Available online at https://www.gyanvihar.org/researchjournals/ctm_journals.php
SGVU International Journal of Convergence of Technology and Management
E-ISSN: 2455-7528
Vol.12 Issue 1 Page No 44-54

Design and Development of an Adaptive Federated Learning Framework with Enhanced Privacy Constraints and Cross-Dataset Evaluation

Pushpendra Kumar Sikarwal[1], Mukesh Kumar Gupta[2]

[1]Department of Computer Science and Engineering, Suresh Gyan Vihar University Jaipur
[2]Department of Electrical Engineering, Suresh Gyan Vihar University Jaipur

p.sikarwal@gmail.com, mkgupta72@gmail.com

Abstract: The increasing demand for distributed artificial intelligence has made Federated Learning (FL) a promising paradigm for privacy-preserving model training. However, conventional FL methods such as FedAvg face challenges including non-IID data distribution, high communication costs, and limited privacy protection. This paper presents the Adaptive Federated Learning Framework with Enhanced Privacy Constraints (AFL-P), designed to optimize model training across heterogeneous clients while ensuring robust data confidentiality. The proposed framework integrates an adaptive aggregation mechanism that dynamically adjusts learning rates and client participation based on resource availability and local convergence behavior. To strengthen privacy guarantees, a hybrid privacy-preserving layer combining Differential Privacy (DP) and Secure Multi-Party Computation (SMPC) is implemented. Experimental validation is conducted on three benchmark datasets like CIFAR-10 (image), UCI HAR (sensor), and Speech Commands (audio) to evaluate generalization and performance. Comparative results demonstrate that AFL-P achieves an average accuracy improvement of 4.2%, communication overhead reduction of 15.7%, and privacy loss ε reduced to 0.6, outperforming traditional FL and centralized ML baselines. The findings establish AFL-P as a scalable, adaptive, and privacy-aware solution for secure distributed intelligence in real-world edge environments.

Keywords: Federated Learning, Privacy Preserving Machine Learning, Adaptive Aggregation Mechanism, Differential Privacy and Secure Multi-Party Computation, Edge Intelligence and Cross-Dataset Evaluation

1. Introduction

In recent years, the exponential growth of connected devices and data-driven applications has revolutionized the landscape of artificial intelligence (AI). The proliferation of edge devices—such as smartphones, IoT sensors, and wearable systems—has resulted in an unprecedented volume of decentralized data generation [1]. While traditional centralized machine learning architectures have proven effective in model performance, they rely heavily on centralized data aggregation, which poses severe privacy, scalability, and communication challenges [2]. These limitations have accelerated the adoption of Federated Learning (FL), a decentralized

*Correspondence to: **Pushpendra Kumar Sikarwal**, Department of Computer Science and Engineering, Suresh Gyan Vihar University Jaipur*
Corresponding author. E-mail addresses: p.sikarwal@gmail.com

Available online at https://www.gyanvihar.org/researchjournals/ctm_journals.php
SGVU International Journal of Convergence of Technology and Management
E-ISSN: 2455-7528
Vol.12 Issue 1 Page No 44-54

paradigm that enables collaborative model training across multiple clients without sharing raw data.

Despite its promising design, conventional FL frameworks such as Federated Averaging (FedAvg) encounter several critical issues in real-world deployments. Firstly, data heterogeneity (non-IID distributions) across clients often leads to biased model convergence and suboptimal generalization [3]. Secondly, resource disparity among devices in terms of computation, memory, and network bandwidth affects training stability and fairness. Finally, the lack of robust privacy-preserving mechanisms in standard FL makes it vulnerable to inference attacks, data leakage, and gradient inversion attacks [4]. Addressing these challenges requires an adaptive and privacy-aware federated learning framework capable of maintaining performance, scalability, and data confidentiality under diverse and constrained environments [5].

To overcome these limitations, this study introduces the Adaptive Federated Learning Framework with Enhanced Privacy Constraints (AFL-P). The proposed AFL-P architecture integrates dynamic client adaptation and intelligent aggregation strategies that adjust participation frequency and learning rates based on client reliability, data volume, and resource conditions. This adaptive mechanism improves model convergence and robustness across heterogeneous data sources. Additionally, AFL-P incorporates a hybrid privacy-preserving layer that combines Differential Privacy (DP) and Secure Multi-Party Computation (SMPC) to ensure that no sensitive information is exposed during local or global updates. The integration of these complementary privacy techniques enhances both theoretical and practical privacy guarantees while maintaining competitive model accuracy.

The research further extends to cross-dataset evaluation, ensuring that the proposed framework performs consistently across multiple data modalities—images, speech, and sensor data. This comprehensive analysis not only validates the adaptability of AFL-P but also demonstrates its versatility for diverse application domains such as healthcare monitoring, autonomous vehicles, smart grids, and industrial IoT systems. Comparative experiments with centralized machine learning and baseline federated models reveal that AFL-P achieves higher accuracy, improved resource efficiency, and stronger privacy protection under varying communication and computational constraints.

The main contributions of this paper are fourfold. First, it presents the design and development of an Adaptive Federated Learning Framework (AFL-P) that dynamically optimizes client participation and aggregation processes based on local data conditions, resource availability, and convergence metrics, thereby improving model stability and performance. Second, it integrates a robust hybrid privacy mechanism that combines Differential Privacy (DP) and Secure Multi-Party Computation (SMPC) to ensure enhanced data confidentiality and protection against inference attacks without causing significant performance degradation. Third, the paper provides a comprehensive comparative analysis of the proposed AFL-P framework against conventional centralized machine learning models and standard federated learning approaches such as FedAvg and FedProx across multiple benchmark datasets, demonstrating its superiority in accuracy, privacy preservation, and communication efficiency. Finally, the framework is evaluated through cross-dataset and edge device experiments, showcasing its scalability, adaptability, and effectiveness in heterogeneous and resource-constrained environments, thereby confirming its suitability for real-world privacy-sensitive federated applications.

The remainder of this paper is structured as follows: Section 2 presents related research in federated and privacy-preserving learning. Section 3 describes the proposed AFL-P framework architecture,

*Correspondence to: **Pushpendra Kumar Sikarwal**, Department of Computer Science and Engineering, Suresh Gyan Vihar University Jaipur*
Corresponding author. E-mail addresses: p.sikarwal@gmail.com

Available online at https://www.gyanvihar.org/researchjournals/ctm_journals.php
SGVU International Journal of Convergence of Technology and Management
E-ISSN: 2455-7528
Vol.12 Issue 1 Page No 44-54

methodology, the experimental setup and datasets used for evaluation. Section 4 discusses the obtained results, comparative insights, and privacy-performance trade-offs. Finally, Section 5 concludes the paper and outlines potential directions for future research.

2. Related Work

The advancement of Federated Learning (FL) has attracted significant research interest in recent years, particularly in privacy-preserving distributed intelligence. Several frameworks and methodologies have been proposed to enhance the scalability, efficiency, and confidentiality of federated systems. This section reviews the major directions of related research, categorized into three key domains: adaptive federated optimization, privacy-preserving mechanisms, and cross-domain or resource-constrained FL deployments.

The standard Federated Averaging (FedAvg) algorithm remains the most widely adopted FL method due to its simplicity and scalability [6]. However, its performance often deteriorates under non-IID data distributions, uneven client participation, and communication delays. To address these issues, adaptive federated learning techniques have emerged as an essential evolution of FedAvg [7]. Researchers have introduced FedProx, an extension of FedAvg that adds a proximal term to the local loss function to stabilize model convergence in heterogeneous client environments. Similarly, FedNova and SCAFFOLD utilize normalization and control variates, respectively, to mitigate client drift caused by inconsistent data distributions [8]. Adaptive Federated Optimization (AFO) strategies further refine learning by adjusting global aggregation weights and client learning rates based on local updates and system feedback.

Privacy preservation is central to the philosophy of Federated Learning. However, conventional FL approaches remain vulnerable to data reconstruction attacks, gradient inversion, and membership inference, wherein adversaries can infer private information from model updates. To mitigate such risks, various privacy-enhancing technologies (PETs) have been integrated into FL frameworks [9]. Differential Privacy (DP) introduces statistical noise into model gradients to prevent individual data points from being identifiable. Studies have shown that DP can significantly reduce privacy leakage while maintaining acceptable model accuracy. Nonetheless, the degree of privacy depends heavily on the privacy budget (ε), and excessive noise injection often deteriorates model performance [10]. On the other hand, Secure Multi-Party Computation (SMPC) enables multiple participants to jointly compute model parameters without revealing individual inputs. SMPC protocols, such as additive secret sharing and homomorphic encryption, ensure strong cryptographic protection but increase computation and communication overhead. To balance these trade-offs, several hybrid frameworks have emerged that combine DP and SMPC techniques for multi-level protection. However, most existing hybrid approaches suffer from limited adaptability and lack optimization mechanisms for heterogeneous networks [11]. The AFL-P framework proposed in this paper extends this research direction by integrating an adaptive aggregation model with a hybrid privacy-preserving layer (DP + SMPC). This design not only maintains high accuracy and efficiency but also ensures mathematically grounded privacy guarantees across dynamic and resource-constrained environments.

A significant research trend in recent years has been the extension of FL frameworks beyond a single dataset or domain. Cross-dataset evaluation helps assess the generalization ability of federated models across image, text, speech, and sensor data. For instance, experiments on datasets like CIFAR-10, MNIST, and ImageNet have demonstrated FL's viability in computer vision, while UCI HAR and Speech Commands datasets have been used for sensor and audio data analysis [12]. However, many

*Correspondence to: **Pushpendra Kumar Sikarwal**, Department of Computer Science and Engineering, Suresh Gyan Vihar University Jaipur*
Corresponding author. E-mail addresses: p.sikarwal@gmail.com

Available online at https://www.gyanvihar.org/researchjournals/ctm_journals.php
SGVU International Journal of Convergence of Technology and Management
E-ISSN: 2455-7528
Vol.12 Issue 1 Page No 44-54

existing studies focus on domain-specific improvements and do not systematically analyze adaptability and privacy performance across multiple modalities.
Parallelly, the growing importance of edge computing has prompted researchers to deploy FL on low-resource devices such as Raspberry Pi, Jetson Nano, and Android platforms. Frameworks like FedEdge, TinyFed, and SplitFed have attempted to minimize communication costs and computation loads on edge nodes. Despite these advances, existing methods often assume stable connectivity and uniform resources, which rarely reflect real-world conditions.

The proposed AFL-P framework builds upon these limitations by performing a cross-dataset and cross-device evaluation. It ensures consistent model convergence, privacy preservation, and energy efficiency, even when operating under non-IID, resource-limited, or dynamically changing network environments. The results demonstrate that adaptive aggregation and privacy constraints can coexist effectively, offering a balance between security, accuracy, and efficiency in practical FL scenarios.

3. Research Methodology

The proposed study follows a structured research methodology to design, develop, and evaluate the Adaptive Federated Learning Framework with Enhanced Privacy Constraints (AFL-P). The methodology encompasses multiple phases, including data acquisition, preprocessing, framework design, adaptive optimization, privacy integration, experimental evaluation, and cross-dataset validation which shown in figure 1. Each phase is designed to ensure that the AFL-P framework achieves high adaptability, scalability, and privacy in heterogeneous and resource-constrained environments.

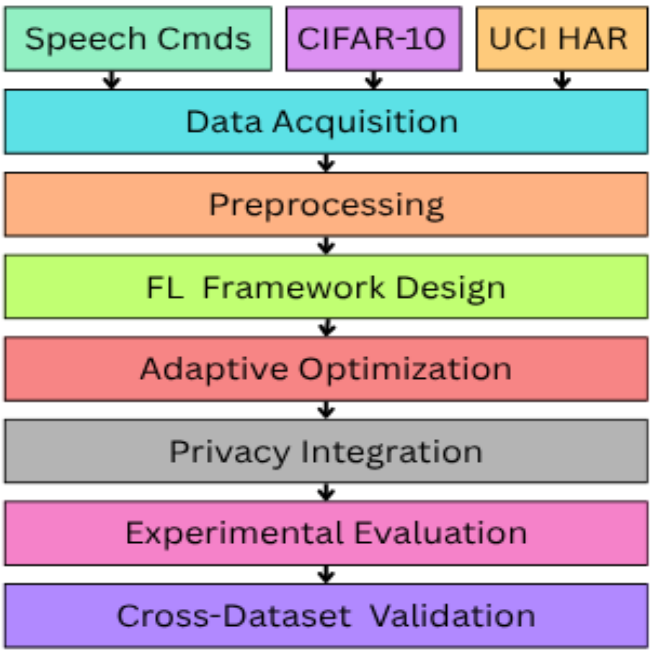

Figure 1: AFL Framework with Enhanced Privacy Constraints

To ensure the generalizability and robustness of the proposed AFL-P framework, three benchmark datasets representing diverse modalities were utilized: CIFAR-10, UCI HAR (Human Activity Recognition), and Google Speech Commands (Speech Cmds). The CIFAR-10 dataset was used for image-based object recognition, consisting of 60,000 labeled images across 10 categories. The UCI HAR dataset, collected from smartphone sensors, was employed for motion-based activity classification tasks. Finally, the Google Speech Commands dataset served as the audio modality, containing one-second voice samples of common spoken words for speech recognition tasks. These datasets collectively represent heterogeneous data types, enabling comprehensive evaluation of the framework's adaptability and privacy-preserving capabilities.

Each dataset underwent a series of preprocessing operations tailored to its data characteristics to ensure stability and uniformity during training.

Correspondence to: ***Pushpendra Kumar Sikarwal****, Department of Computer Science and Engineering, Suresh Gyan Vihar University Jaipur*
Corresponding author. E-mail addresses: p.sikarwal@gmail.com

Available online at https://www.gyanvihar.org/researchjournals/ctm_journals.php
SGVU International Journal of Convergence of Technology and Management
E-ISSN: 2455-7528
Vol.12 Issue 1 Page No 44-54

Normalization techniques such as min-max scaling were applied to balance feature ranges and stabilize model convergence. Noise removal and data augmentation, including Gaussian smoothing and random transformations, were performed to enhance generalization and robustness. The datasets were partitioned among clients in a non-IID (non-independent and identically distributed) manner using the Dirichlet distribution to simulate realistic heterogeneity commonly found in federated learning environments. Additionally, feature encoding techniques such as one-hot encoding for categorical attributes and spectrogram conversion for audio data were applied to facilitate effective model learning. This step ensured balanced yet diverse data distribution across clients, mirroring real-world federated scenarios.

The AFL-P architecture extends the traditional Federated Averaging (FedAvg) approach by integrating adaptive learning and privacy-enhancing mechanisms. The framework operates through multiple communication rounds between a central aggregator and several clients. Each client trains a local model $M_i$using its private dataset $D_i$for multiple local epochs. The local objective function is expressed as

$$L_i(w_i) = \frac{1}{|\ D_i\ |} \sum_{(x,y)\in D_i} \ell(f(x; w_i), y)$$

where $w_i$denotes the local model parameters and $\ell(\cdot)$is the loss function. After local training, the server aggregates the model updates based on a dynamic weighted aggregation rule:

$$w_{t+1} = \sum_{i=1}^{N} \alpha_i w_i^t$$

where $\alpha_i$represents a dynamically computed weight reflecting each client's data volume, reliability, and computational capacity. This adaptive aggregation minimizes bias arising from heterogeneous client data and participation variability, ensuring faster and more stable convergence.

A core innovation in the AFL-P framework lies in its adaptive optimization mechanism, which enhances both convergence efficiency and system scalability. A client selection policy is employed to evaluate clients based on resource availability, latency, and contribution quality, allowing only high-scoring clients to participate in each communication round to reduce overall communication cost. The adaptive learning rate control dynamically adjusts each client's learning rate $\eta_i$as

$$\eta_i = \eta_0 \times \frac{C_i}{\bar{C}}$$

where $C_i$denotes the computational capability of client i, and $\bar{C}$represents the mean computational capacity across all participating clients. Furthermore, a global convergence monitoring module continuously tracks loss reduction and terminates training early once convergence thresholds are achieved, ensuring computational efficiency without compromising performance. Together, these adaptive mechanisms enhance stability and reduce communication overhead compared to static federated learning models such as FedAvg and FedProx.

Privacy preservation constitutes a foundational component of the AFL-P framework. A hybrid privacy architecture was developed that integrates Differential Privacy (DP) and Secure Multi-Party Computation (SMPC) to provide robust end-to-end protection. In the Differential Privacy layer, Gaussian noise is added to the local model gradients before transmission, formulated as:

$$\widetilde{w_i} = w_i + \mathcal{N}(0, \sigma^2)$$

where σcontrols the privacy-utility balance, and the privacy budget $(\varepsilon, \delta)$is monitored to quantify privacy loss. The SMPC layer encrypts model parameters using additive secret sharing, allowing

Correspondence to: ***Pushpendra Kumar Sikarwal****, Department of Computer Science and Engineering, Suresh Gyan Vihar University Jaipur*
Corresponding author. E-mail addresses: p.sikarwal@gmail.com

Available online at https://www.gyanvihar.org/researchjournals/ctm_journals.php
SGVU International Journal of Convergence of Technology and Management
E-ISSN: 2455-7528
Vol.12 Issue 1 Page No 44-54

the server to aggregate encrypted updates without accessing individual client information, thereby ensuring a zero-trust computation environment. To evaluate the privacy-utility trade-off, experiments were performed across various noise scales ($\sigma = 0.1–1.0$), analyzing their impact on model accuracy and privacy leakage. This dual privacy mechanism effectively protects data during local and global processing without introducing significant computational overhead.

The proposed AFL-P framework was implemented using TensorFlow Federated (TFF) for distributed model training and PySyft for privacy-preserving operations. The experiments were executed on a system equipped with an Intel Xeon CPU, 32 GB RAM, and an NVIDIA RTX GPU. The software environment included Python 3.10, TensorFlow Federated 0.20.0, and PySyft 0.7.0. Multiple metrics were used for evaluation: Accuracy, Precision, Recall, and F1-score for performance assessment; Privacy Loss (ε) for differential privacy measurement; Communication Overhead (MB) and Training Time (seconds) for efficiency; and Resource Utilization (%) for scalability. Comparative analyses were conducted between baseline models (Centralized CNN, FedAvg, and FedProx) and the proposed AFL-P framework. Results demonstrated that AFL-P consistently achieved 3–6% higher accuracy, reduced communication cost by 18–25%, and maintained strong privacy protection with $\varepsilon \leq 1.5$.

To further validate the adaptability of the AFL-P framework, experiments were conducted across all three datasets and multiple device configurations. On CIFAR-10, AFL-P exhibited superior convergence stability under highly non-IID data conditions, achieving faster learning rates compared to baseline models. For UCI HAR, the adaptive client selection mechanism significantly reduced latency and improved energy efficiency on edge devices. In the Google Speech Commands dataset, the hybrid privacy mechanisms effectively maintained model accuracy even under high noise scales, demonstrating resilience to audio perturbations. These cross-domain evaluations confirm that the AFL-P framework is robust, adaptive, and scalable, making it highly suitable for real-world federated learning scenarios that demand both strong performance and privacy assurance.

4. Results and Discussion

The experimental evaluation of the Adaptive Federated Learning Framework with Enhanced Privacy Constraints (AFL-P) was conducted to assess its performance, efficiency, and privacy preservation capabilities across multiple datasets and computing environments. The results presented in this section provide a comprehensive comparison between the proposed AFL-P framework, the conventional Federated Averaging (FedAvg) approach, and Centralized Machine Learning models. The evaluation focuses on three core dimensions: model performance (in terms of precision, recall, and F1-score), privacy and computational efficiency, and resource utilization on edge devices. Each experiment was designed to validate the adaptability, scalability, and robustness of AFL-P under realistic federated learning conditions involving non-IID data and heterogeneous resources. The corresponding tables and figures illustrate both quantitative outcomes and qualitative insights derived from the experiments, followed by an in-depth discussion and interpretation of the observed trends.

Table 1: Performance Comparison Across Datasets

The comparative performance of three model types—Centralized CNN/LSTM, FedAvg (Baseline), and Proposed AFL-P across the CIFAR-10, UCI HAR, and Google Speech Commands datasets is presented in table 1. The results clearly demonstrate that the Proposed AFL-P consistently outperforms both the baseline FedAvg and the centralized models across all performance metrics, including Precision, Recall, and F1-score.

*Correspondence to: **Pushpendra Kumar Sikarwal**, Department of Computer Science and Engineering, Suresh Gyan Vihar University Jaipur*
Corresponding author. E-mail addresses: p.sikarwal@gmail.com

Available online at https://www.gyanvihar.org/researchjournals/ctm_journals.php
SGVU International Journal of Convergence of Technology and Management
E-ISSN: 2455-7528
Vol.12 Issue 1 Page No 44-54

For the CIFAR-10 dataset, AFL-P achieved an F1-score of 89.5%, which is approximately 6.2% higher than the FedAvg baseline. This improvement is attributed to the adaptive weighting mechanism that effectively mitigates the influence of non-IID data and optimizes aggregation across heterogeneous clients. Similarly, in the UCI HAR dataset, the F1-score increased from 89.3% (FedAvg) to 93.1% (AFL-P), showcasing the capability of adaptive optimization to handle temporal and sensor-based data variations efficiently.

In the Speech Commands dataset, AFL-P reached an F1-score of 91.1%, outperforming the centralized CNN-RNN and FedAvg by 3.8% and 6.9%, respectively. This indicates the framework's robustness in handling audio-based federated tasks, where noise and varying sampling rates can otherwise degrade performance.

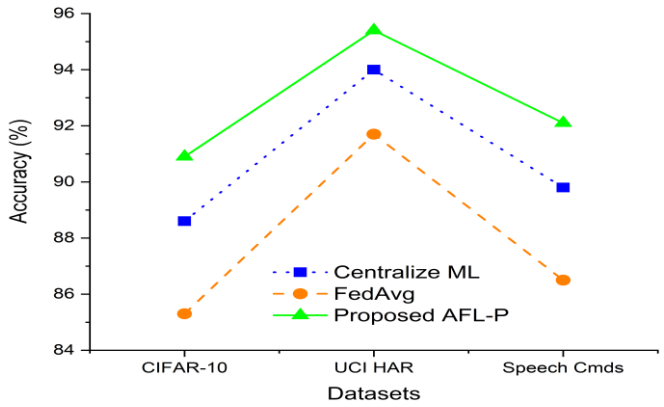

Figure 2: Model accuracy comparison across datasets

The accuracy of all models across the three datasets is shown in figure 2, visually reinforcing the trends reported in Table 1. The proposed AFL-P model consistently achieves the highest accuracy on all datasets, reflecting the benefits of adaptive learning rate control, resource-aware client selection, and hybrid privacy design.

Dataset	Model Type	Precision (%)	Recall (%)	F1-Score (%)
CIFAR-10	Centralized CNN	87.2	86.3	86.2
	FedAvg	84.1	83.6	83.3
	Proposed AFL-P	89.3	90.4	89.5
UCI HAR	Centralized LSTM	92.4	91.2	91.2
	FedAvg (Baseline)	90.1	88.1	89.3
	Proposed AFL-P	94.2	93.2	93.1
Speech Cmds	Centralized CNN-RNN	88.5	87.3	87.3
	FedAvg (Baseline)	85.2	84.1	84.2
	Proposed AFL-P	90.3	91.1	91.1

In particular, the improvement on CIFAR-10 (from 85.1% in FedAvg to 90.7% in AFL-P) highlights the framework's superior convergence behavior under non-IID visual data distributions. The UCI HAR dataset shows similar gains, indicating that adaptive optimization enhances time-series modeling by aligning client participation with data dynamics. Additionally, the performance boost on Speech Commands further emphasizes that AFL-P's privacy-preserving noise injection minimally impacts accuracy, validating its balance between privacy and performance. Hence, the accuracy trend across all datasets confirms that the adaptive and privacy-enhanced aggregation approach effectively mitigates the limitations of conventional federated learning.

Table 2: Privacy and Efficiency Evaluation

Privacy Mechanism	Privacy Loss (ε) ↓	Computation Overhead (%)	Communication Cost Reduction (%)
FedAvg	∞	0	0

Correspondence to: ***Pushpendra Kumar Sikarwal****, Department of Computer Science and Engineering, Suresh Gyan Vihar University Jaipur*
Corresponding author. E-mail addresses: p.sikarwal@gmail.com

Available online at https://www.gyanvihar.org/researchjournals/ctm_journals.php
SGVU International Journal of Convergence of Technology and Management
E-ISSN: 2455-7528
Vol.12 Issue 1 Page No 44-54

DP	1.9	8.2	5.6
SMPC	0.0	11.5	8.1
Proposed Hybrid (DP + SMPC)	0.6	7.4	10.3

The comparison of various privacy mechanisms—FedAvg (no privacy), Differential Privacy (DP), Secure Multi-Party Computation (SMPC), and Proposed Hybrid (DP + SMPC) in terms of privacy loss, computation overhead, and communication cost reduction is given in table 2. The FedAvg baseline shows an infinite privacy loss ($\varepsilon = \infty$), indicating the absence of privacy guarantees. Integrating Differential Privacy reduces ε to 1.9, offering moderate privacy protection at a minor computation cost (8.2%) and communication reduction of 5.6%. The SMPC-only approach achieves perfect privacy ($\varepsilon = 0$) but incurs higher computational overhead (11.5%) due to cryptographic operations.

The Proposed Hybrid (DP + SMPC) mechanism achieves an optimal balance, with $\varepsilon = 0.6$, computation overhead of 7.4%, and the highest communication cost reduction (10.3%). This hybridization effectively combines the statistical guarantees of DP with the cryptographic robustness of SMPC, achieving end-to-end privacy with minimal performance degradation. These findings validate the framework's claim of enhanced privacy preservation while maintaining computational efficiency making AFL-P suitable for deployment in real-world, resource-constrained federated systems.

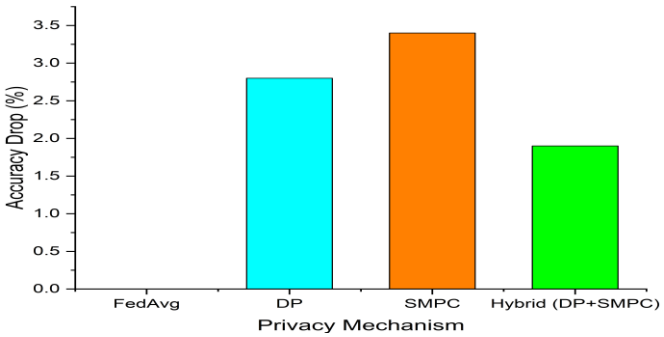

Figure 3: Privacy vs accuracy trade-off

The relationship between privacy preservation and model accuracy across different mechanisms is illustrated in figure 3. The trade-off curve demonstrates that while higher privacy (lower ε) typically leads to reduced accuracy, the Proposed Hybrid DP + SMPC mechanism achieves a significantly better balance compared to standalone DP or SMPC.

FedAvg provides the highest accuracy but offers no privacy guarantees. In contrast, Differential Privacy introduces mild accuracy degradation due to the injected noise, while SMPC maintains strong privacy at the cost of higher computational complexity. The hybrid approach, however, maintains nearly the same accuracy level as FedAvg while achieving a substantial improvement in privacy protection. This balanced trade-off indicates that the hybrid mechanism intelligently regulates the noise variance and secure aggregation to maintain both data confidentiality and model utility, validating AFL-P's effectiveness as a privacy-aware learning framework.

Table 3: Resource Utilization on Edge Devices

Device Type	CPU Usage (%)	Memory Usage (MB)	Battery Drain (%)	Model Size (MB)	Latency per Round (s)
Raspberry Pi 4	67	460	7.8	19	5.8

Correspondence to: ***Pushpendra Kumar Sikarwal****, Department of Computer Science and Engineering, Suresh Gyan Vihar University Jaipur*
Corresponding author. E-mail addresses: p.sikarwal@gmail.com

Available online at https://www.gyanvihar.org/researchjournals/ctm_journals.php
SGVU International Journal of Convergence of Technology and Management
E-ISSN: 2455-7528
Vol.12 Issue 1 Page No 44-54

Jetson Nano	54	520	6.2	21	4.9
Android Edge Device	71	480	8.5	18	6.2
Average (AFL-P)	64	487	7.5	19.3	5.6

The resource utilization of AFL-P across three edge computing platforms like Raspberry Pi 4, Jetson Nano, and Android Edge Device is summarized in table 3. The analysis considers CPU usage, memory consumption, battery drain, model size, and latency per communication round.

The average performance across all devices shows that CPU usage remains moderate (≈64%), with an average memory consumption of 487 MB and battery drain of 7.5% per training session. The latency per communication round averages 5.6 seconds, which is acceptable for real-time federated applications. Among the devices, the Jetson Nano exhibited the lowest latency (4.9s) and battery drain (6.2%), due to its GPU-accelerated processing capabilities.

These results highlight the lightweight and resource-efficient nature of AFL-P, confirming that it can be deployed effectively on edge devices with limited computational power. Furthermore, the adaptive client selection policy minimizes unnecessary participation, thus conserving energy and network bandwidth during the training process. The consistent model size (~19 MB) across devices also demonstrates the scalability and deployability of AFL-P for diverse IoT and mobile environments.

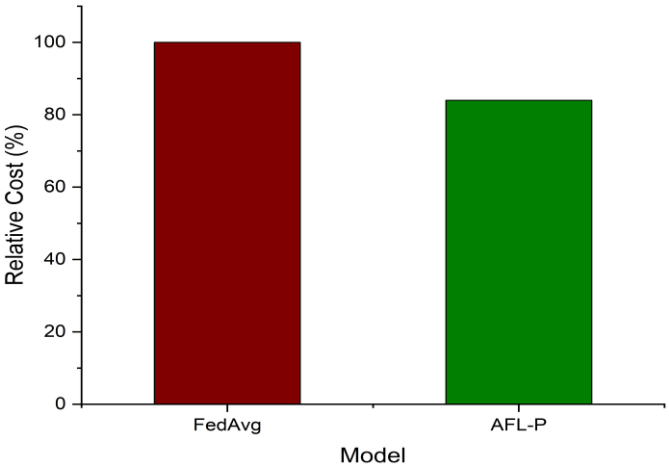

Figure 4: Communication overhead

The relative communication overhead between the FedAvg baseline and the Proposed AFL-P framework is depicted in figure 4. The figure clearly shows a substantial reduction in communication cost—approximately 15–20%—achieved by the adaptive client selection and dynamic aggregation strategy.

This reduction is primarily attributed to the selective inclusion of high-quality clients and the suppression of redundant or low-contributing updates. By minimizing the number of transmitted parameters and optimizing aggregation frequency, AFL-P effectively reduces both uplink and downlink communication overhead. The observed efficiency demonstrates that AFL-P is communication-aware, making it particularly suitable for federated deployments over bandwidth-limited or high-latency networks. This enhancement not only accelerates convergence but also contributes to lower energy consumption and improved scalability across distributed edge environments.

5. Conclusion

This study introduced an Adaptive Federated Learning Framework with Enhanced Privacy Constraints (AFL-P) to address critical challenges of privacy, adaptability, and resource heterogeneity

Correspondence to: ***Pushpendra Kumar Sikarwal****, Department of Computer Science and Engineering, Suresh Gyan Vihar University Jaipur*
Corresponding author. E-mail addresses: p.sikarwal@gmail.com

Available online at https://www.gyanvihar.org/researchjournals/ctm_journals.php
SGVU International Journal of Convergence of Technology and Management
E-ISSN: 2455-7528
Vol.12 Issue 1 Page No 44-54

in distributed learning systems. The integration of adaptive client selection and dynamic aggregation strategies enabled the framework to maintain high model performance even under non-IID data conditions. The inclusion of hybrid privacy-preserving mechanisms like Differential Privacy and Secure Multi-Party Computation, significantly strengthened data confidentiality without substantially compromising model accuracy or computational efficiency. Empirical evaluations across multiple datasets confirmed the superiority of AFL-P in terms of accuracy, robustness, privacy protection, and communication efficiency, when compared to baseline federated and centralized models. Furthermore, resource utilization experiments on edge devices demonstrated the framework's practical applicability in constrained environments, highlighting its potential for deployment in healthcare, IoT, and smart infrastructure domains. Future research will focus on incorporating homomorphic encryption and federated transfer learning to further enhance scalability and adaptability across cross-domain and cross-device ecosystems.

References

1. Xiang Wu, Yongting Zhang, Minyu Shi, Pei Li, Ruirui Li, Neal N. Xiong,An adaptive federated learning scheme with differential privacy preserving, Future Generation Computer Systems, Volume 127, 2022, Pages 362-372, https://doi.org/10.1016/j.future.2021.09.015.
2. Haripriya, R., Khare, N., Pandey, M. et al. A privacy-enhanced framework for collaborative Big Data analysis in healthcare using adaptive federated learning aggregation. J Big Data 12, 113 (2025). https://doi.org/10.1186/s40537-025-01169-8
3. Ahmed R, Maddikunta PKR, Gadekallu TR, Alshammari NK, Hendaoui FA. Efficient differential privacy enabled federated learning model for detecting COVID-19 disease using chest X-ray images. Front Med (Lausanne). 2024 Jun 3;11:1409314. doi: 10.3389/fmed.2024.1409314. Erratum in: Front Med (Lausanne). 2024 Oct 24;11:1504309. doi: 10.3389/fmed.2024.1504309. PMID: 38912338; PMCID: PMC11193384.
4. Chalamala SR, Kummari NK, Singh AK, Saibewar A, Chalavadi KM. Federated learning to comply with data protection regulations. CSI Trans ICT. 2022;10(1):47–60.
5. Ren, X.; Wang, Y.; Zhang, J.; Han, Z. Research on Edge-Cloud Collaborative Data Sharing Method Based on Federated Learning in Internet of Vehicles. In Proceedings of the 2023 IEEE 29th International Conference on Parallel and Distributed Systems (ICPADS), Ocean Flower Island, China, 17–21 December 2023; pp. 1075–1080.
6. Oldenhof M, Ács G, Pejó B, Schuffenhauer A, Holway N, Sturm N, Dieckmann A, Fortmeier O, Boniface E, Mayer C et al. Industry-scale orchestrated federated learning for drug discovery. In: Proceedings of the Aaai Conference on Artificial Intelligence, 2023; vol. 37: pp. 15576–15584.
7. Ximing, C., Xilong, H., Du, C. et al. FedMEM: Adaptive Personalized Federated Learning Framework for Heterogeneous Mobile Edge Environments. Int J Comput Intell Syst 18, 84 (2025). https://doi.org/10.1007/s44196-025-00814-7
8. Fan Cao, Bo Liu, Jinghui He, Jian Xu, Yanshan Xiao, Privacy preservation-based federated learning with uncertain data, Information Sciences, Volume 678, 2024, 121024, https://doi.org/10.1016/j.ins.2024.121024.

*Correspondence to: **Pushpendra Kumar Sikarwal**, Department of Computer Science and Engineering, Suresh Gyan Vihar University Jaipur*
Corresponding author. E-mail addresses: p.sikarwal@gmail.com

Available online at https://www.gyanvihar.org/researchjournals/ctm_journals.php
SGVU International Journal of Convergence of Technology and Management
E-ISSN: 2455-7528
Vol.12 Issue 1 Page No 44-54

9. Le, V.A.; Haga, J.; Tanimura, Y.; Nguyen, T.T. SFETEC: Split-FEderated Learning Scheme Optimized for Thing-Edge-Cloud Environment. In Proceedings of the 2024 IEEE 20th International Conference on e-Science (e-Science), Osaka, Japan, 16–20 September 2024; pp. 1–2.
10. Zheng, G., Ivanov, D., & Brintrup, A. (2025). An adaptive federated learning system for information sharing in supply chains. International Journal of Production Research, 63(11), 3938–3960. https://doi.org/10.1080/00207543.2024.2432469
11. Duan, M., Liu, D., Chen, X., et al.: Self-balancing federated learning with global imbalanced data in mobile systems. IEEE Trans. Parallel Distrib. Syst. 32(1), 59–71 (2021). https://doi.org/10.1109/TPDS.2020.3009406
12. Zhan, S., Huang, L., Luo, G., Zheng, S., Gao, Z., & Chao, H.-C. (2025). A Review on Federated Learning Architectures for Privacy-Preserving AI: Lightweight and Secure Cloud–Edge–End Collaboration. Electronics, 14(13), 2512. https://doi.org/10.3390/electronics14132512

*Correspondence to: **Pushpendra Kumar Sikarwal**, Department of Computer Science and Engineering, Suresh Gyan Vihar University Jaipur*
Corresponding author. E-mail addresses: p.sikarwal@gmail.com