

Federated Ensemble Learning for Privacy-Preserving Intrusion Detection in Industry 5.0-Enabled Smart Factories

Dr. Arjun Mehta

Department of Artificial Intelligence, National
Institute of Advanced Computing, Bengaluru, India

Dr. Kavya Sharma

Department of Intelligent Systems, Institute
of Technology and Innovation, New Delhi, India

VOLUME 03 ISSUE 08 (2026)

Published Date: 27 August 2026 // Page no.: - 01-07

ABSTRACT

Industry 5.0-enabled smart factories integrate industrial Internet of Things (IIoT) devices, cyber-physical production systems, intelligent automation, collaborative robotics, and data-driven decision mechanisms. This interconnected environment improves manufacturing flexibility and operational intelligence but simultaneously enlarges the attack surface and creates substantial challenges for centralized intrusion detection. Conventional centralized approaches require sensitive operational data to be transferred to a common repository, creating privacy, bandwidth, governance, and single-point-of-failure concerns. This paper proposes a conceptual Federated Ensemble Learning (FEL) framework for privacy-preserving intrusion detection in Industry 5.0 smart factories. The framework combines federated learning with heterogeneous ensemble detection so that manufacturing sites can collaboratively improve detection models without directly sharing raw industrial traffic. Blockchain-inspired trust mechanisms are incorporated to support decentralized model accountability, while lightweight identity management is considered for resource-constrained IIoT environments. The framework is theoretically positioned around distributed learning, collaborative intelligence, privacy preservation, and resilient cyber-physical manufacturing. The analysis indicates that federated ensemble learning can provide a more appropriate architectural balance between local data sovereignty and cross-factory threat intelligence than purely centralized detection. However, communication overhead, non-identically distributed industrial traffic, poisoned model updates, computational constraints, and coordination complexity remain significant limitations. The proposed framework therefore represents a research-oriented architecture rather than a claim of experimentally validated performance. Its principal contribution is the integration of privacy-preserving collaborative learning and ensemble-based intrusion detection into an Industry 5.0 manufacturing context.

Keywords: Federated Learning, Ensemble Learning, Intrusion Detection, Industry 5.0, Smart Factory, IIoT Security, Privacy Preservation, Blockchain, Cyber-Physical Systems.

INTRODUCTION

The transition toward Industry 5.0 is characterized by increasingly connected, intelligent, and collaborative manufacturing environments in which industrial machines, sensors, controllers, robots, edge devices, and information systems continuously exchange operational information. Such connectivity creates opportunities for adaptive production, predictive intelligence, and human-machine collaboration, but it also introduces a broader and more heterogeneous cybersecurity surface. In a smart factory, an intrusion may affect not only information confidentiality but also production continuity, equipment behavior, process integrity, and safety.

The security problem becomes particularly difficult when

industrial organizations attempt to detect attacks through centralized machine-learning architectures. A centralized intrusion detection system generally requires network or operational data from multiple manufacturing environments to be transferred to a common analytical infrastructure. Such data may contain sensitive information concerning production processes, equipment configurations, operational patterns, or organizational behavior. Consequently, increasing detection capability through greater data aggregation may conflict with the requirement to preserve industrial data sovereignty.

Distributed architectures offer an alternative. Federated learning enables participating organizations or devices to train models locally and communicate model-related information rather than directly exchanging their raw datasets. In the proposed context, this principle can be combined with ensemble learning, allowing multiple complementary

detection models to contribute to a more robust intrusion-detection decision. The relevance of ensemble-based approaches to Industry 5.0 intrusion detection is also reflected in the hybrid ensemble perspective presented by Govindarajan et al. (2026), which positions heterogeneous learning mechanisms as a means of addressing complex attack patterns in smart manufacturing environments. (Govindarajan et al., 2026)

The broader technological foundation for such a framework can also be understood through the literature on blockchain, decentralized identity, and distributed information management. Blockchain has been investigated as a mechanism for decentralized coordination and trusted transaction processing (Mingxiao et al., 2017), while lightweight blockchain-based identity management has been proposed for IoT environments where device identity and resource limitations are important considerations (Bouras et al., 2021). The decentralized characteristics of blockchain are therefore potentially relevant to the governance and accountability layer of a federated security architecture.

This paper proposes a conceptual Federated Ensemble Learning framework for privacy-preserving intrusion detection in Industry 5.0-enabled smart factories. The primary objective is to establish how local industrial intelligence can be retained while enabling collaborative improvement of intrusion detection. The paper specifically examines the architectural interaction among local learning, ensemble detection, federated aggregation, privacy, identity management, and decentralized trust.

The scope is intentionally architectural and analytical. It does not claim experimentally measured detection accuracy because no experimental dataset, implementation environment, or benchmark results were supplied. Instead, the Results and Findings section evaluates the expected behavior and design implications of the proposed architecture.

LITERATURE REVIEW

The provided literature establishes several technological foundations relevant to the proposed framework, although none individually provides a complete federated intrusion-detection architecture for Industry 5.0 smart factories.

Mingxiao et al. (2017) examine blockchain consensus algorithms and emphasize the importance of consensus mechanisms in decentralized blockchain systems. Their work is relevant to the proposed framework because federated security environments require some mechanism through which participating nodes can establish agreement about shared information. However, blockchain consensus and federated model aggregation

solve different problems. Consensus primarily establishes agreement within a distributed ledger, whereas federated aggregation combines machine-learning model updates. Treating the two as equivalent would therefore be conceptually incorrect. Instead, blockchain principles can potentially complement federated learning by providing an auditable coordination layer.

Guo and Liang (2016) discuss blockchain applications and prospects in the banking industry, demonstrating how distributed ledger technologies can support trusted information exchange among organizational participants. Although their domain is financial rather than industrial manufacturing, the underlying issue of trust among independent entities is relevant. In a multi-factory environment, participating organizations may be unwilling to rely on a centralized authority for security-model coordination. A decentralized trust layer could consequently improve accountability. Nevertheless, financial transactions and industrial telemetry differ substantially, meaning that mechanisms designed for banking cannot be transferred directly to IIoT systems without adaptation.

Huckle et al. (2016) examine relationships among IoT, blockchain, and shared-economy applications. Their work supports the broader theoretical connection between IoT environments and decentralized technologies. For smart factories, this relationship is significant because industrial devices continuously produce distributed data, while security intelligence may need to be coordinated across organizational boundaries. The proposed FEL framework extends this conceptual relationship from decentralized data environments toward decentralized security learning.

Bouras et al. (2021) propose a lightweight blockchain-based IoT identity-management approach. Identity is particularly important for federated intrusion detection because a malicious participant could otherwise submit fraudulent model updates or impersonate legitimate industrial devices. The lightweight emphasis is also relevant to IIoT systems, where sensors and embedded devices may have constrained computational and communication resources. Thus, identity management should be treated as an architectural security layer rather than as an independent feature.

Gordon and Catalini (2018) investigate blockchain technology for healthcare and patient-driven interoperability. Although healthcare is outside the manufacturing domain, their study illustrates how decentralized technologies can facilitate controlled interoperability among entities that need to exchange information while retaining organizational or individual control. This principle maps conceptually onto collaborative industrial security, where factories may need to benefit from collective intelligence without exposing raw operational data.

Ocheja et al. (2019) explore blockchain-based management of

lifelong learning records. Their work further demonstrates how distributed ledgers can support persistent, verifiable records across decentralized participants. For the proposed framework, this concept can be translated into model-governance records, such as model version, participating node, aggregation event, or update provenance. However, an industrial intrusion-detection system requires substantially more dynamic processing than educational record management.

Saberi et al. (2019) investigate blockchain technology and sustainable supply-chain management. Their analysis is particularly relevant to smart manufacturing because modern production environments are increasingly integrated with broader supply-chain ecosystems. Cybersecurity intelligence may therefore need to extend beyond an individual factory toward suppliers, logistics systems, and interconnected production partners. Blockchain-based trust mechanisms can potentially support this distributed ecosystem, although their computational and governance overhead must be carefully controlled.

Nakamoto (2008) establishes the foundational concept of decentralized digital transactions through a peer-to-peer architecture. For the proposed system, the most relevant conceptual contribution is not cryptocurrency itself but decentralized coordination without dependence on a single centralized intermediary. The same architectural principle motivates the use of distributed trust in collaborative intrusion detection.

Finally, Garrison (2016) provides a framework for inquiry-based learning in the context of e-learning. Although it is not a cybersecurity reference, its relevance to this study is indirect: distributed participants can contribute knowledge through structured interaction rather than relying exclusively on a single knowledge source. This provides a conceptual analogy for collaborative learning, but it does not constitute evidence for federated intrusion detection.

The literature therefore reveals a clear research gap. Existing references separately address blockchain consensus, decentralized trust, IoT identity, interoperability, distributed record management, supply-chain coordination, and learning frameworks. The proposed study positions federated ensemble learning as the integrating mechanism that connects these concepts with privacy-preserving intrusion detection. The hybrid ensemble perspective of Govindarajan et al. (2026) further supports the importance of combining multiple detection mechanisms for complex Industry 5.0 security environments. (Govindarajan et al., 2026)

METHODOLOGY

3.1 Research Design

This study adopts a conceptual framework-development methodology. Rather than presenting fabricated experimental measurements, it derives an architectural model from the theoretical relationships identified in the supplied literature. The proposed architecture consists of six logical layers: industrial data acquisition, local preprocessing, local ensemble detection, federated coordination, trust and identity management, and security response.

The design principle is straightforward: raw industrial data should remain at its originating site whenever possible, while useful learning information should be collaboratively exchanged through controlled model updates. This separation attempts to address the tension between collective cybersecurity intelligence and industrial data confidentiality.

3.2 Industrial Data Acquisition Layer

The first layer contains sensors, programmable logic controllers, industrial gateways, robots, manufacturing workstations, and other IIoT endpoints. These devices generate heterogeneous observations such as network packets, connection behavior, protocol characteristics, access events, command sequences, and device-level telemetry.

The heterogeneous nature of industrial traffic is important because normal behavior varies considerably between devices. A robotic controller may exhibit communication patterns that are fundamentally different from those of an engineering workstation. Consequently, a single generalized detector may not adequately represent all operational contexts.

The proposed framework therefore performs initial processing locally. Noise removal, feature transformation, normalization, and event construction occur inside the participating factory. This approach minimizes unnecessary movement of raw operational information.

3.3 Local Ensemble Intrusion Detection

Each participating factory maintains a local ensemble consisting of multiple complementary detection models. Conceptually, the ensemble may contain a behavioral detector, an anomaly detector, and a classification-oriented detector. The exact algorithms can vary according to local computational resources and available training data.

The principal theoretical assumption is that different models capture different aspects of malicious behavior. An anomaly-oriented model may identify deviations from established operational patterns, while a classification-oriented model may perform better when previously characterized attack classes are available.

For example, suppose a smart factory has historically observed stable communication between a programmable controller

and a production gateway. A sudden change in connection frequency may be recognized by a behavioral detector, while a known malicious traffic pattern may be identified by a supervised classifier. Combining these decisions can reduce dependence on the weaknesses of an individual model.

This design is consistent with the hybrid ensemble orientation described by Govindarajan et al. (2026), where multiple learning mechanisms are positioned to address complex intrusion-detection requirements in Industry 5.0 manufacturing. (Govindarajan et al., 2026)

3.4 Federated Learning Layer

After local training, participating factories do not directly upload their raw datasets. Instead, the local learning process produces model parameters or updates. A federated coordinator receives eligible updates and generates an improved global model or ensemble representation.

A simplified aggregation process can be represented as:

$$W_{global} = \sum_{i=1}^{N} \frac{n_i}{\sum_{j=1}^N n_j} W_i$$

where W_i represents the local model update of factory i , n_i denotes the effective local training contribution, and N represents the number of participating factories.

The resulting global model is subsequently distributed to participating sites, where it can be refined using locally observed traffic. This creates a continuous learning cycle.

However, industrial datasets are unlikely to be identically distributed. One factory may manufacture automobiles, another electronic components, and another pharmaceutical equipment. Their network behavior and attack exposure can differ substantially. Therefore, simple aggregation may not always be optimal. The framework should support weighting, local adaptation, or cluster-specific models when substantial heterogeneity exists.

3.5 Privacy-Preserving Mechanism

Federated learning reduces direct exposure of raw data but should not automatically be interpreted as complete privacy protection. Model updates may potentially reveal information about local training conditions. Therefore, privacy preservation should be treated as a layered architectural objective.

The framework recommends limiting the information transmitted from each factory to the minimum required for collaborative learning. Local access controls, secure communication, update validation, and potentially additional privacy-enhancing mechanisms can be

incorporated.

The decentralized perspective of blockchain literature is relevant here because the system must coordinate multiple parties without forcing all operational information into one centralized repository (Huckle et al., 2016; Guo & Liang, 2016).

3.6 Trust, Identity, and Update Validation

Federated intrusion detection introduces an important security paradox: the collaborative learning infrastructure itself can become an attack target. A compromised participant could intentionally submit manipulated updates to degrade the global model.

The proposed architecture therefore introduces a trust layer in which every participating factory and authorized device has a verifiable identity. Lightweight identity management is particularly relevant to resource-constrained IoT environments (Bouras et al., 2021).

Blockchain-inspired records may additionally maintain update provenance. A record could associate a model update with a participating identity, timestamp, version, and validation status. Blockchain consensus concepts provide a theoretical basis for establishing distributed agreement about shared records (Mingxiao et al., 2017). However, the framework does not require every model parameter to be stored on a blockchain. Such a design would introduce unnecessary storage and communication overhead.

3.7 Decision and Response Layer

The final layer converts model outputs into operational security decisions. A multi-stage decision mechanism can assign events to categories such as normal, suspicious, or high-confidence malicious activity.

A hypothetical example illustrates the architecture. If several factories observe variants of the same reconnaissance behavior, each factory can locally train on its observations. The resulting model updates are aggregated, allowing other factories to benefit from the generalized behavioral pattern without receiving the original packet captures.

The framework consequently transforms isolated factory-level observations into collective security intelligence while preserving local data boundaries.

RESULTS

Because this study is conceptual and no experimental dataset or implementation results were supplied, the findings represent architecture-level analytical outcomes rather than measured performance metrics. This distinction is important for research validity: reporting numerical accuracy, precision, recall, latency, or false-positive rates without empirical experimentation would constitute unsupported claims.

The first major finding is that federated ensemble learning provides a stronger conceptual alignment with Industry 5.0 than a purely centralized detection architecture. Smart factories generate highly heterogeneous operational data, and transferring all such information to a central repository creates data-governance and confidentiality concerns. Federated learning allows local data to remain within organizational boundaries while enabling collaborative model improvement. This aligns with the broader decentralized information-management principles represented in blockchain-oriented literature (Huckle et al., 2016; Gordon & Catalini, 2018).

The second finding is that ensemble learning can address the diversity of industrial attack manifestations more effectively than dependence on a single detection model. Different production systems generate different baseline behaviors, and attacks may range from recognizable patterns to previously unseen anomalies. A heterogeneous ensemble provides architectural flexibility because individual models can specialize in different forms of evidence. The hybrid ensemble concept is similarly central to the Industry 5.0 intrusion-detection perspective of Govindarajan et al. (2026). (Govindarajan et al., 2026)

The third finding concerns trust. Federated learning without participant validation creates a potential vulnerability because malicious participants can attempt to influence the collaborative model. Consequently, identity management and update provenance become essential rather than optional features. The lightweight blockchain-based identity approach described by Bouras et al. (2021) provides a useful conceptual foundation for this requirement. Blockchain's decentralized coordination principles also support the argument for auditable participation and distributed trust (Mingxiao et al., 2017).

The fourth finding is that privacy and security cannot be treated as identical objectives. Keeping raw data local improves data sovereignty, but it does not eliminate all information-leakage or model-manipulation risks. Similarly, introducing blockchain does not automatically make the entire learning system secure. The architecture must therefore combine local data protection, participant authentication, update validation, secure communication, and model robustness.

A fifth finding is the existence of an unavoidable efficiency trade-off. More sophisticated ensemble models and decentralized validation mechanisms may increase computational, storage, and communication requirements. This is particularly significant for constrained IIoT devices. The lightweight orientation found in IoT identity research (Bouras et al., 2021) consequently has direct implications for the proposed framework.

Finally, the analysis indicates that decentralized security learning is most appropriate when participating organizations have both a common security interest and sufficient trust to collaborate. Supply-chain research involving blockchain similarly demonstrates the importance of coordination among multiple independent entities (Saberi et al., 2019). Thus, federated intrusion detection should be implemented as a governed collaborative ecosystem rather than simply as a machine-learning algorithm.

DISCUSSION

The proposed architecture has several theoretical implications. First, it shifts the conventional intrusion-detection paradigm from data centralization toward knowledge collaboration. In centralized machine learning, increasing model capability often depends on accumulating larger datasets in a common location. Federated learning reverses this relationship: intelligence can be aggregated while the original data remains distributed. This distinction is particularly important in smart manufacturing because operational data may have strategic or commercially sensitive characteristics.

Second, the framework integrates three concepts that are frequently considered independently: machine-learning collaboration, ensemble-based detection, and decentralized trust. Federated learning addresses data locality; ensemble learning addresses detection diversity; and blockchain-inspired identity and provenance mechanisms address participant accountability. Their integration creates a multilayer security architecture rather than a single detection algorithm.

The theoretical relationship with blockchain requires careful qualification. Nakamoto (2008) established a decentralized peer-to-peer model, while Mingxiao et al. (2017) demonstrate the importance of consensus mechanisms in blockchain systems. These principles can support trusted coordination, but blockchain should not be assumed to solve machine-learning problems automatically. A blockchain ledger cannot independently determine whether a model update is statistically useful or malicious. Therefore, machine-learning validation must remain part of the security architecture.

The same distinction applies to privacy. Federated learning reduces the need to transfer raw datasets, but it does not guarantee absolute confidentiality. This creates an important research direction concerning robust aggregation, update validation, and privacy-enhancing mechanisms. In an Industry 5.0 environment, the challenge is especially significant because compromised models can affect physical processes rather than merely digital records.

A further issue is non-IID data. Industrial factories do not generate identical traffic distributions. Consequently, a global model may perform well for one production environment while underperforming in another. The proposed ensemble

framework should therefore support local specialization alongside global knowledge sharing. A hierarchical architecture could, for example, create industry-specific or production-line-specific model groups before performing broader aggregation.

The framework also has practical implications for cross-organizational manufacturing ecosystems. Saberi et al. (2019) demonstrate the relevance of blockchain to distributed supply-chain relationships, while Guo and Liang (2016) illustrate decentralized trust in organizational information exchange. In a connected manufacturing ecosystem, cybersecurity intelligence can similarly become a shared resource. One organization may encounter a new attack pattern first, while other organizations benefit from the resulting model improvement without obtaining the original sensitive data.

The identity dimension is equally important. If participant identities cannot be established reliably, the collaborative learning process can be manipulated. Bouras et al. (2021) therefore provides a useful foundation for incorporating lightweight identity management into resource-constrained IoT environments. However, identity alone cannot establish behavioral trust; an authenticated device can still be compromised. Consequently, identity verification should be combined with update-quality monitoring and anomaly detection.

The proposed architecture also has limitations. Its conceptual nature means that the expected advantages have not yet been experimentally validated. No numerical improvement over centralized intrusion detection can therefore be claimed. Additional research must evaluate detection accuracy, false-positive rates, communication cost, convergence time, computational overhead, and resilience against adversarial updates.

Another limitation concerns governance. A federated system involving several factories requires agreement about participation rules, model ownership, update validation, liability, and incident response. Decentralized technologies may assist with record integrity, but organizational governance remains necessary. Ocheja et al. (2019) and Gordon and Catalini (2018) illustrate how distributed systems can support controlled information management, but industrial security requires domain-specific governance mechanisms.

Finally, the framework's success depends on the correct balance between local autonomy and global intelligence. Excessive centralization undermines privacy, while excessive decentralization can make coordination inefficient. The principal design challenge is therefore not merely to maximize detection capability but to establish an architecture in which privacy, accuracy, trust,

scalability, and operational efficiency are jointly optimized. This is consistent with the broader hybrid perspective associated with Industry 5.0 intrusion detection in Govindarajan et al. (2026). (Govindarajan et al., 2026)

CONCLUSION

This paper presented a conceptual Federated Ensemble Learning framework for privacy-preserving intrusion detection in Industry 5.0-enabled smart factories. The proposed architecture combines local industrial learning, heterogeneous ensemble detection, federated model aggregation, lightweight identity management, and blockchain-inspired trust and provenance mechanisms.

The central contribution is the architectural integration of these components around a common principle: industrial organizations should be able to collaborate on cybersecurity intelligence without unnecessarily exposing their raw operational data. Federated learning provides the collaborative mechanism, ensemble learning addresses detection diversity, and decentralized trust mechanisms provide accountability and participant governance.

The analysis demonstrates that the approach has substantial theoretical potential, particularly for manufacturing ecosystems characterized by distributed ownership, heterogeneous IIoT environments, and sensitive operational information. However, the framework should not be interpreted as experimentally validated. Communication overhead, non-IID industrial data, compromised participants, model poisoning, computational constraints, and governance complexity require systematic evaluation.

Future research should implement the framework using realistic Industry 5.0 traffic datasets and compare federated ensemble detection with centralized and isolated local-learning baselines. Evaluation should include detection effectiveness, false-positive behavior, communication overhead, convergence, resource utilization, privacy leakage, and robustness against malicious participants. Further work should also investigate adaptive aggregation and factory-specific model personalization. Such extensions could transform the conceptual framework into an experimentally validated security architecture for collaborative and privacy-preserving smart manufacturing.

REFERENCES

1. D. Mingxiao, M. Xiaofeng, Z. Zhe, W. Xiangwei, and C. Qijun, "A review on consensus algorithm of blockchain," in 2017 IEEE international conference on systems, man, and cybernetics (SMC). IEEE, 2017, pp. 2567–2572.
2. M. A. Bouras, Q. Lu, S. Dhelim, and H. Ning, "A lightweight blockchain based iot identity management approach," *Future Internet*, vol. 13, no. 2, p. 24, 2021

3. J. Gordon and C. Catalini, "Blockchain technology for healthcare: facilitating the transition to patient-driven interoperability," *Computational and structural biotechnology journal*, vol. 16, pp. 224–230, 2018.
4. Y. Guo and C. Liang, "Blockchain application and outlook in the banking industry," *Financial Innovation*, vol. 2, no. 1, pp. 1–12, 2016.
5. S. Huckle, R. Bhattacharya, M. White, and N. Beloff, "Internet of things, blockchain and shared economy applications," *Procedia computer science*, vol. 98, pp. 461–466, 2016.
6. S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," *Decentralized Business Review*, p. 21260, 2008.
7. P. Ocheja, B. Flanagan, H. Ueda, and H. Ogata, "Managing lifelong learning records through blockchain," *Research and Practice in Technology Enhanced Learning*, vol. 14, no. 1, pp. 1–19, 2019.
8. S. Saberi, M. Kouhizadeh, J. Sarkis, and L. Shen, "Blockchain technology and its relationships to sustainable supply chain management," *International Journal of Production Research*, vol. 57, no. 7, pp. 2117–2135, 2019.
9. D. R. Garrison, *E-learning in the 21st century: A community of inquiry framework for research and practice*. Routledge, 2016.
10. 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*.
11. K. Ramamurthy, R. K. Konduru and N. Amanmadov, "EvoGraphCoder: An Evolutionary Graph-Reasoning Framework for Self-Adaptive Software Engineering," in *IEEE Access*, vol. 14, pp. 63063–63076, 2026, doi: 10.1109/ACCESS.2026.3686019.
12. Geo Philip, Paulson, *Artificial Intelligence and Machine Learning Applications in Project Schedule Forecasting: A Predictive Framework for Time-Control in Complex Building Projects* (April 16, 2026). Available at SSRN: <https://ssrn.com/abstract=6588119> or <http://dx.doi.org/10.2139/ssrn.6588119>