sustainability, vol. 12, no. 12, p. 4859, Jun. 2020, doi: 10.3390/su12124859.
7. E. Siddhartha and M. C. Lakkannavar, Smart irrigation and crop health prediction, in Proc. Int. Conf. Recent Trends Electron., Inf., Com mun. Technol. (RTEICT), Aug. 2021, pp. 739742, doi: 10.1109/RTE ICT52294.2021.9573542.
8. United Nation Global Population Growth and Sustainable Development. Accessed: Apr. 11, 2022. [Online]. Available: https://www.un.org/development/desa/pd/sites/www.un.org.development.desa.pd/ les/undesa_pd_2022_global_population_growth
9. PwC. Arti cial Intelligence Could Add 232bn to UK GDP by 2030. Accessed: Jun. 28, 2017. [Online]. Available: https://www.pwc.co.uk/ press-room/press-releases/arti cial-intelligence could-add-232bn-to-UK gdp.html
10. WIPO Technology Trends 2019: Arti cial Intelligence, World Intellectual Property Org., Geneva, Switzerland, 2019.
11. Govindarajan, V., Ahmed, F., Faheem, Z. B., Bilal, M., Ayadi, M., & Ali, J. (2026). Aegis-5: A Hybrid Ensemble Framework for Intrusion Detection in Industry 5.0 Driven Smart Manufacturing Environment. ACM Transactions on Autonomous and Adaptive Systems.
12. Ramamurthy, K., Gumber, S., Abdelfattah, W.M. et al. Human AI trust modeling in cognitive systems via ensemble learning and advanced feature engineering. Discov Artif Intell 6, 366 (2026). https://doi.org/10.1007/s44163-026-01255-7
13. Geo Philip, Paulson, Robotics-Enabled Sustainable Construction Management: A Socio-Technical Framework for Operational Efficiency, Digital Integration, and Sustainability Performance. Available at SSRN: https://ssrn.com/abstract=6845167 or http://dx.doi.org/10.2139/ssrn.6845167

.