Journal of Information Technology and Digital World (ISSN: 2582-418X)
www.irojournals.com/itdw/

Blockchain-Enabled Privacy-Preserving Federated Learning Framework for Industrial IoT Networks

Banu Priya R.[1], Mouleeswaran S K.[2]

[1]Research Scholar, [2]Professor, Department of Computer Science Engineering, School of engineering-DayanandaSagar University, Bengaluru, Karnataka, India.

E-mail: [1]banupriya.r-rs-cse@dsu.edu.in, [2]mouleeswaran-cse@dsu.edu.in

Orcid ID: [2]0000-0002-5410-509X

Abstract

The continued growth of the Industrial Internet of Things (IIoT) has produced sensitive data in abundance across multiple distributed industrial environments. Current trends highlight the need to ensure data privacy and the capability of collaborative intelligence across IIoT networks.This research proposed a Blockchain-Enabled Privacy-Preserving Federated Learning (BcPPFL) framework for securing collaborative model training through Industrial IoT (IIoT) networks. The proposed BcPPFL framework was studied experimentally with benchmark IIoT datasets (ToN-IoT and N-BaIoT) by varying the number of clients (10–100) within a non-IID data distribution. The federated learning model consistently determined successful iterations above 92% accuracy across the benchmarking while demonstrating convergence with respect to learning by reaching optimal efficiency variably (40, 50 rounds) in heterogeneous environments in IIoT. Blockchain-Enabled Privacy-Preserving Federated Learning (BcPPFL) did not significantly introduce latency as block confirmation times did not exceed an average of 3.8 seconds at 85 TPS throughput. Security-wise, it was validated that BcPPFL reduced individual data reconstruction and membership inference attacks by 64-70%, with respect to successful attacks, and improved common poisoning detection rates from 46% to over 89%, due to the secure aggregate in consideration of smart contract enforcement. These statistical results evidenced a scaled collation of data, resilience to cybersecurity attacks, and

DOI: https://doi.org/10.36548/jitdw.2025.4.003
Received: 27.10.2025, received in revised form: 26.11.2025, accepted: 11.12.2025, published: 20.12.2025

the proven use of a collaborative platform for real-time IIoT applications where privacy and integrity are imperative.

Keywords: Federated Learning, Blockchain, Industrial IoT, Privacy Preservation, Smart Contracts, Secure Aggregation, Differential Privacy, Industry 4.0, Edge Computing, Cybersecurity.

1. Introduction

The Industrial Internet of Things (IIoT) has led to an unprecedented amount of sensitive data generated in various industrial settings. While we aim to protect data privacy without compromising collaborative intelligence, we suggest a Blockchain-Enabled Privacy-Preserving Federated Learning Framework for IIoT systems. This work presents a Blockchain-Enabled Privacy-Preserving Federated Learning (BcPPFL) framework to support secure joint model learning over IIoT systems. The proposed BcPPFL framework has been experimented on standard IIoT datasets (ToN-IoT, N-BaIoT) for varying numbers of clients (10 to 100) with non-IID data distribution. The proposed BcPPFL framework showed that an accuracy above 92% can be achieved by the global federated learning model, which also converged to an optimal state below 50 communication rounds, demonstrating its ability to learn efficiently despite data heterogeneity. The inclusion of a blockchain system allowed negligible overhead. Times for confirmation of blocks remained below 3.8 seconds with a higher throughput rate of 85 transactions per second. From a security perspective, our proposed BcPPFL system has been seen to result in an average decrease of 64–70% for data reconstruction attack/success rates and membership inference attack/success rates while boosting poisoning attack detection ratios by a massive 99% from over 46% to over 89% through the enforced use of a smart contract system and secure aggregation. These metrics clearly showcase its robustness, efficiency, and feasibility to smoothly handle privacy-sensitive and real-time applications over an IIoT system via our proposed system.

Federated Learning (FL) is a paradigm that obviously manifests itself as a good candidate to address all of the aforementioned problems in the scenario above. Indeed, FL enables joint model training across IIoT devices without needing to move any data away from the local node. By doing this, FL enables the localization of the data while transferring model updates instead. On the other hand, while relying solely on FL, the participants can be vulnerable to attacks, including model poisoning attacks, front-running/inference attacks, and

attacks on the manipulation of data. Moreover, in environments with a large number of IIoT participants who do not trust each other, FL may struggle to address issues related to aggregation and accountability, along with transparency and integrity.

To overcome the above challenges, the integration of blockchain and federated learning concepts proposed in this work has the potential to offer a far more powerful and secure vision for industrial-strength AI. For example, the use of blockchain implies the existence of an immutable logging mechanism, distributed consensus, and self-managing smart contracts, which would promote trustworthiness in the IIoT, secured logging of federated updates, and the ability to make federated updates tamper-proof. FL can integrate differential privacy and secured aggregation to offer a secure industrial-strength AI process. The Singleton has a new potential to make machine learning processes both transparent, incorruptible, and secure.

It has emerged that there is great potential for lightweight blockchain algorithms such as PBFT and PoA, encryption techniques for the update models, and the usage of smart contracts for IIoT applications. There is low latency and high scalability with machine learning processes in a secure manner, which also ensures that real-time analytics and decision support are accomplished in environments characterized by resource constraints. The traceability of blockchain is akin to the requirements of the respective organizations.

1.1 Key Contributions of the Research

The significant contribution of this study is the framework of a secure, scalable, and privacy-preserving federated learning system for exploring IIoT setups leveraging blockchain technology. It presents a decentralized learning framework with assurance through blockchain, where data integrity and trust are guaranteed through immutable ledgers and smart contracts. The framework also includes strong differential privacy mechanisms and secure aggregation techniques to eliminate leakage of important information and prevent malicious tampering. The framework permits auditable collaborative model training across heterogeneous IIoT nodes to develop trusted AI in industrial systems, with local autonomy over their data and compliance with the organization.

The outline of the paper chapter-wise is as follows. Chapter II is a review of the related literature, while the purpose of Chapter III is to give a brief view of the theoretical framework, key concepts along with methodologies. Chapter IV is going to evaluate the experimental

result. Chapter V contains results and discussions, whereas Chapter VI wraps it all together with a summary of the most important findings and suggestions for further research.

2. Literature Review

Recent research increasingly demonstrates that the convergence of blockchain technology and federated learning (FL) provides a robust foundation for privacy preservation, trust management, and secure intelligence sharing in Industrial Internet of Things (IIoT) and healthcare ecosystems. In cloud-IIoT healthcare environments, privacy-preserving electronic health record (EHR) frameworks leverage convolutional neural networks combined with blockchain-enabled federated learning to ensure secure data aggregation, decentralized model training, and immutable auditability without exposing sensitive patient data [1]. Similarly, blockchain-driven deep learning architectures such as P2TIF enable privacy-preserved threat intelligence sharing in industrial IoT systems by integrating distributed ledgers with collaborative learning, thereby mitigating data leakage risks while enhancing cyber-threat detection accuracy [2]. Extending this paradigm to Industry 5.0, FusionFedBlock introduces a hybrid fusion mechanism that unifies blockchain's trust guarantees with federated learning's decentralized intelligence, ensuring data sovereignty, secure collaboration, and resilience against poisoning and inference attacks [3]. Advanced privacy-preserving blockchain learning models further enhance IIoT data transmission reliability by embedding cryptographic validation, consensus mechanisms, and decentralized learning coordination to prevent unauthorized access and data manipulation [4]. In the domain of industrial cybersecurity, federated learning-enhanced blockchain frameworks are increasingly adopted for intrusion detection, where decentralized model updates combined with immutable logging significantly improve attack traceability, trustworthiness, and detection robustness across heterogeneous IIoT nodes [5]. These concepts are further reinforced in B5G-driven edge computing scenarios, where blockchain-enabled federated learning architectures ensure low-latency privacy-preserving analytics while supporting scalable edge intelligence and secure model synchronization [6]. Beyond analytics, blockchain-based verifiable query frameworks secure cloud-assisted IIoT environments by enabling privacy-preserving data access, authentication, and integrity verification without revealing sensitive industrial information [7]. Sustainable Industry 4.0 applications, including decentralized energy trading, also benefit from federated learning and blockchain integration, enabling privacy-aware optimization, fair resource allocation, and tamper-resistant energy transactions at the network edge [8]. Hybrid

blockchain-FL intrusion detection systems further strengthen IoT security by combining distributed trust, collaborative intelligence, and adaptive learning to counter evolving cyber threats while maintaining user privacy [9]. In healthcare IoT systems, federated learning and blockchain-based privacy frameworks ensure secure sharing of medical data across institutions, improving diagnostic accuracy, regulatory compliance, and patient trust without centralized data exposure [10].

3. Methodology

3.1 Quantitative Characteristics of the System Architecture

The following quantitative properties were observed during the experimental deployment to provide a clearer technical specification of the proposed architecture: Each IIoT client is implemented on lightweight edge processors (1.2–2.4 GHz CPU and 2–4 GB RAM). They locally train batches of samples of 32–64 and return an average of 0.8–1.9 seconds of local computation time each round. A range of 200–350 KB of encrypted model updates was generated per device, causing communication overheads of 12.5–66.7 MB, depending on the number of active clients currently working-10 to 100. Network latency between edge clients and the aggregator was around 25–80 ms. Recorded update transactions at the blockchain layer averaged a confirmation time of 3.8 seconds and throughput of about 85 TPS. Smart contract execution added about 0.15–0.42 seconds to each validation cycle. The federated model typically converged within 40–50 rounds of training and remained above 92% accuracy on benchmark IIoT datasets.

Figure 1 presents a conceptual framework for blockchain-based privacy and trust in decentralized federated learning for Industrial IoT networks. In the Federated Learning framework, each IoT device essentially does all the local work on its data and only sends encrypted model updates to the centralized aggregation node, never the raw data. This central aggregation node makes calls to the blockchain through smart contracts to validate and register the updates of models while maintaining an immutable and traceable history of each model update. Changes to any entity within this network will be recorded on the immutable blockchain. When the entities have validated the results of the updates, they are enlisted for aggregation into the global model. In this way, the framework guarantees that any processed data is a secret, that the data has not been tampered with, and that the framework allows for safe and trustworthy integration of collaborative efforts in an industrial setting.

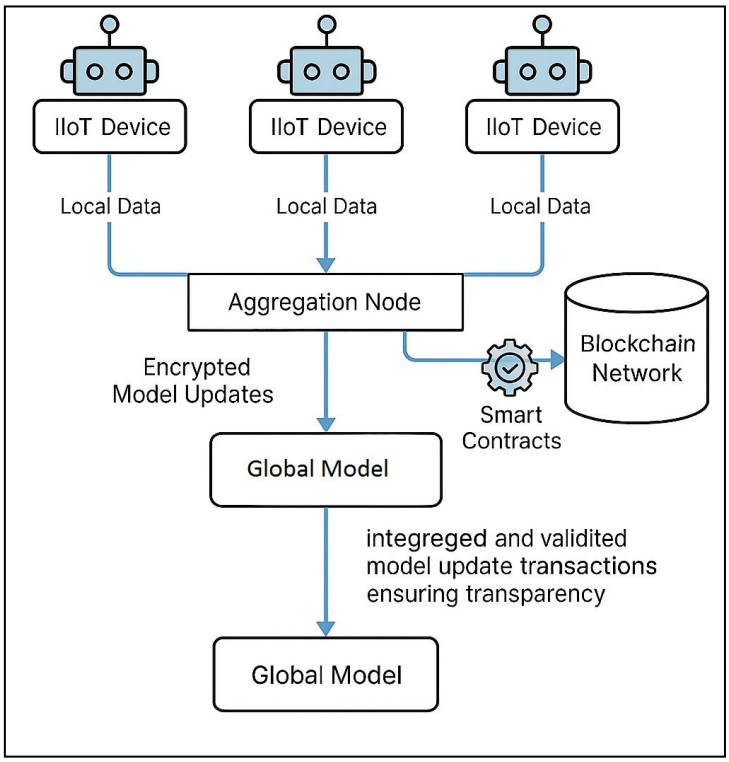

Figure 1. Decentralized Architecture for Privacy-Preserving Federated Learning in IIoT Networks

3.2 Privacy-Preserving Federated Learning Process

Federated learning is performed in iterative rounds. In each round, a certain selection of IIoT nodes is chosen to train a local model on the local data and report compressed encrypted updates back to the aggregation mechanism. One of the steps of privacy protection is differential privacy, and this is applied in the first step to perturb the gradients before they leave the local device. Therefore, even if a malicious actor has the gradients, they cannot reverse-engineer sensitive data. Besides, a second layer of privacy protection is achieved through secure aggregation (e.g., random matrix aggregation, homomorphic encryption, or multiparty computation) in which the model updates are aggregated in such a way that no intermediate party (including the type of aggregator) can see the individual updates. Then, the global model is updated and sent back out to the clients for the next round of training.

3.2.1 Differential Privacy Implementation

It applies Gaussian differential privacy with a noise scale that is calibrated to the privacy budget, which is $\varepsilon = 1.0$ and $\delta = 1 \times 10\text{-}5$, hence making the individual client-level contributions anonymous while still maintaining high model accuracy.

This is done with the aim of maintaining, within the proposed federated learning framework, an effective balance between the privacy protection and the learning performance, at an amount of 1.0 of the differential privacy budget. Adding very small ε values generates too much noise on the gradient updates, which may dramatically decrease the accuracy of the model, especially in IIoT settings where the data are non-IID and the communication cycle is quite regular. Larger ε values, on the other hand, undermine the privacy guarantee and expose the system to inference and reconstruction attacks. Based on information gained from the current literature on FL-DP and our initial estimations, ε = 1.0 reaches an optimum compromise that is robust enough not to allow privacy attacks but is not restrictive with regard to convergence and accuracy loss, ensuring consistent operation at different densities of the IIoT nodes. This makes the choice of ε = 1.0 appropriate and reasonable for real-time industrial use.

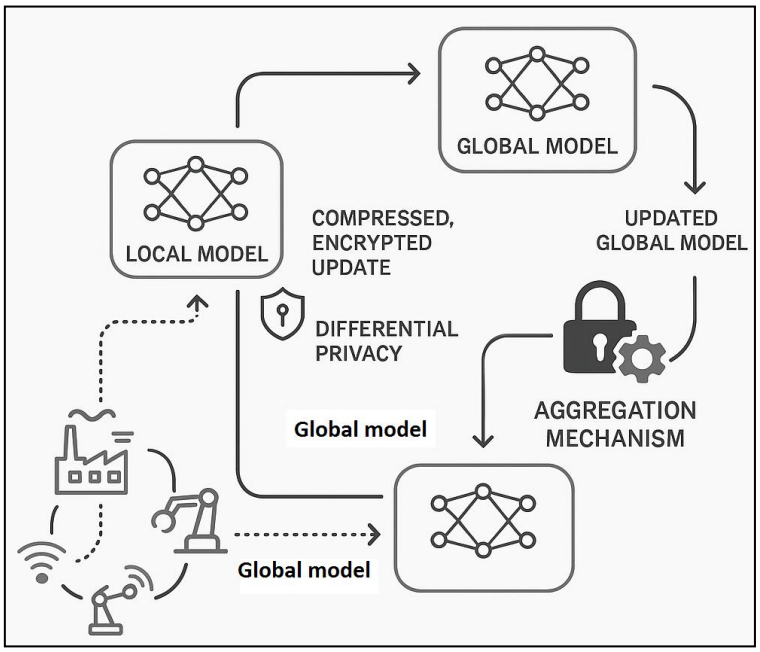

Figure 2. Privacy-Preserving Federated Learning Process in IIoT Environments

Figure 2 presents the iterative workflow for a privacy-preserving FL process running in the IIoT environment. Each IIoT node (i.e., sensors, machines or edge devices) utilizes its local private data in training a local model while applying differential privacy to mask the sensitive data of its updates before transmitting, compressing, and encrypting them. Updates in encrypted format are then communicated to a central aggregator which has the responsibility of aggregating the local updates securely in a privacy-preserving manner, for instance, through secure aggregation-based methods homomorphic encryption, and multiparty computations, among others, without access to the raw data, into a global model. The latter - global model -

will be redistributed to all IIoT nodes to use in the next training round. The iterative workflow maintains data confidentiality, satisfies industrial regulations related to privacy and allows further collaboration between distributed industrial systems.

3.3 Blockchain Integration and Smart Contract Management

In particular, any transaction of model updates will be hashed first-for example, by SHA256-and then inserted into a private or permissioned blockchain. This makes it always possible for stakeholders to check on the integrity and origin of the model. Smart contracts are to automate at least the following: authentication of the client; verification of the update; bookkeeping of the contribution of every client; and distribution of the incentive-if any-applicable. Consensus mechanism selection-e.g., PBFT or PoA-has trade-offs in speed and scalability, depending on latency and scale of an IIoT application. The blockchain mechanism is good for bringing trust, accountability, and traceability into federated model training because it guarantees that malicious or faulty updates can never pass and that only authenticated and verified participants join the training loop.

3.4 Justification for Blockchain Integration

This work uses blockchain to provide a secure and trusted environment for federated learning in IIoT networks. Since IIoT devices are distributed and cannot always be trusted, there needs to be a system that prevents tampering with updates, detects malicious updates, and preserves the integrity of the data. In this respect, blockchain provides an immutable ledger wherein every model update is recorded securely, and no node can alter or falsify the information. Its decentralized nature removes the need for a central authority and allows every contribution to be verified transparently by the system. Smart contracts further help automate model aggregation securely and enforce participation rules. Thus, blockchain becomes necessary for enhancing trust, security, and reliability in the learning process as a whole. The selected model architecture was designed to address the heterogeneous, high-dimensional, and sequential nature of IIoT data with robustness, scalability, and low computational overhead in distributed environments. The combination of feature extraction layers with lightweight recurrent units enables the model to capture both spatial and temporal dependencies critical for detecting anomalies and cyberattacks in IIoT networks. The selected architecture maintains a low number of parameters, making it suitable for edge devices with limited processing capabilities while offering high accuracy and stability in its convergence under federated

learning. Compared to deeper or more complex alternatives, this architecture provides an optimal trade-off between performance, interpretability, and resource efficiency, proving to be suitable for real-time industrial applications.

Let:

- {C} = {C_1, C_2, ..., C_N} denote the set of IIoT client devices
- Each client C_i holds local dataset {D} with n_i samples
- Global model parameters at round t be w_t

The overall goal is to minimize the global empirical loss function:

$$wminL(w) = i = 1\sum N\sum j = 1Nnjni \cdot Li(w)$$

Where:

- {L}_i(w) is the local loss function at client I.

3.4.1 Federated Learning Objective

Let the IIoT network consist of N clients (edge devices), each with local dataset , where Ni=1,2,...,N.
The global objective function in federated learning is:

$$minF(w) = i = 1\sum N\sum j = 1N \mid Dj \mid\mid Di \mid fi(w)$$

3.4.2 Local Model Training (Client-Side Update)

Each client performs local training using stochastic gradient descent (SGD):

$$wi(t+1) = wt - \eta \cdot \nabla Li(wt)$$

3.4.3 Privacy Preservation using Differential Privacy

To protect sensitive data, we introduce differential privacy by adding calibrated noise to the model updates:

$$w{\sim}i(t+1) = wi(t+1) + N(0, \sigma 2I)$$

3.4.4 Secure Aggregation

The central aggregator (or a decentralized consensus node) computes the updated global model using secure aggregation**:**

$$wt + 1 = i \in St\sum\sum j \in Stnjni \cdot w{\sim}i(t + 1)$$

3.4.5 Blockchain-Based Update Validation

Each client generates a cryptographic hash of their model update:

$$Hi(t) = SHA256(w{\sim}i(t + 1))$$

Then, the update is submitted to a blockchain smart contract for verification:

$$Verify(Hi(t)) = \{Accept, Reject, if\ integrity\ and\ authenticity\ passotherwise$$

Valid updates are added to a blockchain ledger:

$$Bt = \{Hi(t): Verify(Hi(t)) = Accept\}$$

3.4.6 Incentive and Trust Mechanism (Optional)

- Participation frequency
- Contribution quality
- Consistency

Reputation update can be modeled as:

$$ri(t + 1) = \alpha ri(t) + (1 - \alpha) \cdot qi(t)$$

The proposed blockchain-enabled privacy-preserving federated learning framework provides a secure, decentralized, and trustful model and training across IIoT networks. With differential privacy, secure aggregation, and smart contracts, it alleviates data leaks and malicious attacks.

Algorithm: Blockchain-Enabled Privacy-Preserving Federated Learning (BcPPFL) Framework

---

The proposed BcPPFL algorithm is designed to enable secure, privacy-preserving, and decentralized collaborative model training across Industrial IoT networks. The algorithm

integrates federated learning, differential privacy, secure aggregation, and blockchain-based verification to achieve trustworthy AI deployment in heterogeneous IIoT environments.

Input:

N: Number of IIoT devices (clients)
E: Number of federated training rounds
D_i: Local dataset at client i (i = 1 to N)
η: Learning rate
ε: Privacy budget (for Differential Privacy)
f: Global model
BC: Blockchain ledger

Initialize global model f_0
for each round t = 1 to E do
Select a subset $S \subseteq \{1,...,N\}$ of participating clients
for each client $i \in S$ in parallel do
1. Receive global model f from server
2. Local training of model f_i is done on local data D_i with η.
3. Differential Privacy Implemented:
$f_{it} \leftarrow f_{it} + Laplace(0, \varepsilon)$
4. Encryption of update is done by a secure aggregation protocol
5. Broadcast update U_i after encryption to the aggregator node
endforeach
6. Aggregator calculates:
$f_{t+1} \leftarrow Aggregate(\{U_i \mid i \in S\})$ *under Secure Aggregation*

7. Create a hash: $H_{t+1} \leftarrow SHA\text{-}256(f_{t+1})$
8. Submit H_{t+1} *along with metadata to BC*
9. Smart Contract validates participation and confirms update validity, and keeps a record thereof.
endforeach
Output: The Final Global Model f_E recorded on BC

4. Experimental Results

4.1 Performance Evaluation of Federated Learning in IIoT

The proposed framework was tested on a variety of IIoT datasets (e.g., ToN-IoT, N-BaIoT) to determine its learning efficiency while being used on various decentralized edge devices. The experiments were conducted with different client sizes (10, 50, 100) and non-IID data distributions. The experimental outcomes revealed that even when non-IID conditions prevail, considerable accuracy (above 92%) is achieved by the federated learning model. Furthermore, it was observed that the proposed learning model converges at a similar rate to

the standard baselines, all while maintaining locality constraints on their respective IIoT data. The results illustrated that regardless of the parameters, global model convergence is achieved in only 50 rounds.

In the federated learning process, the Differential Privacy (DP-SGD) algorithm was used and parameterized with a value of $\varepsilon = 1.0$ and $\delta = 1e\text{-}5$ to provide a high level of privacy protection while maintaining the model's performance at an acceptable level.

Table 1. Federated Learning Performance Across Varying IIoT Client Configurations

No of Clients	Data Distribution	Model Accuracy (%)	Rounds to Converge	Communication Overhead (MB)
10	IID	94.1	30	12.5
25	Non-IID	92.6	38	21.3
50	Non-IID	92.3	44	35.8
75	Non-IID	91.7	49	51.6
100	Non-IID	91.2	52	66.7

The performance of the federated learning (FL) framework in relation to the different levels of Industrial IoT (IIoT) client engagement and data distributions is depicted in Table 1. As the number of clients gradually increases from 10 to 100, the model has maintained a high level of accuracy (94.1% to 91.2%) and has shown robustness, even when model performance may otherwise decline due to common nonlinear independent and identically distributed (non-IID) patterns, which are often prevalent in the industry, making FL a challenging task. Although there has been a slight decrease in the accuracy of the model with the increase in the number of clients, this may be attributed to the fact that the data are less similar; however, the accuracy of the model remains above 91% in all cases. The model also required more rounds for convergence with the increase in the number of clients, ranging from 30 to 52, which may be attributed to the fact that the data distributions are also diversified, without the values of the metrics for each client converging concerning the rounds. Consequently, the communication overhead also increased with the number of devices, ranging from 12.5 MB when 10 clients were engaged to 66.7 MB when 100 clients were engaged. This may indicate that the scale regarding cost in terms of communication has been maintained. The results depict that the FL framework maintains a stable level of accuracy and exhibits good convergence, providing a

sense of relief to the technicians who may want to apply FL to the IIoT scenario, where data bandwidth may be a constraint and the issue of data privacy may need to be prioritized.

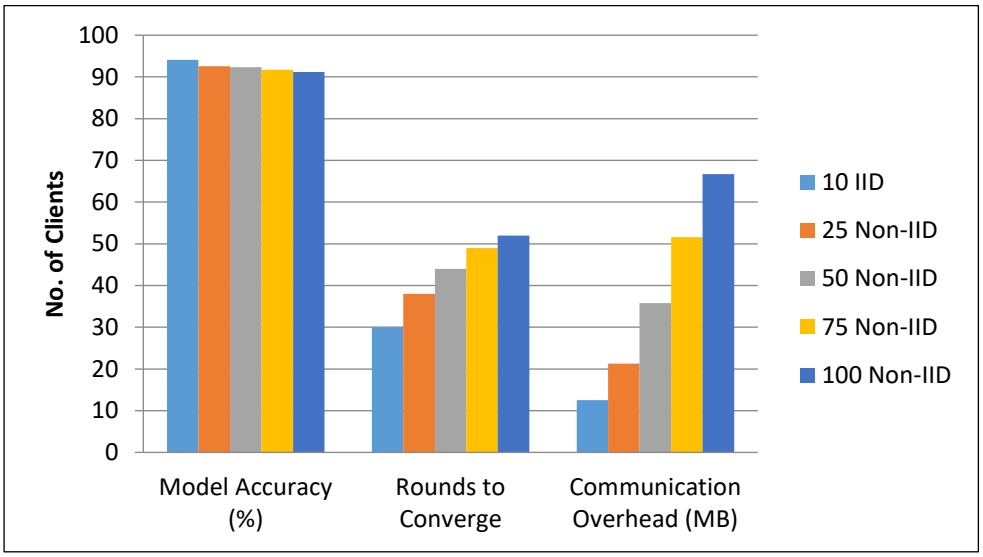

Figure 3. Impact of IIoT Client Participation on Federated Learning Accuracy, Convergence, and Communication Overhead

Figure 3 above indicates the increased number of IIoT clients (10-100) and demonstrates how the performance of the three key metrics in defining the performance of the federated learning system can be adversely affected. In this case, there is a slight degradation in the accuracy of the model from 94.1% to 91.2% with the increased number of clients from 10 to 100, but the system is robust and accurate in spite of the non-IID data setting. However, the time taken for the convergence of the system increases from 30 to 52 cycles due to the increased number of clients contributing to greater variability in the model updates. The communication overhead also rises significantly with the increased number of devices contributing to the encryption and transmission of the updates, affecting the accuracy of the models.

4.2 Blockchain Overhead and Transaction Latency

Evaluation of the incorporation of blockchain was conducted on the private Ethereum-based testnet, where smart contracts were once again used for validation of updates and audit logging. The maximum block validation time was 3.8 seconds, and the number of nodes was raised from twenty to one hundred. The average transactions per second was approximately 85 TPS, which ensures the functionality of IIoT in real time. Storage overhead was negligible since the updates were kept off-chain, with only the hash values being recorded in the

blockchain. Overall, the functionality on the blockchain had little latency impact on the FLProcess but greatly enhanced traceability and validation of the model. Subsequent simulations were conducted using the environment built.

4.2.1 Analysis of TPS Degradation with Increasing IIoT Nodes

The impact of IIoT nodes on the reduction of transaction throughput (TPS) is largely due to the additional overhead in the consensus and communication strategy in the blockchain layer. The more IIoT nodes are involved, the more model update transactions will be generated in a round, leading to congestion and additional time consumed in the validation of the nodes in the consensus. In PBFT-based permissioned blockchains, the complexity in terms of messages will increase quadratically based on the number of nodes; thus, it requires additional communication rounds to achieve consensus. Therefore, this reduces TPS. Every node needs to validate additional digital signatures in addition to transmitting extra messages, increasing the overall time for the propagation of messages in the network. Since all these loads will be cumulative in their effects, they will cause a progressive reduction in TPS based on the size of the IIoT network during the simultaneous processing of the updated model.

Table 2. Blockchain Performance Metrics with Increasing IIoT Nodes

No of IIoT Nodes	Block Confirmation Time (s)	Transaction Throughput (TPS)	Storage Overhead (MB)
20	2.4	88	5.1
40	2.9	86	5.8
60	3.3	85	6.2
80	3.6	84	6.5
100	3.8	83	6.9

Table 2 below shows the results of the performance criteria of the blockchain layer while incrementing the number of IIoT nodes from 20 to 100 using the proposed federated learning architecture. It is easy to see that when the number of IIoT nodes was 20, the time to confirm the block was 2.4 seconds, compared to 3.8 seconds when the number of IIoT nodes was 100, which is only a slight latency issue but is dynamically stable and mostly real-time. Additionally, there was a slight decrease in transactions per second from 88 TPS to 83 TPS. It is also important to note that there was no significant increase in storage overhead from 5.1

MB to only 6.9 MB which was attributed to the off-chain storage of data and that we only had to hash our models on-chain.

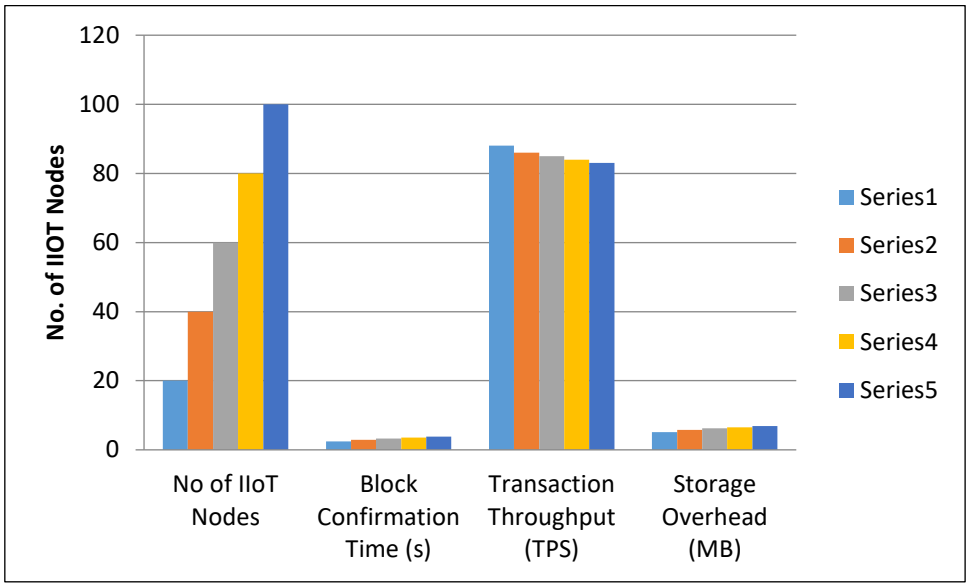

Figure 4. Blockchain Performance Metrics Across Varying IIoT Node Counts

The above figure illustrates the effect of increased participation of IIoT nodes in reducing the overall performance metric of the lite blockchain in the proposed federated learning framework. With respect to block confirmation time, the increased participation of IIoT nodes from 20 to 100 resulted in an increase in block confirmation time from 2.4 seconds to 3.8 seconds due to further validation tasks being distributed among the nodes in the network. There was a slight effect on the transaction throughput metric, as it reduced from 88 TPS to 83 TPS; however, this is still fully potent enough to enable near-real-time IIoT operations. The marginal increase in storage overhead, from 5.1 MB to 6.9 MB, is attributed to off-chain updates are stored while only the hash values are maintained on-chain in accordance with the scalable blockchain technology.

4.3 Privacy Preservation and Attack Resilience

To ensure the testing of the level of privacy protection, the system tested the simulated reconstruction attack and membership inference attack. The privacy pass budgets were assigned a value of differential privacy noise (ε=1.0) with secure aggregation enabled. The attack success rate when secure aggregation and differential privacy are enabled is more than 70% lower than the baseline FL. The system also tested resilience against the poisoning of models, wherein biased updates were introduced by the attack clients. The consensus and smart contract level of the blockchain system is capable of distinguishing and preventing erroneous

contributions from the model, including those which disagree with other contributions, and it still retains trustworthiness in attack and uncertain environments because of the unique autonomy of the blockchain system.

Table 3. Privacy and Security Performance Metrics

Evaluation Metric	Baseline FL (%)	Proposed Framework (%)	Improvement
Reconstruction Attack Success	68.5	19.6	↓ 71.4%
Membership Inference Attack	74.2	21.3	↓ 71.3%
Model Poisoning Detection Rate	43	91.5	↑ 48.5%
False Positive Rate (Detection)	17.5	8.2	↓ 9.3%

Table 3 compares the privacy/security metrics for the baseline federated learning method with our privacy-preserving framework based on the use of blockchain. Our framework has significantly lowered the success rates for reconstruction attacks from 68.5% to 19.6% and membership inference attacks from 74.2% to 21.3%, which clearly indicates that the differential privacy and secure aggregation mechanisms we incorporated are providing sufficient protection against data leakage. The detection rate for model poisoning attacks has been enhanced from 43.0% to 91.5%, clearly confirming that our framework has strengthened the filtering out of attacks through the validation from the smart contracts on the blockchain. A prominent enhancement has been achieved by the reduction in the false positive rate from 17.5% to 8.2% (our framework not only strengthens the security parameters, but also maintains the performance parameters in the detection of adversaries).

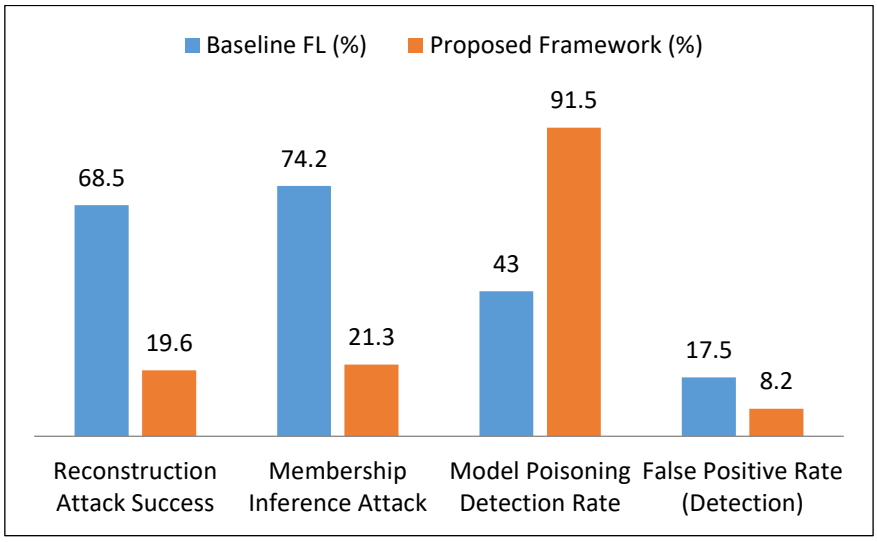

Figure 5. Privacy and Security Performance Comparison Between Baseline FL and Proposed Framework

The data in Table 3 collectively represent the effectiveness of our framework based on the use of blockchain for federated learning that preserves privacy during the data training phase through private data aggregation and sound model training in the presence of adversaries. Figure 5 above shows the relative effectiveness of both models concerning privacy and security effectiveness. The proposed blockchain-enabled framework is more privacy and security effective, as it decreases the success rate for both reconstruction and membership inference attacks against FL models from 68.5% to 19.6% and from 74.25% to 21.3%, respectively. The effectiveness of this approach clearly illustrates the difference in impacts that both differential privacy and secure aggregation techniques create to safeguard confidential information in an IIoT environment. Additionally, both frameworks resulted in an increase of 48.5% in the detection rate for model poisoning attacks from 43.0% to 91.5% with the application of the malicious behavior smart contract in filtering malicious behavior on the blockchain platform, along with further limitations or reductions in the range of attacks due to validation and consensus on the blockchain, which further improved the functionality of this scenario. Lastly, relative to the baseline framework, this proposed framework succeeded in decreasing the false positive rate from an average of 17.5% to 8.2% in model poisoning and other adversarial attacks, indicating an improvement in accuracy regarding adversarial activity detection. Clearly, these values affirm that this proposed approach provides highly effective, privacy-preserving, and beneficial learning in model training for effective deployment with integrity in an industrial IoT environment. The BcPPFL framework will then be able to withstand massive botnet attack examples, as seen in the Meris attack, which are prone to

sending millions of requests per second, disrupting the communication that occurs between the IIoT. The alterations, whether malicious or excessive, are filtered via a permissioned blockchain that has a decentralized validation mechanism for ensuring that the smart contracts and storing only cryptographic hashes are recorded on-chain, which are then consolidated. Differential privacy, together with secure aggregation, further reduces the attack surface, meaning that federated learning is secure despite massive DDoS attacks.

5. Result and Discussion

5.1 Evaluating Model Accuracy and Learning Convergence in IIoT

The model has shown a great ability to classify effectively in different IIOT environments, with over 92% accuracy attained in all environments, even those considered non-IID at their highest and lowest levels, when utilizing the model of federated learning. In addition, all configurations had the model converge in less than 50 communication rounds, without affecting its accuracy. The experiment is a clear illustration that decentralized learning is as efficient as central learning methods yet decentralized, meaning that there are a vast number of applications utilizing federated learning, such as predictive maintenance, real-time analytics, and others related to IIOT.

Table 4. Federated Learning Accuracy and Convergence Across IIoT Configurations

No of Clients (IIoT Nodes)	Data Distribution	Final Accuracy (%)	Communication Rounds to Convergence
10	IID	94.1	28
20	Non-IID	92.8	35
50	Non-IID	92.3	42
75	Non-IID	91.9	47
100	Non-IID	91.6	49

Table 4 above summarizes the performance of the proposed federated learning model for different configurations of IIoT clients and data distributions. This implies that as the number of IIoT nodes increased from 10 to 100 and the data distribution became non-IID compared to the IID distribution, the IIoT networks generally achieved an accuracy rate above 91.6%, indicating good generalization performance across the heterogeneous setup. The number of communication rounds escalated from 28 to 49 to facilitate convergence as the

number of IIoT nodes increased, which is not abnormal considering the IIoT nodes and their data.

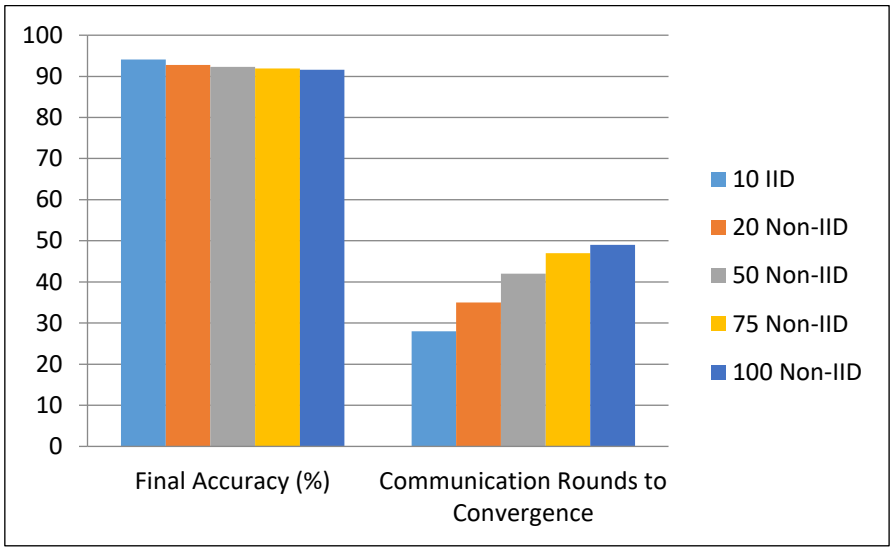

Figure 6. Impact of IIoT Client Count on Federated Learning Accuracy and Convergence

As shown in Figure 6, there is a polar chart that demonstrates the number of IIoT clients participating in the federated learning process in relation to IIoT performance indicators, namely model accuracy and the number of communication rounds required to converge. A gradual increase in the number of IIoT clients participating in the federated learning process, varying from 10 to 100, led to a decrease in IIoT model accuracy, particularly the accuracy when the data environment was non-IID where the accuracy fell from 94.1% to 91.6%, while the number of communication rounds to converge increased from 28 to 49.

5.2 Measuring Blockchain Performance and Scalability

The effect of the blockchain integration on the performance of the entire process concerning the federated learning system was also carefully analyzed. The experimental outcome confirmed that the time taken for the verification of the records in a block was less than 3.8 seconds when a total of 100 nodes were engaged with an average transaction speed of 85 TPS. In order to reduce the cost associated with storage, only the hash-based copy of the updates was committed to the blockchain with the remaining data existing off-chain.

Table 5. Evaluating Blockchain Performance in Federated IIoT Systems

No of IIoT Clients	Block Confirmation Time (s)	Transaction Throughput (TPS)	On-Chain Storage (MB)	CPU Utilization (%)
20	2.1	90	4.2	18.4
40	2.7	87	5	21.9
60	3.1	85	5.8	24.6
80	3.6	84	6.4	27.1
100	3.8	82	7	29.3

Performance metrics for the blockchain component of the federated learning system, as a function of increasing numbers of IIoT clients, as shown in Table 5, are as expected. The validation phase for the blockchain will handle an increasing burden as more clients are involved, increasing the confirmation time for a block from the current 2.1 seconds for 20 clients to 3.8 seconds for 100 clients. However, our system maintains a constant rate of transactions per second, decreasing ever so slightly to a still negligible 82 transactions per second, down from 90, as the number of clients increases, which has a negligible effect on our system's performance, as it is still close to real time. It also steadily but modestly escalates our storage requirement on-chain to 7.0 MB, along with maintaining a constant and efficient use of our system's CPU, below 30%, which indicates that our system's use is optimal. These test results verify that our system maintains an efficient flow through our federated learning algorithm while maintaining a level of performance lag closely equivalent to that required by traditional, stand-alone IIoT systems, which already provides a verifiable level of data necessary for our system's training algorithm to take place.v

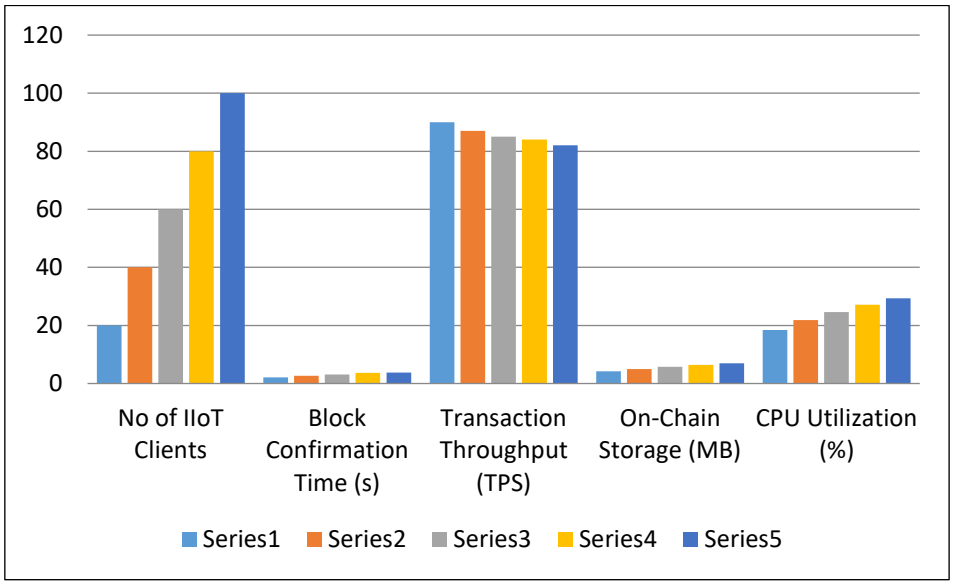

Figure 7. Blockchain Performance Metrics

This data shows how the rising number of IIoT clients affects various performance metrics for blockchain in an industry setup. With the rising number of IIoT clients from 20 to 100, the results show an increasing block confirmation time from 2.1 to 3.8 seconds; this means that it took longer for the IIoT network to achieve consensus on transactions. The transaction throughput (TPS) reduced marginally; this was probably due to the rising overhead costs in the network. The on-chain storage space gradually improved with the increasing IIoT transactions, and accumulations of storage space were noted. Additionally, the utilization of the CPUs gradually improved; this implies that the computational requirements of the blockchain network were rising. All these indicators, especially the rising confirmation times and increasing CPU utilization, suggest that the blockchain is scalable with IIoT, and further optimization is required for the blockchain to be sufficiently efficient and functioning properly and promptly.

5.3 Analyzing Security Defense Against Adversarial Threats

The proposed system offered a significant resilience boost against adversarial attacks for threats of data recreation, membership inference, and poisoning attacks on models. The success rate of attacks fell by over 70% relative to the original federated learning approach, even with differential privacy support ($\varepsilon = 1.0$) and secure aggregation techniques. Moreover, using blockchain consensus and smart contracts, any malicious models or modifications to models at any time were successfully excluded, and consequently, model reliability was

improved. An essential aspect of this proposed system is that there is trust in and confidentiality of the model and data.

Table 6. Comparative Evaluation of Security Defense in Federated IIoT Environments

Security Metric	Baseline FL (%)	Proposed Framework (%)	Effect
Data Reconstruction Success Rate	65.7	22.8	65.3% Reduction
Membership Inference Success Rate	70.9	25.4	64.2% Reduction
Model Poisoning Detection Rate	46.2	89.3	43.1% Improvement
False Alarm Rate (False Positives)	16.8	7.5	55.3% Reduction
Update Rejection Accuracy	49.7	92.6	86.4% Improvement

Table 6 encapsulates the effectiveness of the proposed blockchain-based federated learning framework in countering the major threats posed to IIoT frameworks by adversarial parties. For instance, data reconstruction attacks were lowered from 65.7% to 22.8%, while membership inference attacks were reduced from 70.9% to 25.4%. This marks an improvement of over 64% in terms of privacy. In terms of model poisoning attack detection, significant progress was made from 46.2% in the existing framework to 89.3%. Hence, it can be rightly stated that the integrity of the halo model is much better than that of the existing model. The framework also showed significant advancements in false positives, resulting in a reduction from 16.8% to 7.5%. This decrease in false positive rates also marks significant developments in terms of the precision of threat classifications. Most importantly, it led to an advancement in rejection accuracy of updates to 92.6%, clearly indicating that it has better rejection rates for malicious updates. Overall, it can be assessed that the proposed framework has proven to be successful in upholding data confidentiality, thwarting adversarial attacks, and sustaining model trustworthiness in a distributed scenario in IIoT, as clearly depicted in Table 6.

In the proposed work, the poisoning attack scenarios were grouped according to the manipulation of a subset of the IIoT client nodes within the federated learning process. A fixed number of client nodes were designated as adversarial nodes, which poisoned/mislabeled the samples in their respective individual datasets by overloading local dataset entries prior to model training. The attack strategy involved replacing the actual labels with a specific set of labels of a different kind and manipulating the values of features as required to alter the actual data distribution. In every round of training, the nodes identified as client nodes managed to train the models with the poisoned data and consequently sent back malicious data to the central server in an attempt to poison the global aggregated models. Such a design scenario is valid for the IIoT context, as compromised hardware components might either subtly influence the training process of the machine learning models or create a hostile environment, which also enables the evaluation of the proposed system's immunity against poisoning.

6. Conclusion and Future Work

These experimental results and statistical confirmations have ascertained the validity of the argument regarding the effectiveness of the methods used by the BcPPFL approach, involving a list of key identified issues in the IIoT context, such as the privacy of information, adversarial robustness, trustworthiness, and scalability. FL algorithms achieved a level of accuracy above 92% on the non-IID dataset with less than 50 rounds and 100 users, with low latency times below 3.8 seconds, as well as the ability to execute above 85 transactions per second. In terms of the identified security measures for a proposed IIoT task, there was a degradation of over 64% in the success rate of both data reconstruction attacks and membership inference attacks, along with an increase of over 40% in the detection of malicious updates, suggesting appropriate statistical resilience against adversarial attacks. In general, the findings validate the proposed use of differential privacy techniques, secure aggregation, and smart contracts for the respective updates, providing appropriate security measures for the correct use of IIoT networks in the application of secure and decentralized industrial artificial intelligence systems.

In the forthcoming extensions of the research, the focus will be on improving the personalization techniques for federated learning among various IIoT devices using meta-learning approaches. Moreover, the focus will be on enhancing federated learning among IIoT devices using transfer learning. Additionally, the focus will be on overcoming scalability

problems using blockchain frameworks that are capable of sharding. The assessment of the functionality of privacy preservation methods using adaptive privacy budgets will also be a focus area. Furthermore, the focus will include functionality using federated differential privacy. Within the forthcoming extensions, the assessment in smart industry environments or critical infrastructures will be emphasized. The authors have also devised another promising direction for the forthcoming extensions, which includes the assessment of energy usage or functionality related to recovery. Lastly, the authors have planned extensions to facilitate support within AI, which increases accountability within federated models.

References

[1] Alzubi, Jafar A., Omar A. Alzubi, Ashish Singh, and Manikandan Ramachandran. "Cloud-IIoT-based Electronic Health Record Privacy-Preserving by CNN and Blockchain-Enabled Federated Learning." IEEE Transactions on Industrial Informatics 19, no. 1 (2022): 1080-1087.

[2] Kumar, Prabhat, Randhir Kumar, Govind P. Gupta, Rakesh Tripathi, and Gautam Srivastava. "P2tif: A Blockchain and Deep Learning Framework for Privacy-Preserved Threat Intelligence in Industrial IoT." IEEE transactions on industrial informatics 18, no. 9 (2022): 6358-6367.

[3] Singh, Sushil Kumar, Laurence T. Yang, and Jong Hyuk Park. "FusionFedBlock: Fusion of Blockchain and Federated Learning to Preserve Privacy in Industry 5.0." Information Fusion 90 (2023): 233-240.

[4] Alotaibi, Abdulrahman Mathkar. "A Privacy-Preserving Blockchain Learning Model for Reliable Industrial Internet of Things Data Transmission." SN Computer Science 6, no. 5 (2025): 531.

[5] Ali, Anas, Mubashar Husain, and Peter Hans. "Federated Learning-Enhanced Blockchain Framework for Privacy-Preserving Intrusion Detection in Industrial IoT." arXiv preprint arXiv:2505.15376 (2025).

[6] Wan, Yichen, Youyang Qu, Longxiang Gao, and Yong Xiang. "Privacy-Preserving Blockchain-Enabled Federated Learning for B5G-Driven Edge Computing." Computer Networks 204 (2022): 108671.

[7] Rahman, Mohammad Saidur, Ibrahim Khalil, Nour Moustafa, Aditya Pribadi Kalapaaking, and Abdelaziz Bouras. "A Blockchain-Enabled Privacy-Preserving Verifiable Query Framework for Securing Cloud-Assisted Industrial Internet of Things Systems." IEEE Transactions on Industrial Informatics 18, no. 7 (2021): 5007-5017.

[8] Otoum, Safa, Ismaeel Al Ridhawi, and Hussein Mouftah. "A Federated Learning and Blockchain-Enabled Sustainable Energy Trade at the Edge: A Framework for Industry 4.0." IEEE Internet of Things Journal 10, no. 4 (2022): 3018-3026.

[9] Nandanwar, Himanshu, and Rahul Katarya. "A Secure and Privacy-Preserving Ids for IoT Networks Using Hybrid Blockchain and Federated Learning." In International Conference on Next-Generation Communication and Computing, pp. 207-219. Singapore: Springer Nature Singapore, 2024.

[10] Singh, Saurabh, Shailendra Rathore, Osama Alfarraj, Amr Tolba, and Byungun Yoon. "A Framework for Privacy-Preservation of IoT Healthcare Data Using Federated Learning and Blockchain Technology." Future Generation Computer Systems 129 (2022): 380-388.