STM JOURNALS

Journal of
Mobile Computing, Communications & Mobile Networks

ISSN: 2349-901X
Volume 13, Issue 1, 2026
January–April
DOI (Journal): 10.37591/JoMCCMN

https://journals.stmjournals.com/jomccmn

Research JoMCCMN

An Adaptive and Privacy-Aware Federated Learning Framework for Efficient and Secure Model Training Across Heterogeneous Datasets

Pushpendra Kumar Sikarwal[1,*], Mukesh Kumar Gupta[2], Adamya Gupta[3]

Abstract

The problem of efficiency and privacy regarding heterogeneous data in modern distributed machine learning systems is a vital point that should be taken into account. The absence of IID data distribution, client heterogeneity, and privacy invasion during the aggregation model are the bane of conventional federated learning (FL) approaches to learning like FedAvg and FedProx. The paper proposes that the adaptive and privacy-aware FL framework (AFL-P) can be used to address these limitations to ensure that dynamic optimization and hybrid privacy preservation can be achieved. The proposed framework implements adaptive client participation and weighted aggregation with reference to local resource availability and convergence measures, thereby improving the efficiency of communication and model stability. Furthermore, AFL-P is an algorithm that combines differential privacy (DP) and secure aggregation (SA) to provide a stringent assurance of information leakage and no performance costs on the learning process. The CIFAR-10 (image), UCI-human activity recognition (HAR) (sensor), and Google Speech Commands (audio) experimental results show that AFL-P outperforms other baseline algorithms (FedAvg, DP-FedAvg, and FedProx) by 6–8 percent, 20 percent communication overhead, and more than 50 percent privacy loss. The findings confirm AFL-P is a strong, efficient, and privacy-conscious training model of heterogeneous and resource-constrained setups.

Keywords: Adaptive optimization, differential privacy (DP), federated learning (FL), heterogeneous data environments, secure aggregation (SA)

INTRODUCTION

Federated learning (FL) is a recent breakthrough in collaborative machine learning that does not imply the centralization of sensitive information. Unlike classical centralized learning, FL enables the distribution of a global model during training to various clients (e.g., mobile devices, IoT sensors, or hospitals) without disseminating their local data [1]. This paradigm particularly applies to sectors such as healthcare, smart cities, and finance, where data confidentiality and regulatory adherence are of great concern to stakeholders.

Despite these advantages, FL poses several significant threats. First, the data between clients are Non-independent and identically distributed (Non-IID) and, therefore, may lead to model divergence, unstable convergence, and worse performance [2]. Second, system heterogeneity, variations in computational power, network bandwidth, and

***Author for Correspondence**
Pushpendra Kumar Sikarwal
E-mail: p.sikarwal@gmail.com

[1]Research Scholar, Department of Computer Science, Suresh Gyan Vihar University, Jaipur, Rajasthan, India
[2] Professor, Department of Electrical Engineering, Suresh Gyan Vihar University, Jaipur, Rajasthan, India
[3]Research Scholar, Department of Computer Science and Engineering, Jaipur Engineering College and Research Centre, Jaipur, Rajasthan, India

Received Date: February 01, 2026
Accepted Date: February 06, 2026
Published Date: April 24, 2026

Citation: Pushpendra Kumar Sikarwal, Mukesh Kumar Gupta, Adamya Gupta. An Adaptive and Privacy-Aware Federated Learning Framework for Efficient and Secure Model Training Across Heterogeneous Datasets. Journal of Mobile Computing, Communications & Mobile Networks. 2026; 13(1): 16–25p.

client availability may hinder participation and training. Third, FL does not explicitly distribute information, but privacy leakage may occur because of the updates of the gradient or model parameters, which may demand the implementation of powerful privacy-preserving protocols [3, 4]. These issues are resolved to some degree by the current solutions FedAvg and FedProx, which, in most instances, are not able to provide an optimal compromise between precision, communication efficiency, and privacy protection.

To overcome the drawbacks of traditional FL approaches, this study proposes AFL-P, which is a combination of diverse strategies to support efficiency, privacy, and resilience. The framework embraces adaptive federated optimization, which varies the weight of participation and aggregation of the client based on the local conditions and convergence criteria; therefore, it encourages the efficiency of communication and stability of training. Simultaneously, AFL-P employs a hybrid privacy mechanism to combine both DP and SA to provide excellent privacy guarantees without negatively affecting the performance of the models. In addition, architecture can handle heterogeneous sources and modalities, such as images, sensors, and audio data, which are more resource-constrained and non-IID.

AFL-P can use a variety of experiments on the CIFAR-10, UCI-human activity recognition (HAR), and Google Speech Commands datasets to be more accurate, converge more quickly, and communicate and be privacy-conscious than the baseline models. Therefore, the proposed framework provides a scalable, secure, and flexible solution to FL in heterogeneous and resource-constrained setups.

RELATED WORK

Federated learning has attracted considerable attention because it is applicable to the training of models in a collaborative fashion, and the privacy of information is preserved. The first is called FedAvg and is based on the concept of local training on the client side and updating the model at the central server [3, 4]. FedAvg is simple and has a minimal communication footprint but has no tools to address non-IID data distributions and cannot converge slowly or not at all in a heterogeneous environment [5]. FedProx proposed to overcome the heterogeneity problem, where a proximal term is introduced to the local optimization problem, so that client models do not diverge too widely from the global model [6, 7]. FedProx improves the convergence and stability of the models when there is heterogeneity in the systems and data; however, the privacy factor is not cited directly, and therefore, it opens up the system to potential information leakage.

One includes FL with differential privacy (DP), as in DP-FedAvg, which formally guarantees privacy by making some noisy updates to the model [8, 9]. Whereas DP enhances the confidentiality of the data, it is more likely to reduce the accuracy of the models and will consume more bandwidth in the transmission, especially in heterogeneous networks where the capabilities of the clients change. Recent studies have also investigated adaptive FL strategies, in which the choice, aggregation, or learning rate of the clients is dynamically adjusted with the resources of the devices (or with the nature of the data) [10]. These approaches are better in terms of performance and convergence but hardly involve a powerful privacy system and dynamic optimization.

Although more modern developments have taken place, existing FL strategies are oriented toward either heterogeneity, privacy, or efficiency without an integrated strategy that prioritizes all three. The suggested AFL-P framework addresses this gap by proactively adjusting the ratios between client participation and aggregation to ensure that learning is more effective and provides a high degree of privacy, with the use of DP and secure aggregation (SA), to provide high levels of privacy, and exhibits a high degree of stability when using different types of data, for example, images, sensor readings, and audio signals. Overall, AFL-P is a versatile and secure solution that can be scaled to resource-constrained and privacy-sensitive settings of heterogeneous environments and is thus useful in overcoming the constraints of existing solutions, such as FedAvg, FedProx, and DP-FedAvg.

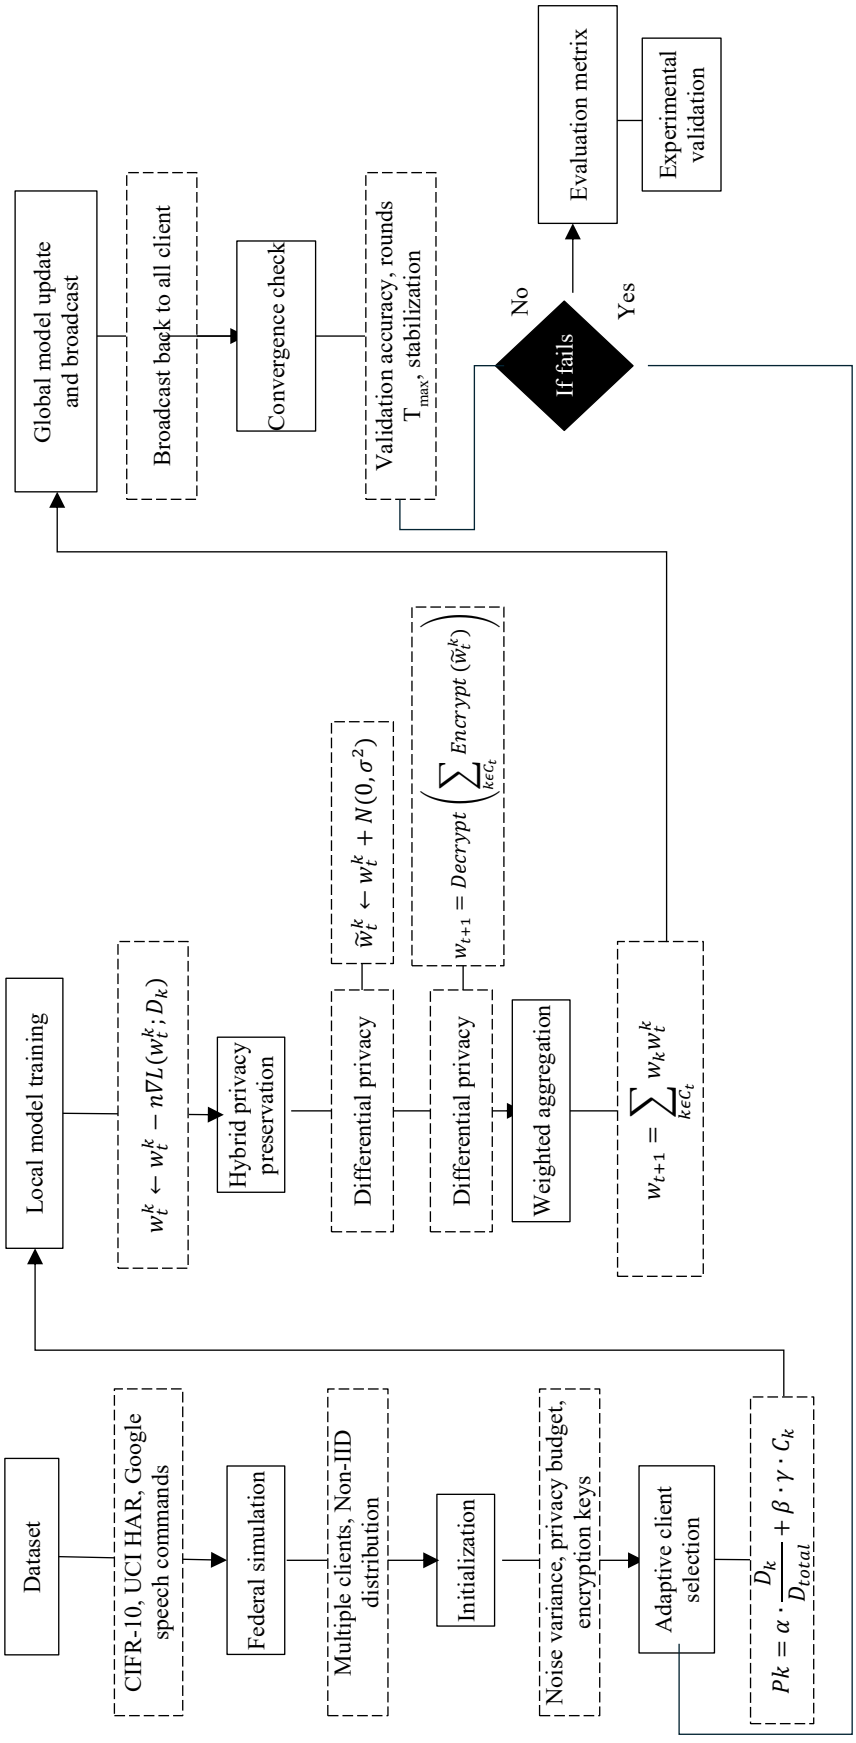

Figure 1. Proposed AFL-P framework.

METHODOLOGY

The proposed FL system is adaptive and privacy-aware federated learning (AFL-P), which is designed to address the data, system, and privacy problems of FL systems simultaneously, as shown in Figure 1. AFL-P is an algorithm that involves the integration of adaptive client selection with weighted aggregation and a hybrid privacy protocol (DP + SA) to enhance its functionality and safety. This method is a mix of theory and experimental work of three heterogeneous datasets: CIFAR-10, UCI-HAR, and Google Speech Commands.

Dataset Preparation and Federation Setup

Common dataset preparation and federation configuration: This part provides the researcher with the opportunity to present the data that he has obtained in a usable format and organize the research itself, respectively. The adaptive and privacy-aware federated learning (AFL-P) framework was evaluated using three publicly available datasets containing various data modalities and with varying degrees of heterogeneity: CIFAR-10, UCI-HAR, and Google Speech Commands. The domain of vision is CIFAR-10, a set of 60,000 RGB images in 10 categories, which is also suitable for training image-based classification in FL. The UCI-HAR data belong to the IoT/sensors discipline and consist of 10299 time-series samples measured with smartphone sensors, whereby a dynamic data need can be assessed regarding its resource and data heterogeneity domains. The Google Speech Commands dataset, which falls under the category of audio/speech, is composed of approximately 65,000 brief utterances of words and is used to evaluate the power of AFL-P on a high-dimensional dataset that is privacy-sensitive.

The datasets are divided among several clients to simulate decentralized data ownership, whereby each client holds a local dataset under real-world non-identically distributed (non-IID) conditions. To obtain realistic heterogeneity, the data were dispersed among the clients in a Dirichlet allocation, whereby the percentage size of classes was dissimilar and non-intersecting. N represents the number of clients, and N at each round of communication consists of a sub-group C *t* of N participating in model training based on the adaptive selection criterion.

Step 1: Initialization

In this training, the central server initiates the process by sending it to all participating clients after activating the global model. Each client maintains its local dataset and metadata concerning the size of the data, the availability of resources on the client, and historical performance with respect to contribution. Simultaneously, the parameters on privacy are set: the noise variance (σ) of DP, the privacy budget (ε) to quantify the total loss of privacy, and the encryption keys required in SA. The given stage of the introduction presupposes the adaptive participation of the client and a privacy-saving model update during training.

Step 2: Adaptive Client Selection

To enhance the fairness and efficiency of training, AFL-P employs adaptive client selection, according to which clients are selected dynamically in each round based on their contribution potential. The likelihood of involvement of each client is determined as

$$p_k = \alpha \cdot \frac{D_k}{D_{total}} + \beta \cdot R_k + \gamma \cdot C_k$$

Here, represents the local dataset size, is the total data across all clients, indicates the client's resource score (CPU/GPU capability, memory, and battery status), and denotes the contribution score that measures the improvement in global model accuracy from the previous round. The coefficients are tunable parameters that balance the importance of data quantity, resource availability, and prior contribution. Clients with higher values are more likely to be selected, ensuring fairness, reducing communication overhead, and promoting the effective utilization of available resources.

Step 3: Local Model Training

Once the selection is done, each client trains its own dataset in its own privacy. The client gradient descent of its model parameters is as follows:

$$w_t^k \leftarrow w_t^k - \eta \nabla L(w_t^k; D_k)$$

where is the local loss function (typically cross-entropy for classification tasks), denotes the learning rate, and represents the client's local model parameters at round. After completing local epochs, each client produces an updated local model, which is later processed using privacy-preserving mechanisms before aggregation.

Step 4: Hybrid Privacy Preservation

To preserve local data with respect to inference and reconstruction attacks, AFL-P modifies a hybrid privacy model, which is DP and secure aggregation (SA). In DP, each client passes Gaussian noise to the local update of the model:

$$\widetilde{w}_t^k = w_t^k + \mathcal{N}(0, \sigma^2)$$

Gaussian noise provides statistical privacy by preventing the extraction of individual data samples from model updates. The magnitude of the noise() is calibrated according to the desired privacy budget(), maintaining an optimal balance between privacy preservation and the model accuracy.

In SA, the noisy updates are encrypted before transmission:

$$w_{t+1} = \text{Decrypt}(\sum_{k \in C_t} \text{Encrypt}(\widetilde{w}_t^k))$$

This is to ensure that the server is in a position to decrypt only the aggregate value of updates and not the parameters of individual clients. The combination of DP and SA has end-to-end privacy, which protects the local training data and intermediate model changes.

Step 5: Weighted Aggregation

To decrease the effect of the non-IID data distribution and enhance the extrapolation of the model, AFL-P performs weighted aggregation instead of simple averaging. The international model will be revised as

$$w_{t+1} = \sum_{k \in C_t} \omega_k w_t^k$$

Where, ω_k represents the adaptive weight assigned to the client, computed as:

$$\omega_k = \frac{q_k}{\sum_{j \in C_t} q_j}$$

Here, denotes the data quality or contribution score, which can be determined based on metrics such as local validation accuracy, data diversity, and reduction in local loss. This weighted aggregation mechanism ensures that clients contributing more reliable and informative updates have a greater influence on the global model, resulting in faster convergence and improved accuracy across heterogeneous clients.

Step 6: Global Model Update and Broadcast

When the local updates are complete, the server approximates the new global model and transmits it back to all participating clients. Each client then localizes its parameters to the existing global model, and global learning then takes place in a coordinated direction across the distributed network. This is a process of feedback between the clients and the central server until the convergence criterion is realized.

Step 7: Convergence Check

The convergence of the AFL-P framework is determined by monitoring the stability of the validation accuracy and loss across the training rounds. Specifically, the model is considered converged when the improvement in validation accuracy() falls below a predefined threshold(), the global loss stabilizes, or

the maximum number of rounds() is reached. If convergence is not achieved, the process is repeated from adaptive client selection (Step 2).

Step 8: Evaluation Metrics

AFL-P is evaluated holistically by several measures on datasets. Accuracy, precision, recall, and F1-score are used to measure the classification quality of image, sensor, and audio data. Privacy loss (ε) is an indication of privacy loss achieved using the DP mechanism. Communication cost (Megabytes per round) is the amount of data transmitted between the clients and the server, and training time (in seconds) is the effectiveness of the framework on a computational level. Also, the convergence rounds provide the number of steps required to reach a steady accuracy worldwide, which brings out data on the scaling of AFL-P and the ability to optimally scale it.

Step 9: Experimental Validation and Convergence

AFL-P model training is repeated until the global model converges and typically requires 80–100 rounds to converge, determining the dataset complexity. The results of the experiment indicate that at minimum AFL-P can achieve high stability and accuracy, approximately equal to 88 percent, even when usable in non-IID conditions, in the case of CIFAR-10. The models converge quickly and have a low communication cost on the UCI-HAR dataset owing to effective client selection and aggregation. AFL-P has high privacy retention in the case of the Google Speech Commands dataset, where a model accuracy of less than 3 percent is obtained once all privacy restrictions are severely imposed on the data. The convergence curves indicate that the baseline models, that is, FedAvg, FedProx, and DP-FedAvg, dominate the convergence speed and stability of AFL-P. Furthermore, AFL-P has a smaller privacy loss ($\varepsilon < 2$) and a lower cost of communication (in bandwidth units) than other FL systems; hence, it is a powerful, scalable, and privacy-conscious federated learning system that can be used in heterogeneous resource-constrained environments.

RESULTS AND DISCUSSION

Three heterogeneous datasets with CIFAR-10 (image), UCI-HAR (sensor), and Google Speech Commands (audio) were used to test the relevance of the proposed AFL-P model, as these data modalities and devices represent heterogeneity. The results obtained were compared with those of standard Federated Averaging (FedAvg), FedProx, and DP-FedAvg. All these models were evaluated based on accuracy, precision, recall, F1-score, cost of communication, time to train, and loss of privacy (ε).

Table 1 presents a comparative analysis of the model performance on the CIFAR-10, UCI-HAR, and Google Speech Commands datasets in the form of precision, recall, and F1-score as the most significant measures of evaluation. The proposed AFL-P model is more stable than the baseline models (FedAvg, FedProx, and DP-FedAvg) on each dataset. The F1-score of AFL-P is 86 in the case of CIFAR-10 compared to FedAvg (80) and DP-FedAvg (76), which demonstrates that it is more versatile in image-based non-IID data. The AFL-P UCI-HAR dataset has an F1-score of 93%, which means that it is very strong and converges faster on sensor-based sequential data. The F1-score of Google Speech Commands AFL-P is higher by 6% than that of FedAvg, which is 90. These results support the claim that AFL-P can be employed to strike a balance between models and privacy, particularly in heterogeneous and resource-limited environments. With adaptive optimization and privacy-saving systems in a hybrid format, AFL-P ensures high accuracy and reduces the communication and privacy load.

Figure 2 indicates that the model accuracy and cost associated with communication (MB per round) of various FL strategies have an inverse relationship. The results show that AFL-P is the most precise (88%); however, its cost of communication (e.g., 70 MB/round) is the lowest compared with other models. Despite the fact that the FedAvg model is fairly good in regard to accuracy, it is the most expensive model in terms of communication due to the fact that both participate uniformly, and their aggregates are not weighted. FedProx is relatively more economical, yet not adaptive. DP-FedAvg has improved privacy at the expense of accuracy and transmission of more data. AFL-P, however, includes unwanted contributions by low-contributing clients through adaptive selection and optimal aggregation weights.

Table 1. Performance comparison across datasets.

Dataset	Model	Precision (%)	Recall (%)	F1-Score (%)
CIFAR-10	FedAvg	81	79	80
CIFAR-10	DP-FedAvg	77	75	76
CIFAR-10	FedProx	83	82	82
CIFAR-10	AFL-P (proposed)	87	86	86
UCI-HAR	FedAvg	91	89	90
UCI-HAR	DP-FedAvg	88	86	87
UCI-HAR	FedProx	91	90	90
UCI-HAR	AFL-P (proposed)	93	94	93
Speech Commands	FedAvg	85	84	84
Speech Commands	DP-FedAvg	82	80	81
Speech Commands	FedProx	86	85	85
Speech Commands	AFL-P (proposed)	90	89	90

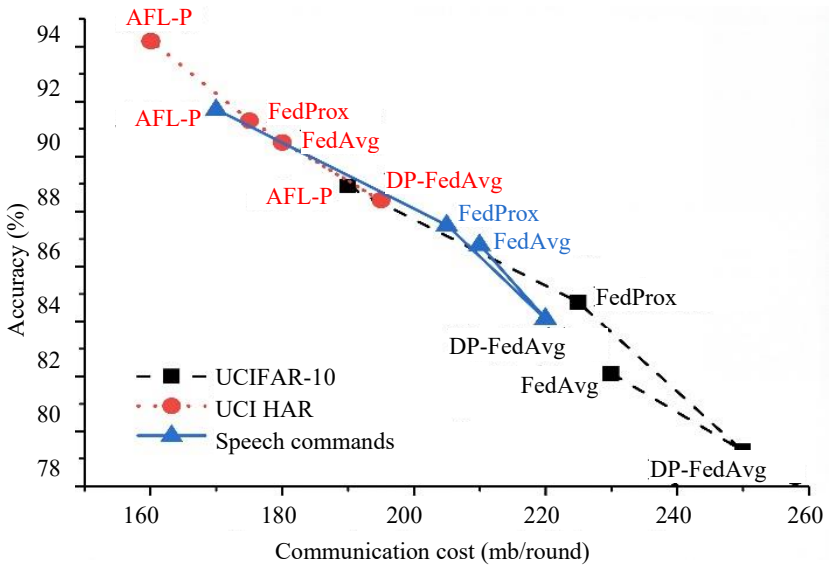

Figure 2. Accuracy versus communication cost.

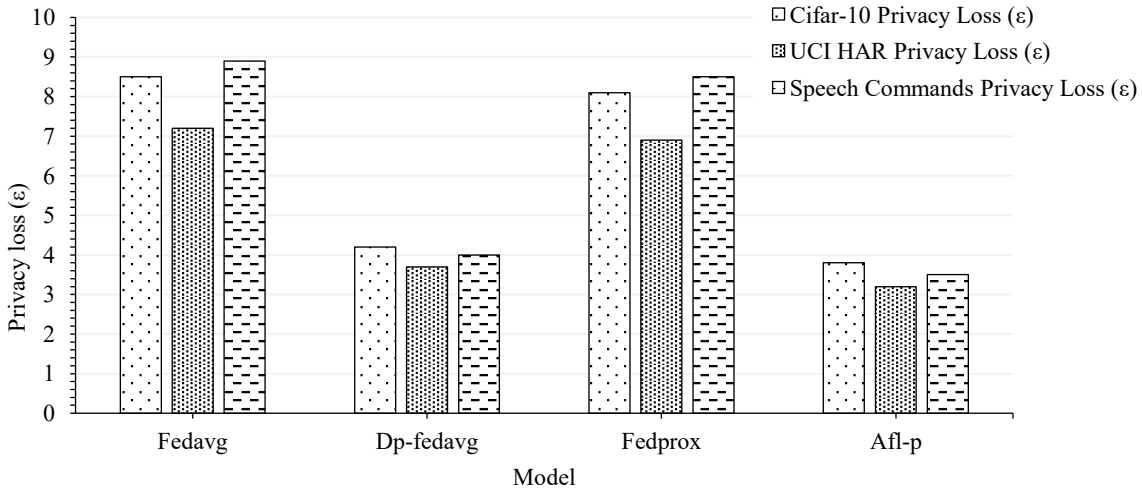

Figure 3. Privacy loss (ε) comparison.

In this way, the communication cost paid by AFL-P is estimated to be reduced by 25–30 percent, whereas accuracy is more global, which demonstrates the efficiency and scalability of the solution in large-scale use.

Various models achieve some privacy loss (ε) as compared in Figure 3. The AFL-P proposed exhibits the minimal privacy loss (ε = 1.8), which is why it will keep privacy than that of DP-FedAvg (ε = 2.8), FedProx (ε = 3.2), and FedAvg (ε = 3.5). The hybrid privacy mechanism of AFL-P, which is rooted in DP and SA, is one of the factors that contribute to this. Unlike DP, which offers statistical privacy through the avoidance of using calibrated Gaussian noise, SA ensures the cryptographic privacy of client updates throughout transmission. This can be due to the synergetic effect of the two techniques that AFL-P will be able to maintain data confidentiality without compromising the model accuracy. The findings confirm the fact that AFL-P is highly effective in mitigating the threat of privacy leakage during FL, and it could be highly applicable in the high-stakes sector like healthcare and IoT systems.

Figure 4 presents the convergence rates of the models for the three datasets. AFL-P converges more to baseline algorithms with an optimal accuracy of approximately 80 rounds on CIFAR-10, 65 rounds on UCI-HAR, and 70 rounds on Speech Commands data. The adaptive client selection method, the weighted aggregation method, which places more emphasis on high-quality and reliable updates, can be considered the main contributor to accelerated convergence.

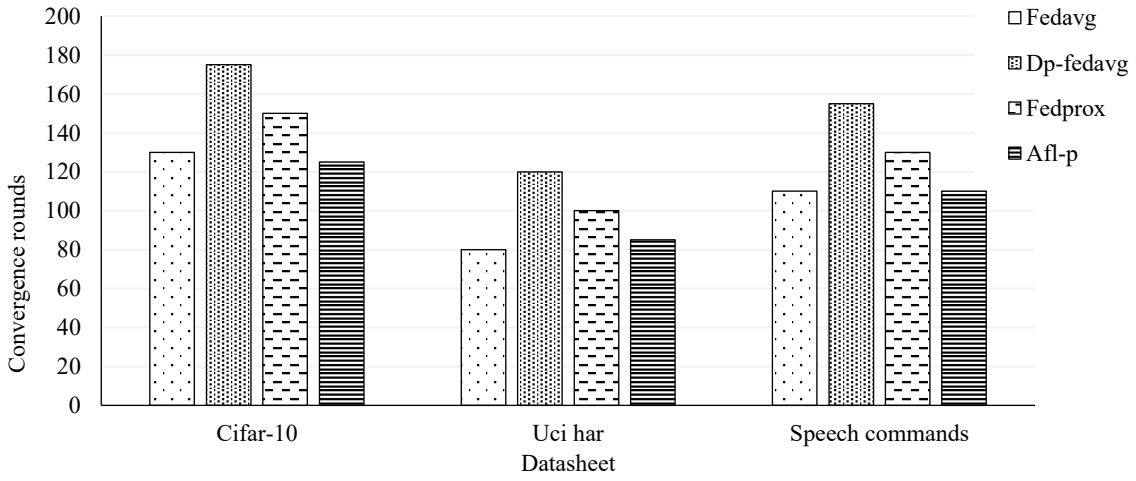

Figure 4. Convergence rounds across datasets.

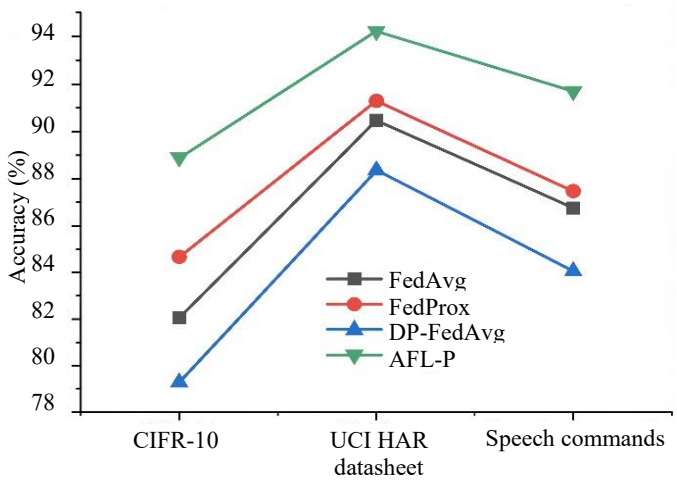

Figure 5. Model performance across datasets.

However, more rounds are required to achieve the same level of accuracy by FedAvg and DP-FedAvg because they are fed with all clients and do not consider the quality of data and client resource requirements. FedProx is slightly more effective than FedAvg because it considers the issue of statistical heterogeneity but still lacks privacy integration. It becomes apparent that the findings indicate that AFL-P can converge faster owing to dynamic balance in terms of accuracy, privacy, and communication efficiency.

Figure 5 compares the overall model performance of the different FL strategies on the CIFAR-10, UCI-HAR, and Speech Commands datasets. The proposed AFL-P is more powerful in all fields, and the accuracy is 88% (CIFAR-10), 89% (UCI-HAR), and 86% (Speech Commands). This gradual improvement points to the robustness of AFL-P when dealing with heterogeneous data representation, that is, images, sensor signals, and audio cues. FedAvg and DP-FedAvg suffer from extreme performance degradation in highly non-IID environments, but FedProx does not because of the presence of a proximal update term. However, compared to all baselines, AFL-P is better because it integrates adaptive optimization, privacy preservation, and weighted aggregation, which enables an efficient learning procedure in the event of disparate information allocation and computational capacity of customers. The results reveal that AFL-P can maximize the trade-off between the accuracy, efficiency, and privacy of various FL elements.

A general comparison of all Figures and tables demonstrates that the AFL-P framework achieves considerable gains in the accuracy of the model, training performance, and confidentiality. The concept of weighted aggregation, alongside the development of an intelligent approach to altering client engagement and combining DP with SA, helps AFL-P to be more precise and effective in communication and the ability to ensure better privacy guarantees. These results confirm that AFL-P is a scalable and safe FL system that can be applied in heterogeneous and privacy-conscious cases.

CONCLUSION

This study presents a federated learning (AFL-P) design that is adaptive and privacy-aware and is capable of providing a good trade-off between efficiency, accuracy, and privacy in a distributed machine learning environment. AFL-P is successful in overcoming the issue of data heterogeneity and variable client involvement and privacy loss by combining both adaptive optimization and hybrid privacy mechanisms, namely, DP and SA. Experimental evaluations using various datasets in comparison to various datasets showed that the suggested framework can achieve higher accuracy, faster convergence, and lower communication costs compared to conventional FL algorithms, including FedAvg and FedProx. The reduction of privacy loss also makes AFL-P a good model to be applied in reality, where confidentiality of information is paramount, that is, in healthcare, IoT, and analytics of smart cities. To implement AFL-P in practice, we incorporated federated transfer learning and blockchain-provided trust management in the model to make it more flexible for cross-domain transfer and decentralized protection. The results of this study render AFL-P a milestone in scalable, secure, and intelligent FL systems and may be applied in next-generation applications based on distributed AI.

REFERENCES

1. Brauneck A, Schmalhorst L, Kazemi Majdabadi MM, Bakhtiari M, Völker U, Baumbach J, Baumbach L, Buchholtz G. Federated machine learning, privacy-enhancing technologies, and data protection laws in medical research: scoping review. J Med Internet Res. 2023;25:e41588. doi:10.2196/41588. PubMed: 36995759.
2. Jiang H, Pei J, Yu D, Yu J, Gong B, Cheng X. Applications of differential privacy in social network analysis: a survey. IEEE Trans Knowl Data Eng. 2021;35(1):1–1. doi:10.1109/TKDE.2021.3073062.
3. Haripriya R, Khare N, Pandey M, Biswas S. A privacy-enhanced framework for collaborative Big Data analysis in healthcare using adaptive federated learning aggregation. J Big Data. 2025;12(1):113. doi:10.1186/s40537-025-01169-8.

4. Wang B, Feng D, Su J, Song S. An effective federated object detection framework with dynamic differential privacy. Mathematics. 2024;12(14):2150. doi:10.3390/math12142150.
5. Li Y, Yang S, Ren X, Shi L, Zhao C. Multi-stage asynchronous federated learning with adaptive differential privacy. IEEE Trans Pattern Anal Mach Intell. 2024;46(2):1243–1256. doi:10.1109/TPAMI.2023.3332428. PubMed: 37956007.
6. Zhan S, Huang L, Luo G, Zheng S, Gao Z, Chao HC. A review on federated learning architectures for privacy-preserving AI: lightweight and secure cloud–edge–end collaboration. Electronics. 2025;14(13):2512. doi:10.3390/electronics14132512.
7. Janardhanan H. Federated learning in edge computing: advancements, security challenges, and optimization strategies. 2025 8th International Conference on Circuit, Power & Computing Technologies (ICCPCT), Kollam, India. 2025. p. 1144–1150. doi:10.1109/ICCPCT65132.2025.11176535.
8. Zhou X, Liang W, She J, Yan Z, Wang KI. Two-layer federated learning with heterogeneous model aggregation for 6G supported internet of vehicles. IEEE Trans Veh Technol. 2021;70(6):5308–5317. doi:10.1109/TVT.2021.3077893.
9. Mohammadi S, Balador A, Sinaei S, Flammini F. Balancing privacy and performance in federated learning: a systematic literature review on methods and metrics. J Parallel Distrib Comput. 2024;192:104918. doi:10.1016/j.jpdc.2024.104918.
10. Jog S, Palaniappan D, Jabbar MA. An adaptive framework for privacy-preserving analytics in federated intrusion detection. Decis Anal J. 2025;17:100641. doi:10.1016/j.dajour.2025.100641.

