

Adaptive Deep Learning-Based Intrusion Detection for Cyber-Physical Systems in Industry 5.0 Smart Manufacturing

Dr. Haruto Nakamura

Department of Artificial Intelligence, Tokyo
Institute of Advanced Computing, Tokyo, Japan

Dr. Yui Tanaka

Department of Intelligent Systems,
Osaka Institute of Technology, Osaka, Japan

VOLUME 03 ISSUE 08 (2026)

Published Date: 26 August 2026 // Page no.: - 98-103

ABSTRACT

Industry 5.0 smart manufacturing environments integrate cyber-physical systems (CPS), industrial communication networks, intelligent machines, connected sensors, edge computing resources, and human-centric automation. This convergence improves operational flexibility and responsiveness but simultaneously expands the attack surface and creates intrusion-detection challenges that conventional static security mechanisms cannot adequately address. This research proposes an adaptive deep learning-based intrusion detection framework for CPS-oriented Industry 5.0 manufacturing environments. The proposed approach conceptually combines multi-source industrial telemetry, adaptive feature representation, deep sequential learning, anomaly-oriented detection, and continuous model adaptation. Particular emphasis is placed on detecting evolving and previously unseen cyber threats while limiting false alarms and computational overhead at the industrial edge. The theoretical foundation is derived from the evolution of advanced wireless and interconnected communication infrastructures, which provide the connectivity required by highly distributed manufacturing systems. Existing work concerning future wireless communications, 5G architectures, and high-frequency communications demonstrates the progressive increase in network capability and heterogeneity, while the Aegis-5 framework further motivates hybrid intelligent intrusion detection for Industry 5.0 smart manufacturing (Govindarajan et al., 2026). The proposed framework is evaluated conceptually through detection capability, adaptability, computational efficiency, and operational resilience. The analysis indicates that an adaptive deep learning architecture can provide a stronger security foundation than fixed-rule or static classification approaches, particularly when manufacturing traffic patterns change dynamically. However, challenges involving data quality, model drift, computational constraints, explainability, and adversarial manipulation remain significant. The study therefore positions adaptive deep learning as a promising but carefully governed security mechanism for next-generation cyber-physical manufacturing.

Keywords: Industry 5.0, Cyber-Physical Systems, Intrusion Detection, Deep Learning, Adaptive Learning, Smart Manufacturing, Anomaly Detection, Industrial Cybersecurity, 5G Networks.

INTRODUCTION

Industry 5.0 represents an evolution of smart manufacturing in which automation, connectivity, intelligent computation, and human participation are increasingly integrated into a unified production ecosystem. Unlike isolated industrial automation, contemporary manufacturing infrastructures depend on continuous communication among sensors, programmable controllers, robotic systems, production platforms, edge devices, and cloud services. The evolution of wireless communication toward highly connected architectures provides an important technological foundation for this transformation. Earlier research on communications beyond third-generation systems

already anticipated increasingly heterogeneous and interconnected wireless environments (Aretz et al., 2001; Otsu, 2003). Subsequent developments in 5G further strengthened the capacity for high-density connectivity and intelligent networked applications (Andrews et al., 2014; Gupta & Jha, 2015).

This connectivity, however, produces a fundamental security dilemma. A manufacturing environment can gain operational intelligence and flexibility while becoming increasingly exposed to unauthorized access, traffic manipulation, denial-of-service behavior, compromised endpoints, and anomalous machine-to-machine communication. In CPS environments, the consequences of intrusion extend beyond information confidentiality. Manipulation of digital commands can influence physical processes, production quality, equipment

safety, energy consumption, and operational continuity. Consequently, intrusion detection must address both cyber-network behavior and the temporal characteristics of physical processes.

Traditional intrusion detection systems generally depend on predefined signatures, static rules, or machine-learning models trained under relatively stable conditions. Such approaches can be effective for known attack patterns but may become less reliable when legitimate industrial behavior changes because of production schedules, equipment states, maintenance operations, network reconfiguration, or newly introduced devices. An Industry 5.0 environment therefore requires an intrusion detection mechanism capable of learning from evolving data rather than treating the operational environment as static.

The growing importance of intelligent intrusion detection in smart manufacturing is reflected in recent work proposing hybrid ensemble approaches specifically for Industry 5.0 environments. Govindarajan et al. (2026) demonstrate the relevance of combining multiple learning perspectives for intrusion detection in Industry 5.0-driven smart manufacturing. Building on this direction, the present research focuses specifically on adaptivity: the ability of a deep learning-based IDS to continuously accommodate changing communication and operational patterns.

The primary objective of this study is to develop a theoretically grounded adaptive deep learning framework for CPS-based smart manufacturing. The study addresses four questions: how industrial telemetry can be represented for deep learning; how temporal and anomalous behavior can be detected; how models can adapt to changing operational conditions; and how detection accuracy can be balanced against latency, computational cost, and false alarms.

The significance of this research lies in connecting the technological evolution of highly connected networks with an intelligent cybersecurity architecture. As communication infrastructures become increasingly capable and heterogeneous, including emerging high-frequency communication environments, security mechanisms must become equally adaptive (Rappaport et al., 2019). The proposed framework therefore considers intrusion detection as a continuous cyber-physical learning process rather than a one-time classification task.

LITERATURE REVIEW

The development of smart manufacturing cannot be separated from the evolution of communication technologies. Aretz et al. (2001) examined wireless communications beyond third-generation systems and

emphasized the future need for advanced communication capabilities. Otsu (2003) similarly discussed systems beyond IMT-2000, illustrating the transition toward increasingly sophisticated wireless environments. These developments are important to CPS security because higher connectivity creates more communication relationships among industrial components.

Zahariadis (2004) discussed migration toward fourth-generation wireless communications, while Andrews et al. (2014) examined the requirements and possibilities associated with 5G. Gupta and Jha (2015) subsequently synthesized 5G network architecture and emerging technologies. Collectively, these works establish a technological progression from earlier mobile systems toward highly connected and flexible network architectures. For Industry 5.0, this progression means that manufacturing systems can increasingly rely on distributed communication rather than isolated industrial networks.

The relationship between communication capability and cybersecurity is therefore indirect but important. Increased connectivity improves accessibility and responsiveness, but it can also increase the number of communication paths that must be monitored. An intrusion detection architecture operating within such an environment cannot rely exclusively on fixed network boundaries. Instead, it must interpret heterogeneous traffic and distinguish normal operational changes from malicious deviations.

Rappaport et al. (2019) extend this technological trajectory by discussing wireless communications above 100 GHz and opportunities associated with 6G and beyond. Their work illustrates that future communication environments are expected to support increasingly demanding applications. In a smart manufacturing context, such developments may intensify the volume and velocity of telemetry generated by connected industrial devices. Consequently, security analysis must increasingly operate close to where data are generated.

The literature also exposes an important methodological gap. The provided communication-oriented studies primarily establish the evolution of connectivity rather than directly addressing adaptive CPS intrusion detection. The Aegis-5 study provides a closer security-oriented positioning by introducing a hybrid ensemble framework for intrusion detection in Industry 5.0-driven smart manufacturing (Govindarajan et al., 2026). Its relevance to the present study lies in recognizing that Industry 5.0 intrusion detection benefits from combining intelligent learning mechanisms rather than depending exclusively on a single detection perspective.

Nevertheless, a hybrid detection architecture and an adaptive deep learning architecture address partially different problems. Ensemble learning can improve robustness by combining complementary models, whereas adaptive deep learning emphasizes representation learning and continuous accommodation of changing distributions. The latter becomes

particularly important when manufacturing traffic changes over time. A model trained on one production state may incorrectly interpret a legitimate operational transition as an attack, while a static model may fail to identify a new attack whose characteristics differ from previously observed samples.

The central research gap is therefore the integration of deep feature learning with controlled adaptation for CPS environments. The proposed research positions adaptive learning between two extremes: a completely static IDS that cannot respond effectively to environmental changes and an unrestricted continuously learning system that may itself learn malicious or corrupted behavior. The desired architecture must adapt selectively, preserve trusted knowledge, and maintain security-oriented decision boundaries.

METHODOLOGY

3.1 Research Design and Framework Architecture

This study proposes a conceptual architecture named the Adaptive Deep CPS Intrusion Detection Framework (ADC-IDF). The framework is designed around five functional layers: data acquisition, preprocessing and representation, deep detection, adaptive intelligence, and security response.

The data acquisition layer collects heterogeneous observations from CPS components. These may include network packets, flow statistics, sensor readings, actuator states, authentication events, controller communications, device status indicators, and application-level industrial telemetry. Instead of treating network traffic as an independent stream, the framework associates cyber observations with operational context. This is important because an unusual packet sequence may be benign during maintenance but suspicious during normal production.

The second layer performs preprocessing and feature representation. Raw industrial data can contain missing values, redundant attributes, inconsistent sampling intervals, and highly imbalanced attack classes. Normalization and temporal alignment are therefore necessary before model training. Rather than depending entirely on manually engineered features, the proposed architecture uses deep representation learning to derive higher-level behavioral characteristics.

3.2 Deep Learning Detection Mechanism

The core detection layer combines spatial and temporal learning. A convolutional component can identify local patterns in structured traffic or feature sequences, whereas a recurrent or attention-oriented component can capture temporal dependencies. This is particularly

suitable for CPS because malicious activity may not be distinguishable from a single observation. Instead, an attack can emerge through a sequence of small deviations.

The model can be expressed conceptually as:

$$H_t = f_{\theta}(X_{t-k:t}) \quad H_t = f_{\theta}(X_{t-k:t})$$

where $X_{t-k:t}$ represents a window of observations over a defined temporal interval, f_{θ} denotes the deep representation model, and H_t represents the learned behavioral representation at time t .

The resulting representation is passed to a detection layer that produces a classification or anomaly score:

$$S_t = g_{\phi}(H_t) \quad S_t = g_{\phi}(H_t)$$

where S_t represents the estimated degree of maliciousness. A threshold mechanism then determines whether the observation is normal, suspicious, or malicious.

This structure differs from purely signature-based IDS operation. Instead of asking whether a packet exactly matches a known attack signature, the model estimates whether the observed behavior is consistent with learned operational patterns.

3.3 Adaptive Learning Mechanism

Adaptivity is the central methodological contribution. Industrial environments exhibit concept drift because their legitimate behavior changes over time. New machines may be introduced, production lines may be reconfigured, communication loads may change, and maintenance activities may produce temporary anomalies.

The proposed framework therefore maintains a monitored distribution of recent observations. Let $P_{train}(X)$ represent the training distribution and $P_t(X)$ represent the current operational distribution. When the difference between these distributions exceeds a predefined tolerance, an adaptation process is triggered:

$$D(P_{train}, P_t) > \delta \quad D(P_{train}, P_t) > \delta$$

where D is a distribution-difference measure and δ is an adaptation threshold.

However, adaptation must not occur blindly. Automatically retraining on all incoming traffic could allow compromised observations to influence the model. The framework therefore introduces a trusted adaptation buffer. Only observations with sufficiently high confidence and verified operational consistency are considered for incremental learning. Suspicious observations are isolated for further analysis rather than immediately incorporated into the model.

This controlled adaptation principle is essential for industrial cybersecurity because the system must remain responsive without sacrificing model integrity. Govindarajan et al. (2026) provide a useful conceptual basis for intelligent intrusion detection in Industry 5.0 smart manufacturing, while the present architecture extends this perspective toward continuous behavioral adaptation.

3.4 Edge-Oriented Deployment

Industrial intrusion detection cannot always transfer all telemetry to a centralized cloud because latency, bandwidth, privacy, and operational continuity may impose constraints. The proposed architecture therefore supports hierarchical deployment.

An edge-level detector performs rapid screening close to industrial devices. More computationally intensive analysis can be executed at an intermediate industrial server or centralized platform. This arrangement creates a division between immediate detection and deeper investigation.

The communication evolution described by Andrews et al. (2014) and Gupta and Jha (2015) demonstrates why such distributed architectures are increasingly practical. At the same time, increased network capacity should not be interpreted as eliminating computational constraints. Industrial security systems must still prioritize low latency and predictable behavior.

3.5 Evaluation Framework

Because the present study is a methodological framework rather than a report of a newly executed industrial experiment, evaluation is defined analytically rather than through fabricated numerical results. The principal performance dimensions are accuracy, precision, recall, F1-score, false-positive rate, detection latency, adaptation time, computational overhead, and resilience to distribution changes.

A balanced evaluation should distinguish between conventional detection accuracy and adaptive performance. A model may achieve high accuracy on a stationary benchmark while failing when production behavior changes. Therefore, evaluation should include sequential time windows, changing operational states, and previously unseen attack conditions.

The F1-score can be represented as:

$$F1 = \frac{2PR}{P+R}$$

where PP denotes precision and RR denotes recall. In CPS environments, both are important. Low recall can allow attacks to remain undetected, while low precision can

generate excessive alarms and overwhelm security operators.

RESULTS

The analytical findings indicate that the proposed adaptive architecture provides a stronger conceptual fit for Industry 5.0 CPS environments than a static intrusion detection design. The principal finding is that adaptation should be treated as a controlled security function rather than an unrestricted machine-learning process. Manufacturing systems continuously change, and a detector that assumes stationary traffic distributions is likely to experience increasing false positives or declining sensitivity as operational conditions evolve.

A second finding concerns temporal representation. CPS attacks cannot always be characterized by isolated observations. Sequential changes in controller communications, device states, and network flows can provide more meaningful behavioral evidence than individual packets. A deep architecture capable of learning temporal dependencies is consequently better aligned with the operational structure of cyber-physical environments. This complements the broader evolution toward increasingly interconnected communication systems described by Aretz et al. (2001), Otsu (2003), and Andrews et al. (2014).

The third finding is that adaptive learning creates a security trade-off. Continuous adaptation can improve performance under concept drift, but it simultaneously creates a potential poisoning pathway. If malicious traffic is incorporated into the training set, the IDS may gradually redefine abnormal behavior as legitimate. The proposed trusted adaptation buffer addresses this problem by separating high-confidence operational observations from suspicious events.

A fourth finding concerns deployment. A centralized detector may provide greater computational resources, but edge-oriented analysis is more compatible with time-sensitive CPS operations. The increasing capacity of advanced wireless networks, including future high-frequency communication systems, can support large-scale connectivity; nevertheless, security decisions that directly affect industrial operations should remain capable of being executed locally (Rappaport et al., 2019).

The proposed approach also complements hybrid intelligent IDS research in Industry 5.0. Aegis-5 demonstrates the relevance of hybrid ensemble learning for smart manufacturing intrusion detection (Govindarajan et al., 2026). The present framework adds an adaptive dimension by emphasizing temporal representation, distribution monitoring, and controlled incremental learning.

However, no claim is made that deep learning automatically guarantees superior security. Performance remains dependent on representative training data, feature quality, computational resources, threshold selection, and model governance. In

addition, explainability remains a practical concern because manufacturing operators may require understandable evidence before accepting an automated security decision. The findings therefore support adaptive deep learning as a promising architectural direction rather than as a universally optimal solution.

DISCUSSION

The central theoretical implication of this research is that intrusion detection for Industry 5.0 should be understood as a dynamic cyber-physical inference problem. Conventional IDS formulations often emphasize the classification of network events, whereas CPS environments require interpretation of interactions among network, computational, and physical states. A change in communication behavior becomes meaningful only when considered in relation to production context.

This argument is consistent with the technological progression identified in the literature. As communication systems evolved from third-generation architectures toward 4G, 5G, and emerging beyond-5G/6G environments, connectivity became more distributed and application-oriented (Zahariadis, 2004; Andrews et al., 2014; Rappaport et al., 2019). Industry 5.0 extends this trajectory into manufacturing, where communication is no longer merely an information-transfer mechanism but an operational component of production.

The practical implication is that an effective IDS should not operate as an isolated cybersecurity appliance. It should exchange contextual information with manufacturing systems while preserving strict security boundaries. For example, a sudden increase in traffic could represent an attack, but it could also result from a legitimate production-line reconfiguration. Adaptive learning can reduce such ambiguity when the model continuously observes trusted operational states.

The comparison with Govindarajan et al. (2026) is particularly relevant. Their Aegis-5 framework emphasizes hybrid ensemble intelligence for Industry 5.0 intrusion detection. The proposed ADC-IDF does not reject ensemble learning; rather, it provides a complementary direction in which adaptive deep representations can serve as a continuously updated behavioral layer. A future implementation could therefore combine adaptive deep learning with ensemble decision mechanisms to increase robustness.

A significant trade-off concerns adaptability versus stability. High adaptation rates allow rapid response to legitimate changes but increase the risk of model instability. Conversely, conservative adaptation improves stability but may delay recognition of genuine environmental changes. This suggests that adaptation

should be confidence-driven and policy-controlled rather than triggered solely by time.

Another limitation concerns data imbalance. Industrial attack datasets may contain substantially fewer malicious observations than normal operational observations. A detector optimized only for overall accuracy could therefore achieve apparently strong performance while missing rare but critical attacks. Evaluation should consequently emphasize recall, F1-score, false-positive rate, and detection latency rather than accuracy alone.

Explainability is another unresolved challenge. Deep learning models can identify complex patterns but may provide limited reasoning in forms immediately understandable to plant operators. In high-consequence manufacturing environments, a useful IDS should ideally provide contextual evidence such as affected device, abnormal communication sequence, temporal deviation, and confidence level.

Finally, the framework remains vulnerable to adversarial manipulation and poisoned training data. Adaptivity is beneficial only when the learning pipeline itself is trustworthy. Future implementations should therefore incorporate model-version control, trusted data validation, adaptation auditing, rollback mechanisms, and human verification for high-impact updates.

CONCLUSION

This research presented an adaptive deep learning-based intrusion detection framework for cyber-physical systems operating within Industry 5.0 smart manufacturing environments. The proposed architecture addresses a fundamental weakness of static IDS approaches: their limited ability to accommodate continuously changing industrial behavior. By integrating heterogeneous telemetry, deep temporal representation, anomaly scoring, controlled model adaptation, and edge-oriented deployment, the framework establishes a conceptual foundation for security systems capable of responding to evolving cyber threats.

The literature demonstrates a clear progression toward increasingly connected communication infrastructures, from systems beyond third generation to 4G, 5G, and emerging 6G-oriented environments (Aretz et al., 2001; Otsu, 2003; Gupta & Jha, 2015; Rappaport et al., 2019). Industry 5.0 builds upon this connectivity while introducing more complex interactions between digital intelligence and physical production. Consequently, intrusion detection must evolve from static network classification toward context-sensitive and adaptive cyber-physical intelligence.

The study further positions adaptive deep learning as complementary to hybrid intelligent IDS approaches such as Aegis-5 (Govindarajan et al., 2026). Its principal contribution is the emphasis on controlled adaptation: the detector should learn from changing legitimate behavior without allowing

unverified or malicious observations to redefine its security model.

Future research should implement the framework using real industrial datasets and evaluate it under controlled concept-drift scenarios, zero-day attack conditions, heterogeneous device populations, and edge-computing constraints. Particular attention should be given to explainable deep learning, secure incremental learning, adversarial robustness, and model rollback. Such developments can help transform adaptive IDS from a conceptual architecture into a dependable security component for resilient, intelligent, and human-centric Industry 5.0 manufacturing.

REFERENCES

1. Andrews, j. G., Buzzzi, s., choi, w., Hanly, s. V., Lozano, a., Soong, a. C., & Zhang, j. C. (2014). What will 5g be? Ieee journal on selected areas in communications, 32(6), 1065- 1082.
2. Andrews, J. G., Buzzzi, S., Choi, W., Hanly, S. V., Lozano, A., Soong, A. C., & Zhang, J. C. (2014). What will 5G be? IEEE Journal on Selected Areas in Communications, 32(6), 1065-1082.
3. Govil, and j. Govil, on the investigation of transactional and interoperability issues between ipv4 and ipv6, ieee eit may 2007. • development of 3g mobile services, OECD report, sept. 2004.
4. Gupta, a., & Jha, r. K. (2015). A survey of 5g network: architecture and emerging technologies. Ieee access, 3, 1206-1232.
5. Gupta, A., & Jha, R. K. (2015). A survey of 5G network: Architecture and emerging technologies. IEEE Access, 3, 1206-1232.
6. K. Aretz, m. Haardt, w. Kornhauser, and w. Mohr, "the future of wireless communications beyond the third generation", international journal computer and telecommunications networking (computer networks), vol. 37, 2001, pp. 83-92.
7. T. Otsu, "the challenge of systems beyond imt 2000—approach from wireless", international telecommunications union journal, vol. 33, no. 3, 2003, pp. 26-30.
8. Rappaport, t. S., Xing, y., kanhere, o., Ju, s., Mahanayake, a., mandal, s., & wang, g. (2019). Wireless communications and applications above 100 ghz: opportunities and challenges for 6g and beyond. Ieee access, 7, 78729-78757.
9. Rappaport, T. S., Xing, Y., Kanhere, O., Ju, S., Mahanayake, A., Mandal, S., & Wang, G. (2019). Wireless communications and applications above 100 ghz: Opportunities and challenges for 6G and beyond. IEEE Access, 7, 78729-78757
10. T. Wang, x. Shan, x and y. Ren, "a new approach for evaluating clipping distortion in ds cdma systems". Ieee transactions on communications. Vol. E88-b, no. 2, 2005 pp.
11. T. B. Zahariadis, "migration towards 4g wireless communications," ieee wireless communications, vol. 11, no. 3, 2004, pp. 6-7.
12. 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.
13. Ramamurthy, K. (2023). AI-Driven Test Automation Frameworks for the Modern Software Quality Engineering. International Journal of Emerging Trends in Computer Science and Information Technology, 4(4), 257-269.
14. Philip, P. G. (2024). Artificial Intelligence-Driven Project Risk Prediction Models: Enhancing Decision-Making Accuracy in Large-Scale Infrastructure Projects . American Journal of Technology, 3(1), 52-69. <https://doi.org/10.58425/ajt.v3i1.571>.