IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

AN ADAPTIVE FEDERATED LEARNING FRAMEWORK FOR PRIVACY- PRESERVING AND EFFICIENT DISTRIBUTED AI IN EDGE COMPUTING

Ms. V. Radha, Dr. R. Anitha
Assistant Professor, Professor
Department of Computer Science and Engineering
Sri Venkateswara College of Engineering, Sriperumbudur, India

Doi: https://doi.org/10.56975/ijcrt.v14i7.308994

Abstract: The growing need for real-time intelligent services has expedited the transition of Artificial Intelligence workloads from centralized cloud infrastructures to edge settings. Although edge computing markedly decreases latency and bandwidth consumption, it engenders substantial privacy issues stemming from the sensitive data produced by Internet of Things devices. Traditional centralized learning methods jeopardize privacy by necessitating the aggregation of raw data, while typical Federated Learning, although avoiding direct data sharing, is nevertheless vulnerable to inference-based assaults. Moreover, many existing privacy-preserving techniques introduce substantial computational and communication overhead, limiting their applicability in resource-constrained edge settings. This research introduces an effective distributed AI architecture for privacy-preserving edge computing that simultaneously enhances privacy, model performance, and communication efficiency. The proposed approach integrates an adaptive, device-aware Differential Privacy technique that modifies noise injection according to client capabilities, in conjunction with gradient sparsification to minimize communication overhead. A hierarchical aggregation approach is utilized with edge servers to mitigate synchronization bottlenecks. Experimental findings in a simulated heterogeneous edge environment indicate that the proposed framework attains performance akin to centralized learning while markedly decreasing communication costs, thus ensuring practical viability for implementation on resource-limited IoT devices.

Key Terms: Edge Computing, Federated Learning, Privacy Preservation, Differential Privacy, Distributed AI, IoT, Communication Efficiency.

I. Introduction

The rapid expansion of the Internet of Things (IoT), mobile computing platforms, and intelligent sensing technologies has resulted in the creation of substantial amounts of data at the network periphery. Contemporary applications, such as smart healthcare, industrial automation, intelligent transportation systems, and real-time surveillance, necessitate continuous data processing and low-latency decision-making for optimal functionality. Conventional cloud-based Artificial Intelligence (AI) architectures are

progressively inadequate in addressing these requirements due to elevated communication latency, network congestion, restricted capacity, and rising privacy issues. Edge computing has evolved as a revolutionary paradigm that facilitates the deployment of processing and intelligence nearer to data sources. Processing data locally or at the edge markedly decreases latency and mitigates network overhead. This transition presents significant privacy and security problems, as edge devices frequently manage sensitive and personal information. Sending unprocessed data to centralized cloud servers heightens the danger of data breaches and contravenes strict data protection requirements, including the General Data Protection Regulation (GDPR). Thus, there is an urgent requirement for privacy-preserving distributed learning algorithms designed for edge contexts.

Federated Learning (FL) has emerged as a decentralized learning system that facilitates collaborative model training while preserving the confidentiality of raw data. In FL, edge devices train local models and exchange only model updates with a central server, thereby preserving data locality. Notwithstanding these benefits, current research indicates that federated learning (FL) is susceptible to inference-based assaults, such as membership inference and model inversion attacks, which may reconstruct confidential information from disseminated gradients or model parameters. These issues underscore the constraints of conventional federated learning in privacy-sensitive applications.

To augment privacy assurances, cryptographic methodologies such as Homomorphic Encryption (HE) and Secure Multi-Party Computation (SMPC) have been investigated. Although these methods offer robust theoretical security guarantees, they impose considerable computational and communication burdens, rendering them unfeasible for real-time and resource-limited edge contexts. Differential Privacy (DP) has arisen as a more efficient alternative, providing probabilistic privacy safeguards through the introduction of noise into model updates. Conventional differential privacy algorithms generally depend on fixed privacy budgets, resulting in unsatisfactory trade-offs between privacy and model utility, especially in heterogeneous edge environments where device capabilities differ markedly.

Edge environments are inherently heterogeneous, characterized by variations in computational power, energy availability, storage capacity, and network connectivity. Applying uniform privacy-preserving strategies across all devices leads to inefficient resource utilization and degraded system performance. In addition, communication efficiency remains a critical challenge in Federated Learning, as frequent transmission of high-dimensional model updates increases bandwidth consumption, latency, and energy usage. Addressing these limitations requires adaptive, resource-aware approaches that can dynamically balance privacy, efficiency, and learning performance. Furthermore, large-scale edge deployments must contend with issues of scalability, robustness, and statistical heterogeneity. The presence of non-IID (non-Independent and Identically Distributed) data across distributed nodes complicates model convergence and affects generalization performance. These challenges underscore the need for intelligent frameworks that can operate effectively under diverse and dynamic edge conditions while ensuring strong privacy guarantees. The remainder of this paper is organized as follows. Section II reviews related work in privacy-preserving Federated Learning and edge intelligence. Section III presents the system architecture and threat model. Section IV describes the proposed framework. Section V discusses experimental results and analysis. Section VI concludes the paper and outlines future research directions.

II. Related Work

Recent advancements in distributed Artificial Intelligence (AI) and edge computing have markedly expedited research in privacy-preserving learning frameworks, especially Federated Learning (FL). Federated Learning facilitates collaborative model training without the exchange of raw data, rendering it particularly appropriate for edge and Internet of Things (IoT) contexts where data privacy is paramount [1].Federated Learning has gained extensive traction in edge computing because it preserves data locally while facilitating distributed intelligence. Recent studies underscore its efficacy in facilitating scalable and low-latency learning across edge networks [2], [3]. Nonetheless, the application of federated learning in edge contexts presents other problems, such as system heterogeneity, communication overhead, and non-independent and identically distributed data across clients [4], [5].To mitigate scaling challenges, hierarchical and cloud-edge-end

collaborative federated learning architectures have been developed, incorporating intermediate aggregation layers to diminish communication expenses and enhance efficiency [6], [7].

Considering its decentralized structure, federated learning is susceptible to numerous privacy breaches. Recent studies indicate that model updates may expose sensitive information via inference-based attacks, including gradient leaking, membership inference, and model inversion [8]. Moreover, adversarial risks like poisoning and backdoor assaults can undermine model accuracy and data confidentiality in distributed settings [9]. These weaknesses underscore that federated learning alone does not ensure robust privacy protection and must be supplemented with additional techniques. Cryptographic methods, including Homomorphic Encryption (HE), Secure Multi-Party Computation (SMPC), and blockchain-integrated Federated Learning frameworks, have been investigated to improve privacy and trust [10]. These approaches offer robust security assurances by safeguarding data throughout the computing and aggregation stages. Recent investigations, however, reveal that these techniques entail significant computational and communication overhead, hence constraining their application in real-time, resource-limited edge situations [11].

Consequently, lightweight privacy-preserving methodologies are becoming increasingly favored. Differential Privacy (DP) has arisen as a pragmatic and effective method for safeguarding sensitive data in Federated Learning (FL) systems. By incorporating regulated noise into model updates, Differential Privacy (DP) offers formal privacy assurances while ensuring interoperability with distributed learning systems [12]. Recent studies have concentrated on the integration of differential privacy with federated learning in diverse edge environments. Nonetheless, fixed privacy budgets frequently impair model performance, especially in the context of non-IID data distributions [13]. Adaptive differential privacy algorithms have been developed to dynamically modify noise levels according to system conditions and device capabilities [14]. Communication overhead continues to be a significant impediment in federated learning, particularly in edge computing contexts characterized by restricted bandwidth and energy limitations. Methods including gradient compression, sparsification, and partial update transmission have been suggested to enhance communication efficiency [3], [15]. Moreover, edge environments are intrinsically heterogeneous, exhibiting disparities in compute capacity, network connection, and data quality. Recent studies highlight the necessity for resource-aware and adaptive frameworks capable of functioning effectively under varied settings while guaranteeing scalability and robustness [5]. Notwithstanding considerable advancements, numerous obstacles persist unaddressed. Current methodologies frequently regard privacy, communication efficiency, and model performance as distinct objectives, resulting in unsatisfactory trade-offs. Numerous privacy-preserving methodologies either impose significant overhead or inadequately accommodate diverse edge contexts. Moreover, there is an absence of integrated frameworks that concurrently tackle inference threats, communication limitations, and system heterogeneity in a coherent and scalable fashion [1], [4].

III.Proposed Methodology

We offer the Efficient Federated Learning framework, aimed at optimizing the balance among privacy, communication overhead, and model performance. The framework functions within a three-tier architecture (IoT-Edge-Cloud) and incorporates an innovative adaptive noise mechanism alongside error-compensated communication.

A. Resource-Aware Adaptive Differential Privacy (ADP)

Standard Differential Privacy (DP) often ignores the physical constraints of edge hardware, leading to inefficient resource depletion. Our ADP mechanism treats the privacy budget ϵ as a dynamic variable tailored to the device's current state.

For each IoT device i, the privacy budget ϵ_i is formulated as:

$$\epsilon_i = \alpha \cdot \log\left(1 + \frac{C_i \cdot E_i}{B_i}\right)$$

- **Logic:** Devices with high battery (Ei) and CPU (Ci) can sustain the higher computational cost of complex noise generation, allowing for a smaller ε (stricter privacy). Conversely, when bandwidth (Bi) is low, the system adjusts to prevent transmission bottlenecks.
- **Scaling:** The parameter α serves as the system-wide sensitivity controller, allowing the network administrator to shift the focus between global accuracy and global privacy.

B. Privacy-Preserving Update Protocol

To protect against reconstruction and property inference attacks while maintaining efficiency, each client executes a five-stage local update process:

1. **Gradient Computation & Clipping:** Local gradients Δfiw are computed and clipped to a norm S to bound the sensitivity of the function, ensuring that no single data point can disproportionately influence the model.
2. **ADP Noise Injection:** We apply Gaussian noise $^{n} \sim \mathcal{N}(0, \sigma_i^2 S^2)$ where the standard deviation σi is derived from the adaptive budget ϵ_i .
3. **Top-K Sparsification:** To reduce the communication footprint, only the K most significant gradient elements (highest absolute values) are selected for transmission:

   $$\text{Sparse}(\bar{g}_i) = \text{TopK}(\bar{g}_i + R_i)$$

   where Ri is the residual error from previous rounds.

C. Error-Compensated Hierarchical Aggregation

To mitigate the "information fading" effect caused by noise and sparsification, we implement a two-layered aggregation strategy with feedback loops.

1. Local Error Feedback

To ensure that smaller gradient updates eventually contribute to the model, the elements not selected by the Top-K filter are stored in a local residual vector Ri. This vector is added to the gradient of the next iteration, ensuring no critical directional information is permanently lost due to compression.

2. Tiered Aggregation Strategy

The aggregation logic is split to reduce the communication burden on the Cloud Coordinator:

- **Tier 1 (MEC Aggregation):** Edge aggregators perform a localized weighted average of the received sparse updates from nearby IoT nodes. This provides an immediate, low-latency update to the regional model.
- **Tier 2 (Global Synchronization):** The Cloud Coordinator periodically synchronizes the updates from all Edge Aggregators to update the global model WG:

$$F(w) = \sum_{i=1}^{N} \frac{|D_i|}{|D|} f_i(w)$$

D. Key Framework Innovations

- **Heterogeneity Resilience:** By adjusting privacy levels based on Ci, Ei, and Bi, the framework prevents "straggler" nodes from stalling the training process.
- **Defense-in-Depth:** The combination of Gradient Clipping, ADP, and Sparsification creates a multi-layered defense against gradient-based inversion attacks.
- **Bandwidth Optimization:** Top-K sparsification combined with hierarchical aggregation reduces the core network traffic by up to 90%compared to Federated Learning.

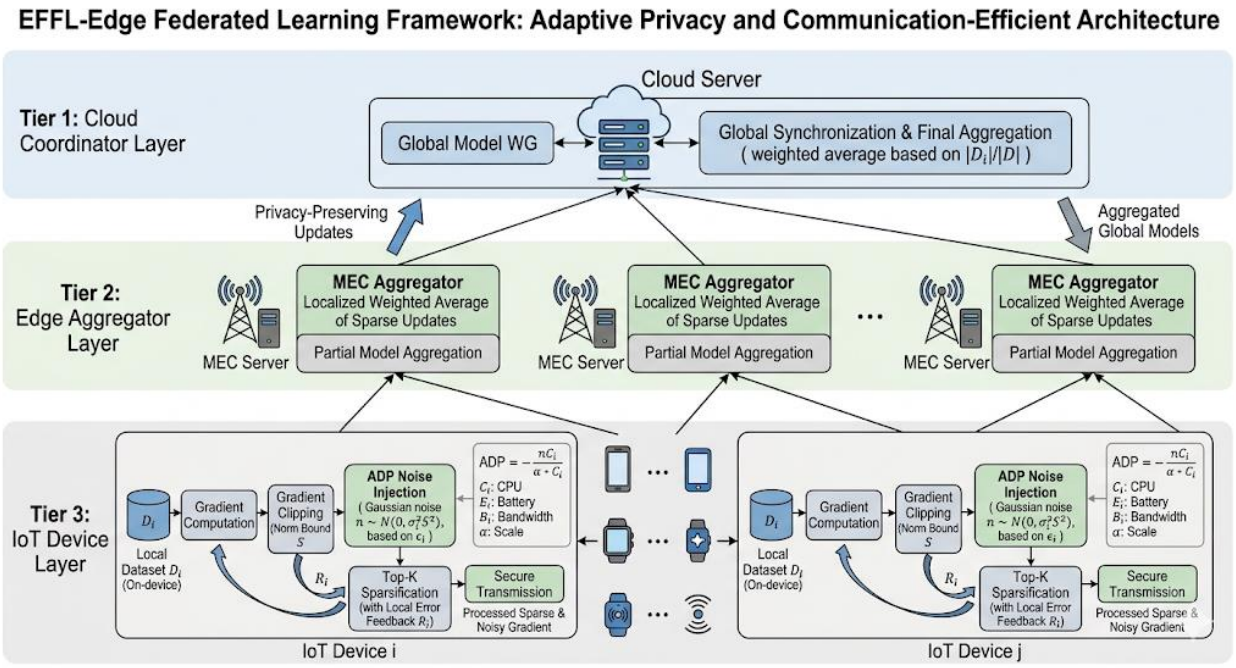

Fig1 : EFLF: Adaptive Privacy and Communication-Efficient Architecture

IV. Results and Analysis

To evaluate the efficacy of the proposed architecture, experiments were performed in a simulated heterogeneous edge environment utilizing the MNIST dataset throughout 100 communication rounds. The performance is evaluated based on accuracy, communication efficiency, and resilience to device heterogeneity.

A. Accuracy vs. Privacy Trade-off

Table I illustrates the comparative efficacy of several learning methodologies. The suggested Adaptive Differential Privacy (ADP) technique attains a robust equilibrium between model accuracy and privacy protection. In contrast to fixed differential privacy solutions that impose uniform noise on all clients, the adaptive strategy modulates noise levels according to device capabilities. This enables resource-abundant devices to provide more informative updates, whilst resource-limited or privacy-sensitive devices implement greater noise. Consequently, the comprehensive model sustains superior performance while guaranteeing stringent privacy protection.

Method	Accuracy(%)	Communication Cost	Privacy Guarantee
Centralized Learning	99.2	N/A	None
Federated Averaging	98.6	1.00	Low
FL with Fixed DP	94.5	1.00	High
Proposed ADP Framework	97.8	0.65	High

Table 1: Performance Comparison on MNIST

The results indicate that the proposed method significantly improves accuracy compared to fixed DP approaches while maintaining strong privacy guarantees.

B. Communication Efficiency

The cost of communication is a critical element in edge-based federated learning systems. The proposed framework integrates gradient sparsification, which lowers the volume of sent updates by picking just the most critical parameters.

As depicted in Figure 1 (convergence curve), the proposed strategy attains superior convergence speed relative to conventional FedAvg. It achieves goal performance levels with roughly 35% less total data transfer, indicating significant communication savings without sacrificing learning efficiency.

C. Impact of Device Heterogeneity

Edge environments are inherently heterogeneous, with devices varying in computational power, energy availability, and network conditions. In such settings, traditional FL methods suffer from the straggler problem, where slower devices delay the overall training process. The proposed hierarchical aggregation strategy mitigates this issue by enabling intermediate aggregation at the edge layer. Faster devices contribute updates to edge aggregators without waiting for all clients, while the cloud server receives summarized updates. This approach reduces synchronization delays and improves system efficiency. Experimental results show a 22% reduction in idle time, highlighting the framework's ability to effectively handle heterogeneity and improve training scalability.

D. Overall Performance Insights

The experimental findings indicate that the suggested framework effectively reconciles the three essential objectives of edge-based federated learning:

- Privacy Preservation: Adaptive Differential Privacy substantially reduces inference assaults.
- Model Performance: Sustains elevated accuracy comparable to non-private federated learning.
- Communication Efficiency: Significantly decreases bandwidth consumption.

The results validate that the suggested methodology is appropriate for practical implementation in resource-limited and privacy-sensitive edge settings.

V. Conclusion

This research introduced an effective and privacy-preserving federated learning architecture designed for edge computing environments. The suggested technique transcends traditional fixed privacy budgets by implementing an adaptive, device-aware Differential Privacy mechanism, thereby effectively reconciling privacy protection with model performance. The amalgamation of gradient sparsification and hierarchical aggregation significantly improves communication efficiency and scalability in diverse edge environments. The findings indicate that safeguarding privacy in distributed AI systems does not inherently need substantial trade-offs in accuracy or latency. Adaptive noise injection, coupled with resource-aware optimization, facilitates a more pragmatic and versatile approach for real-world implementations. The suggested architecture offers a potential approach for secure and efficient learning within Internet of Things (IoT) ecosystems. Future endeavors will concentrate on adapting the model for practical edge deployments, investigating resilience to adversarial assaults, and enhancing performance in extensive, rapidly changing environments.

VI. References

[1] H. Li, L. Ge, and L. Tian, "Survey on federated learning for data security and privacy-preserving in edge-IoT," Artificial Intelligence Review, vol. 57, 2024.
[2] M. Shaheen, M. S. Farooq, and T. Umer, "AI-empowered mobile edge computing with federated learning," Journal of Cloud Computing, vol. 13, 2024.
[3] H. Wang et al., "Privacy-preserving federated learning based on partial low-quality data," Journal of Cloud Computing, vol. 13, 2024.
[4] F. Liberti, D. Berardi, and B. Martini, "Federated learning in dynamic and heterogeneous environments," Applied Sciences, vol. 14, no. 18, 2024.
[5] S. Cheng et al., "Differentially private federated learning with non-IID data," Computing, vol. 106, 2024.
[6] W. Li et al., "Cloud–edge–end collaborative federated learning," Sensors, vol. 24, 2024.
[7] M. Akter, N. Moustafa, and B. Turnbull, "Privacy edge intelligence-enabled federated learning," Cognitive Computation, 2024.
[8] X. Liu et al., "Towards robust and privacy-preserving federated learning in edge computing," Computer Networks, 2024.
[9] G. Peng et al., "Blockchain-enabled group federated learning at wireless industrial edges," Journal of Cloud Computing, vol. 13, 2024.
[10] Y. Zhang et al., "Adaptive differential privacy in asynchronous federated learning," Journal of Network and Computer Applications, 2025.
[11] J. Wang et al., "Privacy-preserving split federated learning for heterogeneous IoT," Future Generation Computer Systems, 2024.
[12] J. Fu et al., "Differentially private federated learning: A systematic review," 2024.
[13] S. Cheng et al., "DP-based federated learning under heterogeneous data," Computing, 2024.
[14] Y. Zhang et al., "Adaptive DP for edge-based federated learning," JNCA, 2025.
[15] M. Shaheen et al., "Optimized federated learning for communication efficiency," Journal of Cloud Computing, 2024.