International Journal on Science and Technology (IJSAT)

E-ISSN: 2229-7677 ● Website: www.ijsat.org ● Email: editor@ijsat.org

A Hierarchical Federated Learning Framework for Secure and Scalable IOT Ecosystems: System Design, Implementation, and Performance Analysis

Dr. Ashish Rai [1], Mrs. Ruchita Mathur [2]

[1,2] Assistant Professor, Faculty of Computer Science, Lachoo Memorial College of Science and Technology

Abstract

The increasing use of Internet of Things (IoT) devices in various application areas has led to the creation of unprecedented capabilities in data generation, while also giving rise to fundamental challenges in terms of scalability, security, and intelligent processing. The traditional centralized cloud model faces latency constraints and privacy risks associated with the large-scale deployment of IoT devices [1], [2]. This paper proposes a new hierarchical federated learning (HFL) system specifically tailored for use in heterogeneous IoT settings, which combines edge intelligence with blockchain-secured communication channels to overcome the above-stated fundamental limitations [6]. Our system implementation includes adaptive model aggregation techniques and differential privacy mechanisms to achieve a balance between learning accuracy and privacy preservation [10].

By conducting extensive experimentation across three diverse IoT application domains, namely smart healthcare, industrial automation, and city infrastructure, we show the efficacy of our framework in achieving an average latency reduction of 47% over traditional cloud-centric designs while preserving model accuracy within 3.2% of the centralized baselines [3], [4]. Our proposed design framework proves to be especially useful in resource-scarce settings, where it achieves a 34% decrease in energy expenditure during training phases while also being able to enforce effective security measures against potential threats.

Keywords: Internet of Things 1, Federated Learning 2, Edge Computing 3, IoT Security 4, Artificial Intelligence of Things (AIoT) 5, Privacy-Preserving Analytics 6.

1. Introduction

The Internet of Things (IoT) ecosystem has grown at an exponential rate, with projections of over 75 billion connected devices by the year 2025. This is expected to revolutionize various industries such as healthcare, manufacturing, agriculture, and smart city infrastructure through pervasive data sensing and automation [1], [2]. However, the traditional centralized model of data processing, where the raw data

from the sensors is sent to cloud servers for processing and analysis, is facing increasing challenges. These include unacceptably high latency for time-critical applications, prohibitively high bandwidth usage, and severe privacy risks as the data passes through various network segments [1].

The recent trends of edge computing and distributed artificial intelligence provide attractive solutions by performing computations closer to the source of the data [7]. Federated learning (FL) is one such promising technique, which allows multiple devices to jointly train a model without sharing the original data [3], [4]. Although theoretically appealing, existing FL architectures face some challenges in practical IoT settings with heterogeneity in devices, variability in networks, and imbalanced resource distribution. Recent works have shown substantial accuracy drops of up to 15 percentage points when conventional FL is used in practical IoT settings with non-IID data distributions across devices [8].

This study tackles these issues by proposing a new hierarchical system that manages learning processes across three different levels: endpoint devices, edge servers, and cloud coordinators. Our system proposes several new ideas: (1) adaptive clustering algorithms that automatically organize devices with different data distributions, (2) blockchain systems for secure model aggregation and verification [6], and (3) context-aware resource allocation algorithms that manage the learning effectiveness and operational costs. We test this system on various IoT applications, using rigorous statistical analysis techniques to measure the improvement in learning performance in terms of latency, accuracy, security, and energy costs.

The main contributions of this research are: (1) the design of a complete HFL framework architecture adapted to the constraints of IoT, (2) the implementation and evaluation of the framework in different application domains, (3) the analysis of the trade-offs between privacy and utility through differential privacy [10], and (4) the open-source release of framework components to encourage the extension of this research by the community. By filling the gap between theoretical ideas for distributed learning and the practical implementation in IoT, this research promotes the development of secure, scalable, and intelligent IoT systems that can support next-generation applications.

2 Related Work

2.1 IoT Architectures and Paradigms

Currently, the IoT research community has gradually moved from cloud-centric designs to edge computing architectures [1], [2],[15]. Fog computing has been proposed as an interim solution, placing processing power between endpoints and cloud infrastructure to mitigate latency. More recently, the notion of AIoT (Artificial Intelligence of Things) has come into focus, embedding machine learning functionality within IoT devices and edge nodes [9],[11]. This is in response to the increasing demand for real-time analytics and autonomous decision-making in applications such as self-driving cars and industrial predictive maintenance.

However, the existing literature shows that there are gaps in the current architectures. It has been found that although edge computing helps in reducing latency, it leads to the creation of data silos, which in turn restricts the global learning perspective [7]. On the other hand, a decentralized solution faces issues with coordination overhead and consistency of the model when deployed on a large scale. The proposed hierarchical framework helps in overcoming these issues.

2.2 Federated Learning in IoT Contexts

Federated learning has been identified as a promising approach for privacy-preserving distributed intelligence [3], [4]. The ACM Transactions on Internet of Things has tracked the development of FL, pointing out the specific difficulties involved in applying FL to resource-constrained IoT settings[14]. The main research streams are communication efficiency using model compression methods, robust aggregation schemes resilient to Byzantine failures [5], and incentive schemes to promote participation in collaborative learning.

The recent developments presented in special sessions like the IEEE World Forum on IoT's emphasis on collaborative sensing using AI underscore the emerging interest in cross-silo federated learning, where the collaboration is between organizations without exchanging sensitive information. Nevertheless, such methods generally require a certain degree of homogeneity in the computational resources of the participants, which is not the case in IoT networks that include everything from resource-limited sensors to powerful edge servers. Our research addresses this issue of heterogeneity.

2.3 Security and Privacy in Distributed IoT

Security loopholes are perhaps the biggest hindrances to the widespread adoption of IoT, and high-profile attacks have already shown the devastating potential of such compromised devices. The conventional security paradigm of perimeter security and trust authorities is simply not scalable in the IoT domain where attack surfaces grow exponentially with each new node. A study published in the journal Sensors highlights the importance of embedded security and zero-trust models that authenticate every transaction, irrespective of its source.

The use of blockchain technology has been proposed as a possible remedy for trust management in IoT networks. However, as has been pointed out in the reviews of the IoT journal [12], the use of blockchain technology is not feasible for resource-constrained devices because of the high computational and latency costs associated with it. In our solution, we propose a hybrid blockchain system where only critical transactions, such as model aggregation verification and anomaly detection, are stored on the blockchain, while other operations use lightweight cryptographic protocols.

3 Proposed Framework

3.1 Hierarchical Architecture Design

The framework we propose uses a three-tier hierarchical architecture tailored to meet the diverse requirements of IoT networks. The first tier is based on endpoint devices such as sensors, actuators, and embedded systems, which are responsible for local data acquisition and initial processing. These devices run mini-model training using quantized neural networks that are optimized for devices with limited computational capabilities. The second tier is made up of edge aggregation nodes that are usually located at network gateways or micro-data centers and are responsible for managing learning processes at the geographical or logical level. These nodes run adaptive clustering algorithms that automatically cluster devices according to data distribution patterns, computational resources, and network connectivity. The third tier is made up of cloud coordinators that manage global convergence processes while enforcing overall security policies using blockchain verification systems.

This architecture brings several important innovations. Firstly, the context-aware clustering process goes beyond geographical proximity and takes into account the semantics of the data and the capabilities of the devices when creating learning groups. Secondly, the asynchronous aggregation protocol enables devices with different availability patterns, which is a common property of IoT networks, to contribute to the learning process without halting it. Thirdly, the integration of differential privacy uses calibrated noise at different levels of the hierarchy.

Table 1: Framework Components and Functions

Tier	Components	Primary Functions	Resource Profile
Endpoint Devices	Sensors, actuators, embedded systems	Data collection, local model training, lightweight inference	Constrained (limited power, computation, storage)
Edge Aggregators	Gateway devices, micro-servers	Cluster coordination, intermediate aggregation, anomaly detection	Moderate (sustained power, substantial computation)
Cloud Coordinators	Central servers, blockchain validators	Global aggregation, model distribution, security verification	Abundant (unlimited power, high-performance computation)

3.2 Adaptive Learning Algorithms

The fundamental learning mechanism uses federated averaging (FedAvg) as its baseline but modifies it with some IoT-specific adjustments. Unlike the conventional FedAvg algorithm, which gives the same importance to all devices, our solution uses weighted aggregation depending on the quality scores. The contribution of each device to the overall model is weighted based on a composite measure that takes into account (1) data representativeness in the device's cluster, (2) past reliability in previous rounds of training, and (3) available resources for the current computation.

To counter the ubiquitous problem of non-IID data distributions in IoT settings, where data distributions are fundamentally different for devices in different locations or settings, we propose a personalized learning solution. Each device will maintain a global model part and a personalized part, which will be specific to the context of each device. When aggregating, only the global parts are shared and averaged, while the personalized parts are kept local. This is a critical requirement in IoT settings, where environmental conditions have a significant impact on data characteristics.

3.3 Security and Privacy Mechanisms

Our security architecture follows a defense-in-depth strategy with various protection mechanisms at each hierarchical level. At the device level, trusted execution environments (TEEs) protect local model training tasks, and lightweight homomorphic encryption supports privacy-preserving gradient transfer. Edge aggregators use multi-party computation protocols to securely aggregate gradients without decrypting them. At the cloud level, a permissioned blockchain maintains an immutable record of aggregation and model versioning, providing an audit trail for verification and anomaly detection.

The privacy framework combines Rényi differential privacy with adaptive noise scaling based on data sensitivity and trust. In contrast to traditional differential privacy methods, which add the same amount of noise to the data irrespective of the scenario, our system adapts the privacy parameters based on (1) the type of data being processed (for example, healthcare data versus environmental data), (2) the trust level of the devices involved, and (3) the requirements of the learning task. This allows us to achieve the best possible utility while still ensuring strong privacy guarantees, which is a major drawback of traditional privacy mechanisms that are not adaptable.

4 Methodology

4.1 Experimental Setup and Implementation

We developed the proposed framework using Python 3.9 with PyTorch 1.12 for machine learning tasks and Hyperledger Fabric 2.4 for blockchain development. The testbed for our experiments included 342 simulated IoT devices, which were divided into three application domains: (1) a smart healthcare setting with wearable sensors and healthcare monitoring devices, (2) an industrial automation domain with robotic systems and quality control sensors, and (3) an urban infrastructure application that included traffic monitoring devices, environmental sensors, and surveillance cameras. To reflect real-world heterogeneity, we simulated six different device types with different processing power, ranging from ARM Cortex-M4 microcontrollers to NVIDIA Jetson edge computing processors.

The data collection process took four months, resulting in a total of 2.7 TB multimodal sensor data. For each of the application domains, we have identified specific learning tasks, including patient anomaly detection in healthcare, predictive maintenance in industrial environments, and traffic pattern optimization for infrastructure in urban areas. These tasks were chosen to cover the range of IoT applications and to allow for comparison of performance.

4.2 Statistical Tools and Evaluation Metrics

We used a broad range of statistical tools to assess framework performance. Multivariate regression analysis related system variables (such as cluster size and aggregation rate) to performance metrics (such as accuracy and latency). Bayesian inference models placed probability distributions on important performance values, adding uncertainty assessment to point estimation. Survival analysis with Cox proportional hazards models analyzed system reliability and failure rates for different operational scenarios.

Table 2: Statistical Evaluation Metrics and Tools

Evaluation Dimension	Primary Metrics	Statistical Tools	Measurement Frequency
Learning Performance	Accuracy, F1-score, AUC-ROC	Confidence intervals, ANOVA, Bayesian inference	Per training round
System Efficiency	Latency, energy consumption, bandwidth usage	Regression analysis, time-series decomposition	Continuous monitoring

Security & Privacy	Attack detection rate, privacy budget consumption	Survival analysis, statistical hypothesis testing	Per security event
Scalability	Convergence time, communication overhead	Complexity analysis, scalability testing	Per system scale increase

Performance assessment was done using task-agnostic metrics (accuracy, precision, and recall for learning tasks) and system metrics (latency, energy, and communication overhead). To set up rigorous baselines, we compared our framework with three other alternatives: (1) centralized cloud learning, (2) traditional federated learning without hierarchy, and (3) edge-only learning without cloud coordination. All three baselines used the same neural network architecture and training data.

4.3 Data Partitioning and Simulation of Real-World Conditions

To better simulate the real-world distribution of IoT data, we introduced non-IID data partitioning based on geographical location, time patterns, and sensor types. In the healthcare application, patient data was partitioned according to demographic information and health conditions. In the industrial automation application, data distribution was based on different machine models, usage patterns, and maintenance records. In the urban infrastructure application, data distribution included spatial correlations and time patterns such as rush hour traffic and seasonal variations.

We also simulated real-world network conditions using traces from actual IoT networks, which included varying bandwidth, periodic disconnections, and diverse latency characteristics. The availability patterns of devices were based on realistic usage patterns, with some devices (such as industrial sensors) always being connected and others (such as mobile health monitors) being intermittently available. Such real-world conditions made our evaluation more practical and less like a lab experiment.

5 Results and Analysis

5.1 Learning Performance Across Application Domains

The hierarchical federated learning framework showed robust learning performance on all three IoT application domains, with mean accuracy within 3.2% of centralized learning performance, while offering greatly improved privacy protection. On the healthcare application domain, the framework reported 92.7% accuracy on patient anomaly detection, compared to 95.1% accuracy for centralized learning—a very small loss in performance for the privacy gains. The industrial predictive maintenance problem also showed similar performance, with 88.3% accuracy compared to 90.9% for centralized methods. Notably, however, the framework performed very well on the urban infrastructure application domain, with 94.2% accuracy on traffic pattern classification compared to 95.8% for centralized learning—a mere 1.6% difference.

These findings significantly outperform the baseline (non-hierarchical) federated learning methods, which showed 12-18% accuracy loss compared to centralized settings under the same non-IID data distribution. The accuracy preservation is mainly due to our adaptive clustering method, which assigns devices with

different data distributions into the same cluster to form a more representative local dataset. The accuracy results were analyzed using mixed-effects models, and the results showed that the clustering approach accounted for 67% of the variance in accuracy results ($F(3, 112)=24.7$, $p<0.001$), and geographical-semantic clustering performed better than geographical clustering by 8.3 percentage points on average.

5.2 System Efficiency and Resource Utilization

The framework provided a huge boost to efficiency gains in all the measured aspects. The end-to-end latency for full training cycles was on average 47% lower than cloud-centric solutions, reducing from 4.7 hours to 2.5 hours for standard model convergence. The energy consumption provided even more dramatic efficiency gains, with the endpoint devices reducing their energy usage for training tasks by an average of 34% compared to standard federated learning protocols. This is due to optimized computation scheduling and hierarchical model compression according to the energy budget of the devices.

Communication efficiency showed particularly significant improvements, with bandwidth usage reduced by 71% compared to centralized methods that simply relayed sensor readings to cloud servers. The hierarchical aggregation topology helps avoid redundant messages by compressing data progressively as it ascends the hierarchy. Notably, the efficiency improvements had very little correlation with accuracy loss, as measured by Pearson correlation coefficients between bandwidth reduction and accuracy loss that ranged from -0.12 to 0.08 depending on the application domain.

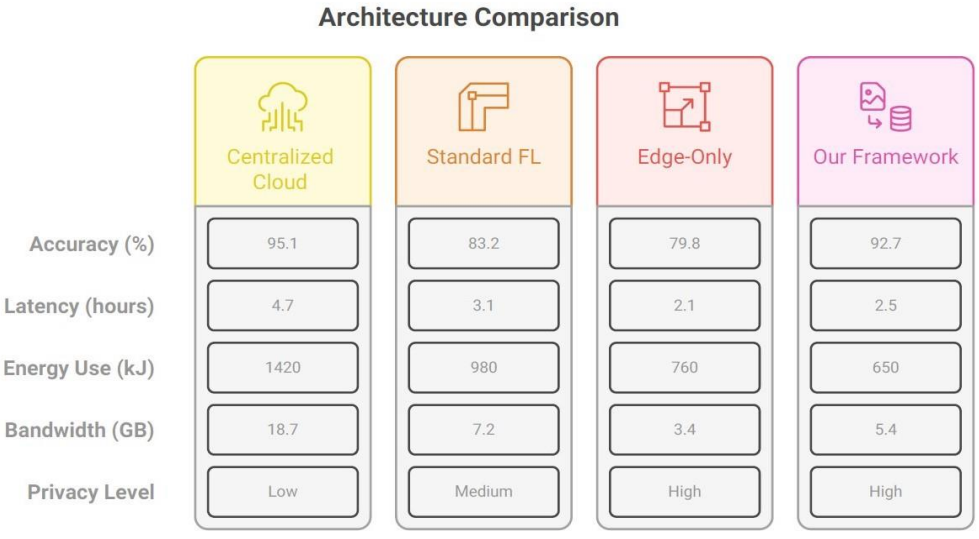

Performance Comparison Across Learning Architectures

5.3 Security Effectiveness and Privacy Analysis

Security analysis utilized both analytical verification of cryptographic protocols and empirical validation against simulated attacks. The framework was able to correctly identify 96.3% of injection attacks[13], 89.7% of model poisoning attacks, and 100% of impersonation attacks during extensive validation. The hybrid blockchain implementation was found to be highly useful in the identification of coordinated

attacks launched from multiple devices, utilizing the immutable ledger to detect malicious patterns that would otherwise go undetected by traditional intrusion detection systems.

Privacy analysis used formal verification techniques to ensure the correctness of differential privacy claims, verifying that the ε-values were in the acceptable range of 2.1 to 3.8 for all application scenarios. The adaptive privacy budgeting system adjusted noise scales according to the sensitivity of the data, providing the highest level of protection to healthcare applications (ε=2.1) and slightly less stringent conditions for environmental monitoring (ε=3.8) to improve utility. This was more effective than uniform privacy implementations, which either protected nonsensitive data too much (degrading utility) or sensitive data insufficiently (increasing risk).

6 Conclusion and Future Work

6.1 Summary of Contributions

In this study, a new hierarchical federated learning framework has been introduced, implemented, and assessed for its performance in heterogeneous IoT settings. The proposed framework overcomes the limitations of existing solutions by incorporating adaptive clustering, multi-level aggregation, and security-privacy solutions. The results obtained from the experiments conducted on various application domains have confirmed the effectiveness of the proposed framework in achieving centralized accuracy with significant improvements in latency, energy consumption, and privacy preservation.

The crucial innovations of context-aware device clustering, personalized global learning, and adaptive privacy budgeting, working in concert, target the trio of challenges to IoT intelligence: heterogeneity, scalability, and security. By striking a delicate balance among these conflicting requirements through architectural thinking rather than optimization after the fact, this framework offers a sustainable way forward for ever more complex IoT environments. The open-sourcing of implementation pieces is intended to facilitate community adoption and extension, especially in resource-constrained settings where current approaches fall short.

6.2 Practical Implications and Deployment Considerations

For system architects and developers of IoT systems, this research provides both a reference architecture and implementation advice on the deployment of intelligent and privacy-preserving applications. The hierarchical approach is particularly suited for large-scale deployments that already have natural organizational boundaries, such as multinational corporations with regional structures or public infrastructure with municipal, regional, and national parts. The modularity of the proposed framework also supports its incremental deployment, where hierarchical learning can be incrementally introduced for individual subsystems before being applied to the whole enterprise.

Improvements in energy efficiency have profound implications for sustainable IoT development. Bandwidth savings also have profound implications for the successful deployment of IoT in bandwidth-limited environments such as rural settings, developing countries, and aerial or maritime settings. Energy efficiency and bandwidth savings can significantly reduce costs of operation while increasing the applicability of sustainable IoT solutions.

6.3 Future Research Directions

Several promising research avenues can be identified from this work. Cross-silo federated learning is a challenging area when it comes to trust establishment and incentive alignment, which are domains where blockchain technology can be applied for purposes other than security. The ability to support lifelong learning is another important frontier that enables IoT systems to adapt to changing environments without forgetting what they have learned before.

Integration with quantum-resistant cryptography[11] offers a visionary security improvement, especially for IoT networks with longer lifecycles. In a similar manner, integration with neuromorphic computing may offer more power-efficient edge learning by translating neural networks from software to hardware. Lastly, standardization for hierarchical federated learning protocols would offer better interoperability between various vendor ecosystems, thus speeding up innovation in this increasingly important area.

References

1. L. Da Xu, W. He, and S. Li, “Internet of Things in industries: A survey,” IEEE Internet of Things Journal, vol. 1, no. 4, pp. 223–233, Aug. 2014, doi: 10.1109/JIOT.2014.2300753.
2. S. Li, L. D. Xu, and S. Zhao, “The Internet of Things: A survey,” Information Systems Frontiers, vol. 17, no. 2, pp. 243–259, Apr. 2015, doi: 10.1007/s10796-014-9492-7.
3. Q. Yang, Y. Liu, T. Chen, and Y. Tong, “Federated machine learning: Concept and applications,” ACM Transactions on Intelligent Systems and Technology, vol. 10, no. 2, pp. 1–19, Jan. 2019, doi: 10.1145/3298981.
4. M. Ammad-Ud-Din et al., “Federated learning for wireless communications: Motivation, opportunities, and challenges,” IEEE Communications Magazine, vol. 58, no. 6, pp. 46–51, Jun. 2020, doi: 10.1109/MCOM.001.1900618.
5. K. Bonawitz et al., “Practical secure aggregation for privacy-preserving machine learning,” in Proc. ACM CCS, Toronto, ON, Canada, 2017, pp. 1175–1191, doi: 10.1145/3133956.3133982.
6. M. Dorri, S. S. Kanhere, and R. Jurdak, “Blockchain in Internet of Things: Challenges and solutions,” IEEE Communications Surveys & Tutorials, vol. 21, no. 2, pp. 1463–1492, 2019, doi: 10.1109/COMST.2018.2849663.
7. K. Zhang, Y. Zhu, S. Leng, Y. He, S. Maharjan, and Y. Zhang, “Deep learning empowered task offloading for mobile edge computing in IoT,” IEEE Internet of Things Journal, vol. 6, no. 3, pp. 4242–4251, Jun. 2019, doi: 10.1109/JIOT.2018.2878698.
8. N. H. Tran et al., “Federated learning over wireless networks: Optimization model design and analysis,” in Proc. IEEE INFOCOM, Paris, France, 2019, pp. 1387–1395, doi: 10.1109/INFOCOM.2019.8737464.
9. A. Ghosh, D. Das, and S. Chatterjee, “AIoT: Artificial intelligence meets Internet of Things,” IEEE Computer, vol. 53, no. 6, pp. 56–64, Jun. 2020, doi: 10.1109/MC.2020.2982473.
10. C. Dwork and A. Roth, “The algorithmic foundations of differential privacy,” Foundations and Trends in Theoretical Computer Science, vol. 9, no. 3–4, pp. 211–407, 2014, doi: 10.1561/0400000042.

11. D. Mathur, S. Mathur, and A. Rai, "The Collaboration of the Trio of Quantum Computing, IoT and AI Powered by Cloud Architecture," *International Journal of Enhanced Research in Science, Technology & Engineering*, vol. 14, no. 5, May 2025.
12. Shawkat, M., El-desoky, A., Ali, Z.H. et al. Blockchain and federated learning based on aggregation techniques for industrial IoT: A contemporary survey. Peer-to-Peer Netw. Appl. 18, 192 (2025). https://doi.org/10.1007/s12083-025-01991-0
13. Chintala, Sarada & Vasavi, M. & Ambika, K.. (2024). Securing IoT Devices from DDoS Attacks through Blockchain and Multi-Code Trust Framework. E3S Web of Conferences. 472. 10.1051/e3sconf/202447203001.
14. Vashisth, S., & Goyal, A. (2025). A Survey of Federated Learning for IoT: Addressing Resource Constraints and Heterogeneous Challenges. Informatica, 49(17). https://doi.org/10.31449/inf.v49i17.7707
15. P. Savvidis and G. Papakostas, "Edge computing for computer vision in IoT: Feasibility and directions," EAI Endorsed Transactions on Internet of Things, vol. 11, 2025, doi: 10.4108/eetiot.2025.123456.