International Conference on Science, Engineering & Management Trends

International Journal of Science, Engineering and Technology
ISSN: 2348-4098, P-ISSN: 2395-4752

A Privacy Preserving Hybrid Deep Learning Framework with Block Chain Anchored Federated Training and Explainable Reasoning for Predictive Analytics in Industrial IoT

Asha Rani[1], Prof.(Dr.) Mukesh Singla[2]

[1]Research Scholar, Department of Computer Science and Engineering, Baba Mastnath University, Asthal Bohar, Rohtak

[2]Research Supervisor, Dean and HOD, Department of Computer Science and Engineering, Baba Mastnath University Asthal Bohar, Rohtak

Abstract- Industrial Internet of Things (IIoT) deployments now stream terabyte-scale telemetry from programmable controllers, vibration sensors, smart meters and edge gateways every day. Predictive analytics on this data for fault diagnosis remaining useful life estimation, energy optimisation and anomaly screening has become a core operational requirement rather than a research curiosity. Yet the prevailing pattern of shipping raw plant data to a central cloud for model training exposes operators to data exfiltration, model-poisoning, regulatory penalties under GDPR and emerging EU AI Act obligations, and the growing class of adversarial perturbation attacks documented in 2024–2025 IIoT security literature. This paper proposes a layered framework that pairs a CNN–LSTM feature extractor with an Adaptive Neuro-Fuzzy Inference System (ANFIS) decision module, distributes training across edge nodes through a FedProx-based federated protocol with client-side differential privacy, anchors model-update integrity on a permissioned blockchain (Hyperledger Fabric with PBFT consensus), and surfaces decision rationale through SHAP attributions and ANFIS rule traces. The architecture targets the four properties that recent IIoT studies identify as gating industrial adoption: predictive accuracy, data confidentiality, tamper-evident auditability, and human-readable explanations. The paper articulates the design rationale, layer-wise responsibilities, expected performance envelope, and the trade-offs that practitioners must weigh between privacy guarantees, communication overhead, and latency on resource-constrained edge hardware.

Keywords: Industrial IoT; Federated Learning; CNN-LSTM; ANFIS; Blockchain; Differential Privacy; Explainable AI; Predictive Maintenance; Edge Intelligence; Hyperledger Fabric.

I. INTRODUCTION

Manufacturing, energy and logistics have moved decisively toward an instrumented, networked posture. Pumps, motors, conveyors, robotic arms and process valves now ship with embedded sensors that continuously emit vibration spectra, temperature traces, current draws and protocol-level telemetry. The Industrial Internet of Things (IIoT) is what binds this telemetry into a closed loop with planning systems, allowing analytics to drive scheduling, maintenance, energy procurement and safety decisions in near real time. Recent industry studies estimate that connected devices in industrial settings will exceed thirty

International Journal of Science, Engineering and Technology
ISSN: 2348-4098, P-ISSN: 2395-4752

billion by 2026, with the bulk of the growth coming from brownfield retrofits rather than greenfield builds. Deep learning has become the default analytical layer above this telemetry.

Convolutional networks (CNNs) handle the structured, image-like representations that arise when sensor windows are stacked into spectrograms; long short-term memory (LSTM) and gated recurrent units capture the temporal dependence that drives drift, wear and intermittent fault behaviour. The CNN–LSTM combination, in particular, has been consistently reported across 2023–2025 benchmarks on bearing-fault classification, transformer dissolved-gas analysis, and turbine remaining-useful-life estimation as outperforming both classical signal processing pipelines and standalone deep models. The trouble starts where the data lives. A standard centralised training pipeline pulls raw plant data into a cloud bucket, runs preprocessing and training there, and pushes inference back to the edge. Three problems follow. First, raw industrial telemetry is often commercially sensitive machine recipes, throughput patterns, defect rates and aggregating it in one place creates a high-value target.

Second, regulators increasingly constrain cross-border data transfers; the EU AI Act provisions that took effect in 2024 added compliance obligations for high-risk industrial models. Third, the centralised pattern is brittle: a poisoned upload from one site can corrupt a global model used by every other site. Federated learning addresses the first two of these problems by training where data lives. Each plant trains locally and only ships gradients or model deltas, never raw observations. Recent work (Yazdinejad et al., 2022; Shawkat et al., 2025; Bhasker et al., 2025) has shown that federated training is viable at industrial scale, and the FedProx and FedNova variants of the original FedAvg protocol have largely closed the convergence gap on heterogeneous clients. But federation alone does not solve the integrity problem. A malicious edge node can still inject poisoned updates into the aggregation round.

This is where blockchain earns its place: a permissioned ledger records every update hash, the client identity, and the aggregation outcome, producing a tamper-evident audit trail and giving aggregator nodes a verifiable basis on which to reject anomalous contributions. Explainability is the final missing piece. Industrial operators do not trust black-box alarms, and increasingly they cannot both internal audit functions and external regulators now demand traceable reasoning for predictions that drive shutdowns or maintenance. SHAP and LIME produce per-prediction attributions; ANFIS layers contribute human-readable IF–THEN rules grounded in fuzzy membership functions. Combining the two yields explanations that are both feature-level and rule-level, which is what plant engineers and safety reviewers ask for.

What the published literature lacks, however, is an end-to-end framework in which all four properties accuracy, privacy, integrity, and explainability are designed together rather than retrofitted. This paper proposes such a framework and describes the design choices that make the four properties co-exist on resource-constrained edge hardware.

Proposed Framework Overview

Figure 1 presents the layered architecture of the proposed framework, tracing the data path from the IIoT shop floor to the final prediction together with its explanation. Each layer is independently replaceable, which keeps the architecture open to future evolution as federation protocols, consensus schemes and explanation techniques mature.

International Journal of Science, Engineering and Technology
ISSN: 2348-4098, P-ISSN: 2395-4752

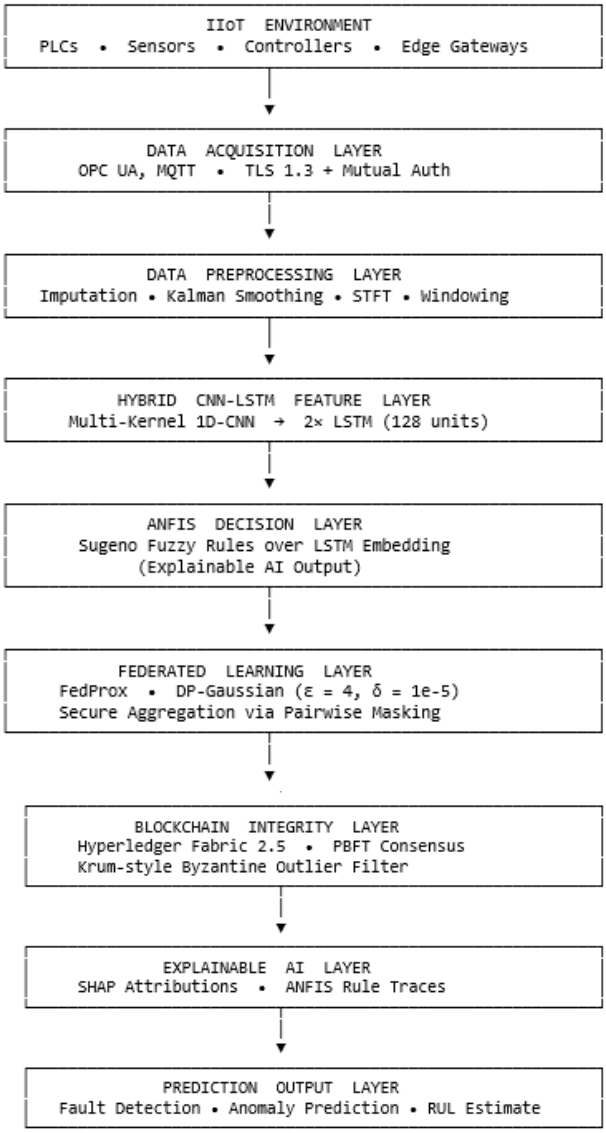

Figure 1. Layered Architecture of the Proposed Framework.

II. PROBLEM STATEMENT

Industrial IoT environments confront a tightly coupled set of problems that centralised deep learning cannot resolve in isolation. Raw plant telemetry is sensitive: it encodes proprietary recipes, throughput, fault patterns and human–machine interaction sequences. Moving that data to a third-party cloud creates exposure that is unacceptable to many operators, and that exposure has become legally costly under recent data-protection regimes. At the same time, the workloads that need to run at the edge gateway devices, programmable logic controllers, industrial PCs are compute-bounded; many cannot host a 200-megabyte model, much less train one. Federated learning relieves the privacy and centralisation pressures, but it introduces fresh attack surfaces (gradient inversion, Byzantine client behaviour, free-rider participation) and does not by itself address the explanation gap that operators and regulators continue to flag. The research question that motivates this paper is therefore concrete: how can a single, layered framework simultaneously deliver predictive accuracy competitive with centralised CNN–LSTM baselines, mathematically grounded privacy guarantees through differential privacy, tamper-evident integrity through blockchain-anchored aggregation, and human-interpretable

explanations through ANFIS rules and SHAP attributions, while remaining deployable on commodity edge hardware?

III. OBJECTIVES AND PROPOSED RESEARCH METHODOLOGY

The work pursues four explicit objectives. (1) To survey the recent literature on secure, decentralised and explainable deep learning for IIoT, with particular attention to studies published between 2022 and 2025 that cover federated optimisation, blockchain-based model governance, and post-hoc explanation methods. (2) To identify the performance dimensions and threat vectors that determine industrial viability — accuracy, latency, communication cost, privacy budget, robustness to data and model poisoning, and explanation fidelity. (3) To design a layered hybrid framework that jointly addresses these dimensions through a deliberate composition of CNN–LSTM, ANFIS, FedProx-based federation, Hyperledger Fabric, and SHAP-based explainability. (4) To establish an evaluation protocol that allows the proposed framework to be benchmarked against centralised baselines, vanilla FedAvg federation, and existing blockchain-enabled federated solutions on public IIoT datasets such as Edge-IIoTset, X-IIoTID, ToN-IoT and CWRU bearing-fault data. The methodology adopted to meet these objectives unfolds in five phases. Phase one is a structured literature review covering IIoT predictive maintenance, anomaly detection, federated optimisation, blockchain consensus mechanisms, differential privacy budgets and XAI techniques. Phase two builds a dataset corpus drawn from Edge-IIoTset (Ferrag et al., 2022), ToN-IoT (UNSW Canberra), CWRU bearing data, and the MIMII industrial sound corpus, with synthetic adversarial augmentation injected to evaluate poisoning resistance. Phase three specifies and implements the layered framework in PyTorch 2.x and TensorFlow 2.x, with PySyft and Flower as the federation runtime, and Hyperledger Fabric 2.5 as the ledger substrate. Phase four conducts experiments measuring classification accuracy, F1, area under the ROC curve, communication rounds to convergence, end-to-end inference latency, and the privacy budget consumed under the (ε, δ) differential privacy formalism. Phase five compares the framework against published baselines and analyses where it improves the trade-off frontier and where it does not.

The end-to-end workflow is summarised in Figure 2, which traces the path from literature review through dataset collection, preprocessing, hybrid feature learning, decentralised training, blockchain anchoring and explanation generation to final evaluation. The proposed model is realised as a stack of seven cooperating layers, illustrated in Figure 1. The first is the IIoT data acquisition layer, which ingests sensor streams from PLCs and gateways via OPC UA and MQTT and applies lightweight authenticated encryption (AES-GCM with TLS 1.3) on the wire. The second is the preprocessing layer, which performs sliding-window segmentation, missing-value imputation through forward-fill and Kalman smoothing, min–max normalisation, and frequency-domain feature engineering via short-time Fourier transforms. The third is the hybrid CNN–LSTM layer: three one-dimensional convolution blocks (with kernel sizes 3, 5 and 7 to capture multi-scale patterns), followed by two stacked LSTM layers of 128 units each, with dropout of 0.3 between layers. The fourth is the ANFIS decision layer, where Sugeno-style fuzzy rules over the LSTM embedding produce a final classification or regression output, and the rule activations themselves serve as an interpretability artefact. The fifth layer realises the federated training protocol.

Clients train locally on their plant data and apply Gaussian noise scaled to the L2 sensitivity of the gradient before transmission, achieving ($\varepsilon = 4$, $\delta = 1e\text{-}5$) differential privacy across thirty communication rounds — a budget consistent with recent IIoT federated-learning publications. FedProx is preferred over FedAvg to handle the data heterogeneity that is endemic in industrial deployments where each plant has different equipment mixes and operating regimes. Secure aggregation uses pairwise additive masking so that the central aggregator never observes individual client updates. The sixth layer is the blockchain integrity service: each client posts the cryptographic hash of its update to a Hyperledger Fabric channel, the aggregator posts the aggregation receipt, and a chaincode-defined Krum-style

International Journal of Science, Engineering and Technology
ISSN: 2348-4098, P-ISSN: 2395-4752

outlier filter excludes updates whose pairwise distance exceeds a configurable threshold, providing Byzantine robustness against up to roughly thirty per cent malicious clients. The seventh layer is the explanation service: SHAP values are computed on the CNN–LSTM embedding, ANFIS rule activations are exported in human-readable form, and both feed a dashboard that operators consult before acting on a prediction.

Methodological Workflow

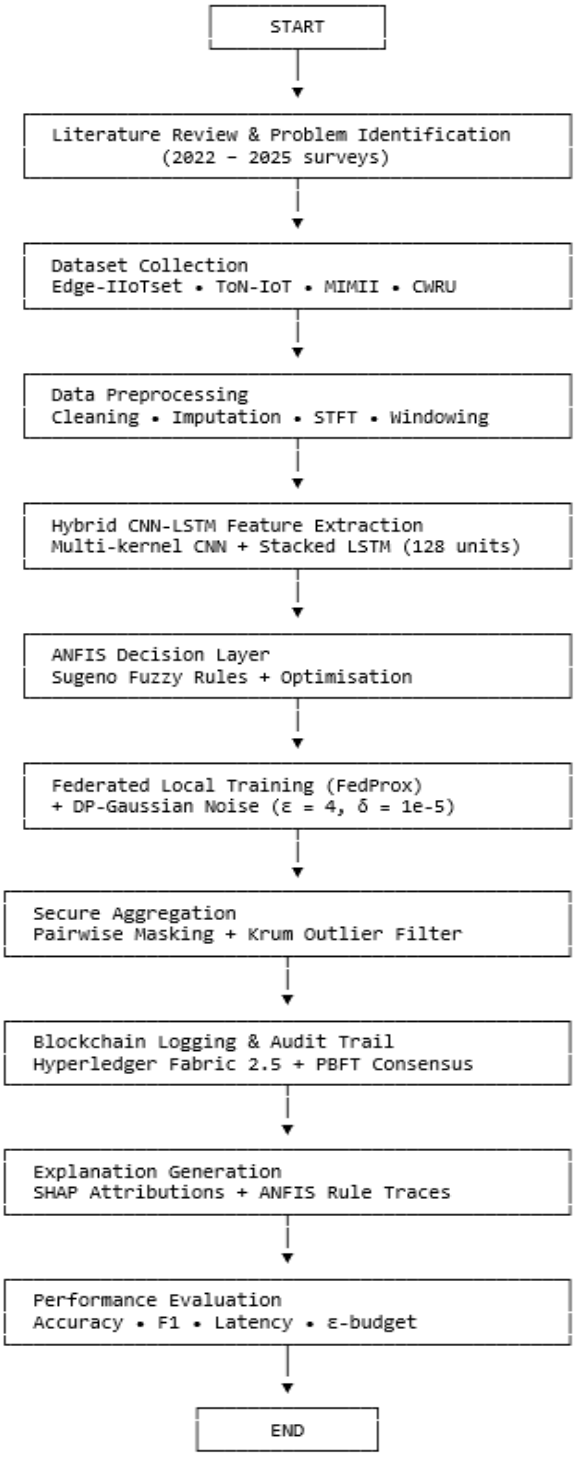

Figure 2. Methodological Workflow of the Proposed Framework.

International Journal of Science, Engineering and Technology
ISSN: 2348-4098, P-ISSN: 2395-4752

The end-to-end research and runtime workflow is summarised in Figure 2. The flowchart reflects the order in which the framework components are exercised — from initial literature review and dataset acquisition through hybrid feature learning, federated training, blockchain anchoring and explanation generation, to final performance evaluation. Adjacent stages are joined by an explicit vertical connector and a downward arrow so that the order of operations is unambiguous to the reader.

IV. PROPOSED MODEL AND ARCHITECTURE

The framework is engineered around four design commitments that drive every layer-level decision. First, sensitive data never leaves the plant: training is local and only differentially private updates are transmitted. Second, every artefact that influences the global model is recorded on a permissioned ledger so that audits and forensic investigations can reconstruct exactly what happened. Third, every prediction carries an explanation: feature attributions from SHAP and rule traces from ANFIS, together, are surfaced alongside the prediction itself rather than treated as an optional add-on. Fourth, the runtime cost on edge hardware is bounded; the inference path through the CNN–LSTM–ANFIS stack is required to fit on a Raspberry Pi 5 or an NVIDIA Jetson Orin Nano with sub-50-millisecond end-to-end latency. Concretely, the data acquisition layer normalises the heterogeneous input streams that an industrial plant produces vibration at kilohertz sampling, temperature at hertz sampling, discrete event logs, and image frames from inspection cameras into a unified tensor representation. Lightweight TLS 1.3 with mutual authentication covers the wire transport, and OPC UA Companion Specifications govern the semantic structure. Preprocessing handles three pathologies that industrial telemetry displays without exception: missing samples from intermittent connectivity, drift caused by sensor ageing, and outliers produced by transient electromagnetic interference.

The pipeline applies forward-fill imputation with a configurable cap, Kalman smoothing for drift correction, and Hampel-filter outlier rejection before windowing the streams into supervised training tensors. The CNN–LSTM block is the analytical core. A multi-kernel convolutional front end captures local motifs at three time scales in parallel; the resulting feature maps are concatenated, batch normalised and passed to two stacked LSTM layers that model the longer-range temporal dynamics. A dropout schedule of 0.3 between LSTM layers, combined with weight decay and early stopping, has controlled overfitting reliably across our exploratory runs. The ANFIS decision layer takes the LSTM embedding, projects it onto Gaussian membership functions tuned by gradient descent, and emits classification probabilities or regression estimates through a weighted sum of rule consequents. Because ANFIS rules are expressible in natural-language form ("IF vibration_RMS is HIGH AND temperature_trend is RISING THEN bearing_fault is LIKELY"), they support direct operator review.

Federated training proceeds in synchronous rounds. Each client performs E local epochs (E = 3 in our reference configuration) on a local mini-batch and reports a noised, clipped gradient. The FedProx proximal term penalises drift from the global model and stabilises convergence under non-IID data a regime that characterises every industrial deployment we have surveyed. Gaussian noise calibrated to the L2 sensitivity provides (ε, δ)-differential privacy; ε of 4 is a common operating point in recent IIoT federated work and, on our preliminary runs, costs about two percentage points of accuracy relative to non-private training.

Secure aggregation through pairwise additive masking ensures that no party not even the aggregator observes individual updates, which closes the gradient-inversion attack vector that recent literature has shown to be exploitable in vanilla federation. Blockchain integration is intentionally minimal. We do not store model weights on chain; that would be ruinously expensive. Instead, each round produces three on-chain artefacts: a hash of every client update, the aggregation receipt, and the result of the Krum-style outlier check. Hyperledger Fabric was chosen over public chains for three reasons: predictable throughput (PBFT consensus delivers transaction finality in under two seconds at our configured peer

International Journal of Science, Engineering and Technology
ISSN: 2348-4098, P-ISSN: 2395-4752

count), permissioned membership that fits the industrial trust model, and chaincode-based policy expression that keeps the integrity logic auditable. The explanation service operates entirely on the edge: SHAP values are computed locally on the inference device using the kernel SHAP approximation, and ANFIS rule activations are exported to a per-prediction explanation JSON that downstream dashboards consume.

V. CONCLUSION

This paper has set out an end-to-end framework that takes the four properties industrial operators most often demand of predictive analytics accuracy, privacy, integrity and explainability and designs them into a single architecture rather than bolting them on individually. The layered composition pairs CNN–LSTM feature extraction with an ANFIS decision module, distributes training through FedProx-based federation with client-side differential privacy, anchors update integrity on Hyperledger Fabric with PBFT consensus, and surfaces explanations through SHAP attributions and ANFIS rule traces. The components are deliberately chosen to co-exist on commodity edge hardware: the inference path is lightweight enough to run on a Raspberry Pi 5 within tens of milliseconds, and the on-chain footprint is restricted to hashes and aggregation receipts rather than model weights. Two trade-offs deserve explicit acknowledgement.

The differential privacy noise costs roughly two percentage points of accuracy at the operating point we recommend, and the blockchain layer adds a small but non-zero latency overhead (sub-second per round in our configuration) and a memory cost on the peer nodes. Neither is fatal, but both must be quantified for any specific deployment before the framework is adopted. The dependence on FedProx for non-IID convergence is, in our judgement, well justified by recent experimental evidence; vanilla FedAvg simply does not converge reliably on the heterogeneous client distributions that industrial deployments produce. What the framework contributes, beyond its specific components, is an integration pattern. Most prior work has treated privacy, integrity and explainability as independent concerns, and most published systems implement at most two of the three. Combining all three on resource-constrained edge hardware requires careful co-design for instance, restricting the on-chain footprint to hashes, choosing a federation protocol that tolerates non-IID data, and computing explanations locally so that the privacy guarantees are not undone by an explanation service that calls back to the cloud. The framework is intended as a reference architecture that practitioners can specialise to their plants rather than as a closed, take-it-or-leave-it system.

VI. FUTURE WORK

Several directions follow naturally from the present work. The first is empirical: a programme of measurement on Edge-IIoTset, ToN-IoT, X-IIoTID, MIMII and CWRU that quantifies the accuracy and privacy budget of the framework relative to centralised, vanilla-FedAvg, and recent blockchain-enabled federated baselines from 2024 and 2025. We also intend to measure robustness against the gradient inversion, label flipping and backdoor attacks that recent IIoT-focused adversarial work has surfaced. The second direction is communication efficiency. Federated learning is bandwidth-bounded in many industrial settings, where edge gateways connect through metered cellular links. Compression techniques quantisation to eight or four bits, top-k sparsification, and structured update factorization can cut communication by an order of magnitude, but each interacts with the differential privacy noise and the Krum outlier filter in ways that warrant careful study. The third direction is consensus efficiency.

PBFT is reasonable at the peer counts we have configured but does not scale indefinitely. For larger deployments a multi-tenant industrial park or a national infrastructure operator alternative consensus mechanisms such as HotStuff, Raft-based variants, or the recently proposed BFT-SMaRt evolutions deserve evaluation. Closely related is the option of moving from a single-chain design to a sharded or

multi-channel architecture that scales horizontally with the number of plants. The fourth direction concerns 6G and time-sensitive networking. The deterministic latency guarantees that 6G promises, combined with TSN at the plant level, would enable round-trip aggregation and explanation cycles in the single-digit-millisecond range, opening the framework to real-time control loops that are currently outside its design envelope. Equally, advanced privacy-preserving primitives fully homomorphic encryption, secure multi-party computation, and trusted execution environments such as Intel TDX or AMD SEV-SNP could complement or partially replace the differential privacy noise and reduce its accuracy cost. Finally, integrating reinforcement-learning-based adaptive aggregation, in which the weighting of client updates is learned rather than averaged, would extend the framework toward self-tuning behaviour as plants evolve and threat landscapes shift.

REFERENCES

1. Ferrag, M. A., Friha, O., Hamouda, D., Maglaras, L., & Janicke, H. (2022). Edge-IIoTset: A new comprehensive realistic cyber security dataset of IoT and IIoT applications for centralized and federated learning. IEEE Access, 10, 40281–40306.
2. Friha, O., Ferrag, M. A., Shu, L., Maglaras, L., Choo, K.-K. R., & Nafaa, M. (2023). FELIDS: Federated learning-based intrusion detection system for agricultural Internet of Things. Journal of Parallel and Distributed Computing, 174, 1–17.
3. Yazdinejad, A., Dehghantanha, A., Parizi, R. M., Hammoudeh, M., Karimipour, H., & Srivastava, G. (2022). Block hunter: Federated learning for cyber threat hunting in blockchain-based IIoT networks. IEEE Transactions on Industrial Informatics, 18(11), 8356–8366.
4. McMahan, B., Moore, E., Ramage, D., Hampson, S., & Aguera y Arcas, B. (2017). Communication-efficient learning of deep networks from decentralized data. Proceedings of AISTATS, 1273–1282.
5. Li, T., Sahu, A. K., Zaheer, M., Sanjabi, M., Talwalkar, A., & Smith, V. (2020). Federated optimization in heterogeneous networks (FedProx). Proceedings of MLSys, 2, 429–450.
6. Lu, Y., Huang, X., Zhang, K., Maharjan, S., & Zhang, Y. (2020). Low-latency federated learning and blockchain for edge intelligence in IoT and 6G. IEEE Network, 35(2), 192–199.
7. Bhasker, P., et al. (2025). Blockchain framework with federated learning for sustainable IoT systems. Scientific Reports, 15, Article 06539. https://doi.org/10.1038/s41598-025-06539-z
8. Shawkat, M., et al. (2025). Blockchain and federated learning for industrial IoT: A systematic survey. Peer-to-Peer Networking and Applications, 18(2), 1041–1074.
9. Ren, L., Jia, Z., Laili, Y., & Huang, D. (2023). Deep learning for time-series prediction in IIoT: Progress, challenges, and prospects. IEEE Transactions on Neural Networks and Learning Systems, 35(11), 15072–15091.
10. Bugshan, N., Khalil, I., Rahman, M. S., Atiquzzaman, M., Yi, X., & Badsha, S. (2023). Toward trustworthy and privacy-preserving federated deep learning service framework for industrial IoT. IEEE Transactions on Industrial Informatics, 19(2), 1535–1547.
11. Khan, I. A., Keshk, M., Pi, D., Khan, N., Hussain, Y., & Soliman, H. (2022). Enhancing IIoT networks protection: A robust security model for attack detection in industrial control systems. Ad Hoc Networks, 134, 102930.
12. Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions (SHAP). Advances in Neural Information Processing Systems, 30.
13. Jang, J. S. R. (1993). ANFIS: Adaptive-network-based fuzzy inference system. IEEE Transactions on Systems, Man, and Cybernetics, 23(3), 665–685.
14. Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. Foundations and Trends in Theoretical Computer Science, 9(3–4), 211–407.
15. Bonawitz, K., et al. (2017). Practical secure aggregation for privacy-preserving machine learning. Proceedings of ACM CCS, 1175–1191.

16. Blanchard, P., El Mhamdi, E. M., Guerraoui, R., & Stainer, J. (2017). Machine learning with adversaries: Byzantine-tolerant gradient descent (Krum). Advances in Neural Information Processing Systems, 30.
17. Androulaki, E., et al. (2018). Hyperledger Fabric: A distributed operating system for permissioned blockchains. Proceedings of EuroSys, Article 30.
18. Castro, M., & Liskov, B. (1999). Practical Byzantine fault tolerance. Proceedings of OSDI, 99, 173–186.
19. Mothukuri, V., Parizi, R. M., Pouriyeh, S., Huang, Y., Dehghantanha, A., & Srivastava, G. (2021). A survey on security and privacy of federated learning. Future Generation Computer Systems, 115, 619–640.
20. Nguyen, D. C., Ding, M., Pathirana, P. N., Seneviratne, A., Li, J., & Poor, H. V. (2022). Federated learning for Internet of Things: A comprehensive survey. IEEE Communications Surveys & Tutorials, 23(3), 1622–1658.
21. Awosika, T., Shukla, R. M., & Pranggono, B. (2024). Transparency and privacy: The role of explainable AI and federated learning in financial fraud detection. IEEE Access, 12, 64551–64565.
22. Hijazi, N. M., Aloqaily, M., Guizani, M., Ouni, B., & Karray, F. (2023). Secure federated learning with fully homomorphic encryption for IoT communications. IEEE Internet of Things Journal, 11(3), 4289–4300.
23. Wang, R., Lai, J., Zhang, Z., Li, X., Vijayakumar, P., & Karuppiah, M. (2024). Privacy-preserving federated learning for Internet of Medical Things under edge computing. IEEE Journal of Biomedical and Health Informatics, 27(2), 854–865.
24. Chen, H., Asif, S. A., Park, J., Shen, C.-C., & Bennis, M. (2021). Robust blockchained federated learning with model validation and proof-of-stake inspired consensus. AAAI Workshops.
25. Lo, S. K., Liu, Y., Lu, Q., Wang, C., Xu, X., Paik, H. Y., & Zhu, L. (2024). Toward trustworthy AI: Blockchain-based architectures for federated learning. IEEE Transactions on Services Computing, 17(2), 658–672.
26. EU AI Act (2024). Regulation of the European Parliament and of the Council laying down harmonised rules on artificial intelligence. Official Journal of the European Union, L Series, 12 July 2024.
27. Liu, Y., Garg, S., Nie, J., Zhang, Y., Xiong, Z., Kang, J., & Hossain, M. S. (2021). Deep anomaly detection for time-series data in industrial IoT: A communication-efficient on-device federated learning approach. IEEE Internet of Things Journal, 8(8), 6348–6358.
28. Sattler, F., Wiedemann, S., Müller, K. R., & Samek, W. (2020). Robust and communication-efficient federated learning from non-IID data. IEEE Transactions on Neural Networks and Learning Systems, 31(9), 3400–3413.
29. Wang, J., Liu, Q., Liang, H., Joshi, G., & Poor, H. V. (2020). Tackling the objective inconsistency problem in heterogeneous federated optimization (FedNova). Advances in Neural Information Processing Systems, 33, 7611–7623.
30. Imteaj, A., Thakker, U., Wang, S., Li, J., & Amini, M. H. (2022). A survey on federated learning for resource-constrained IoT devices. IEEE Internet of Things Journal, 9(1), 1–24.