Article

Not peer-reviewed version

Privacy-Enhanced Federated Learning for Distributed Heterogeneous Data

Tianze Kang *, Haifeng Yang, Linyan Dai, Xin Hu, Junliang Du

Posted Date: 14 July 2025

doi: 10.20944/preprints202507.1160.v1

Keywords: federated learning; differential privacy; adaptive optimization; multi-center collaboration

Article

Privacy-Enhanced Federated Learning for Distributed Heterogeneous Data

Tianze Kang [1,*], Haifeng Yang [2], Linyan Dai [3], Xin Hu [4] and Junliang Du [5]

[1] San Francisco Bay University, Fremont, USA
[2] Northeastern University, Boston, USA
[3] University of California, Davis, Sacramento, USA
[4] Hofstra University, Hempstead, USA
[5] Shanghai Jiao Tong University, Shanghai, China
* Correspondence: tkang403@student.sfbu.edu

Abstract

This paper proposes a heterogeneous federated learning method for collaborative modeling, addressing data heterogeneity and privacy protection in multi-center scenarios. Based on the traditional federated learning framework, the method introduces a regularization term to mitigate performance instability caused by inconsistent data distributions across clients. It also designs a weighted aggregation strategy based on dynamic client features to enhance model convergence and robustness in multi-source heterogeneous environments. To ensure data privacy, the method applies differential privacy during local training. Noise is added to the model updates to prevent leakage of sensitive information. During development, the method is evaluated using a real federated learning dataset. It is compared with several mainstream approaches to validate its advantages in key performance metrics such as Accuracy, AUC, and F1-Score. Experimental results show that the proposed method effectively balances the personalized needs of local models with the goal of global consistency. It adapts to differences across clients and achieves better federated modeling performance. The overall architecture demonstrates strong practicality, security, and generalization ability. It is well-suited for real-world applications where centralized data storage is not feasible.

Keywords: federated learning; differential privacy; adaptive optimization; multi-center collaboration

I. Introduction

In today's world, where data-driven approaches have become the underlying logic of intelligent societies, the integration and efficient use of data resources have emerged as a key driver of artificial intelligence (AI) development. However, in practice, data sharing across organizations or institutions faces many restrictions due to data security regulations, privacy protection requirements, and industry barriers. In fields such as computer vision [1,2], finance [3], risk control [4], and manufacturing, data is often stored in isolated silos [5], creating information barriers. This "data immobility" severely limits the capacity for global modeling, reduces the generalization performance of algorithms, and hinders the realization of multi-center collaborative intelligence. Enabling cross-center intelligent mining without moving data out of its original location has become a critical challenge in the current development of AI.

Federated learning, as an emerging paradigm for distributed collaborative modeling, introduces the concept of "model federation, data locality" and provides a novel approach to addressing the above challenges. This method allows multiple data holders to collaboratively train models without directly exchanging data. It achieves joint modeling through local training and global model aggregation, theoretically balancing data security and collaboration efficiency. However, traditional federated learning often assumes a homogeneous environment where participants have similar data distributions, computational capabilities, and model structures. This assumption does not hold in real-world multi-center applications. In practice, participants usually face heterogeneity issues such as non-independent and identically distributed (Non-IID) data, varying hardware capabilities, and diverse task requirements. These issues lead to poor model convergence or degraded performance. Therefore, federated learning strategies for heterogeneous scenarios require in-depth research to enhance system adaptability and collaboration.

Moreover, with increasingly strict privacy protection regulations, such as global legal requirements for handling personal data, the "data remains local" strategy in federated learning is no longer sufficient to meet higher-level privacy needs. Since model parameters themselves may leak sensitive information, stronger privacy mechanisms must be incorporated into training and communication processes. Differential privacy, one of the mainstream theoretical approaches, can enhance privacy by injecting noise during model updates. This suppresses the influence of individual data samples while maintaining model utility. Deep integration of differential privacy with federated learning can create secure and efficient collaborative modeling frameworks. This is especially important in applications involving highly sensitive data.

The stability and efficiency of federated learning systems also depend on the design of their optimization mechanisms [6]. Under Non-IID and heterogeneous conditions, traditional model averaging methods, such as FedAvg, are prone to issues like model oscillation and performance degradation. As a result, adaptive optimization strategies have gained attention. These include dynamic weighting of local training, loss compensation mechanisms, and personalized scheduling. Such methods are crucial for improving the effectiveness of federated learning. By applying adaptive local updates and global aggregation mechanisms, it is possible to alleviate conflicts caused by heterogeneity, enhance model generalization, and improve overall convergence efficiency. Especially in multi-center collaborative tasks, adaptive optimization helps the system achieve flexible strategies and robust performance in collaborative modeling.

In summary, for real-world multi-center collaborative mining tasks, building a heterogeneous federated learning framework that integrates differential privacy and adaptive optimization has significant theoretical and practical value. On one hand, it responds to the growing societal demand for secure and intelligent data fusion, promoting the practical use of privacy-preserving computation in AI. On the other hand, it offers new solutions to modeling difficulties and low training efficiency in heterogeneous environments. This approach holds promise for breakthroughs in critical areas such as healthcare, smart manufacturing, and intelligent finance. The proposed research is a key step in making federated learning practical, controllable, and trustworthy, and contributes to the development of a secure and reliable AI ecosystem.

II. Related Work and Foundation

Recent advances in federated learning (FL) have addressed fundamental barriers in cross-center modeling, but practical deployment still faces challenges related to data heterogeneity and privacy assurance. A central concern is the tension between preserving local data privacy and achieving robust, generalizable models in environments where data distributions vary widely among clients. This study extends the field by integrating adaptive optimization and differential privacy mechanisms into a heterogeneous FL framework.

Differential privacy remains one of the most influential approaches for securing federated updates. El Ouadrhiri and Abdelhadi offered a comprehensive survey of differential privacy techniques in deep and federated learning, highlighting the importance of privacy budget management and the practical difficulties of balancing privacy with utility in collaborative models. Their analysis underlines our choice to use lightweight, noise-injection-based privacy at the local model-update level [7]. Similarly, Fu et al. provided a systematic review of differentially private federated learning, calling attention to the challenge of maintaining model convergence under strict privacy constraints—an issue directly addressed by the adaptive regularization in our approach [8]. In high-stakes application domains, such as healthcare, Adnan et al. demonstrated that FL can facilitate secure, privacy-preserving collaboration across institutions, supporting the practical relevance of our privacy-preserving design [9]. Aziz et al. further investigated privacy in FL by exploring homomorphic encryption in combination with differential privacy, an approach that inspires our security threat model, though our method favors a more efficient DP-centric solution [10].

Handling data heterogeneity and personalization in federated learning is another active research direction. Wei et al. tackled this by introducing personalized FL with differential privacy and convergence guarantees, providing a foundation for our dynamic, weighted aggregation strategy to accommodate non-IID client data [11]. Zhang et al. contributed a distributed protocol for secure collaboration in cross-domain FL, which reinforces the need for adaptive aggregation and personalization that our method implements through dynamic client weighting [12].

Adaptive optimization and dynamic scheduling have emerged as effective means to boost FL robustness. Works by Huang et al. and Liu et al. have demonstrated how reinforcement learning (RL)-based techniques—such as Q-learning and A3C—can optimize data mining and risk control in dynamic, distributed settings, an idea we extend by incorporating a reinforcement-style weighting mechanism for aggregation [13][14]. Sun et al. and Huang et al. also showed that RL-driven adaptive scheduling, such as DQN-based approaches, can enhance real-time system responsiveness, supporting our model's capacity for robust aggregation in dynamic environments [15][16].

Model efficiency and feature alignment are essential for scalable FL. Wang et al. developed a feature-alignment-based knowledge distillation approach for efficient model compression, echoing our goal of reducing communication overhead and supporting the scalability of our aggregation process [17]. Zheng et al.'s selective knowledge injection and structured gradient guidance approaches directly inspire our regularization strategy for harmonizing client updates while preserving privacy and stability [18][19]. Furthermore, Zhang et al.'s unified instruction encoding framework for multi-task models highlights the need for flexible aggregation rules—a principle reflected in our weighted client strategy [20].

Finally, recent advances in context-aware modeling and domain adaptation, such as those by Liu et al. and Zhao et al., demonstrate that personalized local models and advanced feature mining can significantly improve global model robustness in heterogeneous or noisy environments [21][22]. These insights support the multi-center adaptability and anomaly detection robustness of our proposed framework. In summary, current literature consistently highlights the necessity of combining adaptive, privacy-aware regularization with dynamic aggregation strategies to address the persistent challenges of privacy, heterogeneity, and efficiency in real-world federated learning. By integrating differential privacy directly into the optimization process and adopting reinforcement-inspired client weighting, our framework achieves stable and secure collaborative modeling even when client data are highly diverse. These advances contribute not only to improved accuracy and robustness in distributed anomaly detection but also to the practical deployment of federated learning systems in complex, privacy-sensitive environments.

III. Method

This study focuses on the core challenges in heterogeneous federated learning and builds a multi-center collaborative modeling framework that combines differential privacy mechanisms with adaptive optimization strategies. The model architecture is shown in Figure 1.

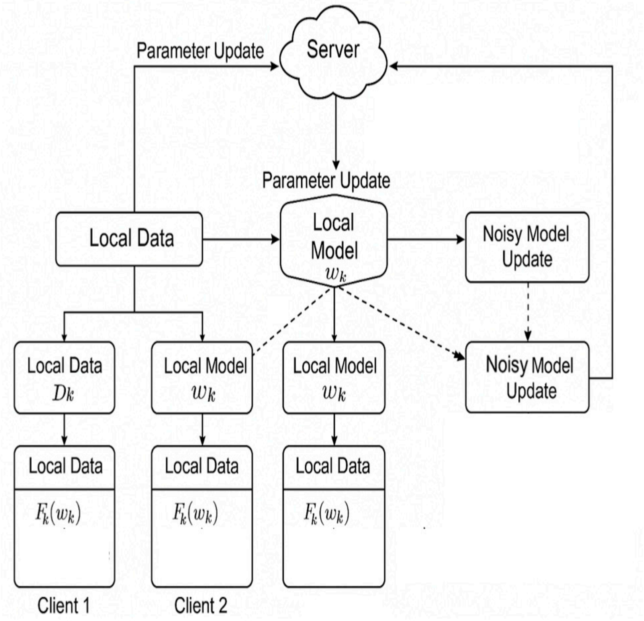

Figure 1. Overall model architecture diagram.

The architecture presents a federated learning framework where multiple clients perform local model updates on their private data and transmit differentially private gradients to a central server. The server aggregates these noisy updates using adaptive weighting to accommodate the heterogeneity among clients and iteratively refines the global model. This design effectively combines collaborative optimization with robust privacy guarantees in decentralized, non-IID data environments.

In the basic federated learning process, there are K clients, each client $k \in \{1,2,\ldots,K\}$ has a local dataset D_k, and the overall goal is to minimize the global loss function without sharing the original data:

$$\min_w F(w) = \sum_{k=1}^{K} \frac{n_k}{n} F_k(w) \text{ (1)}$$

w represents the model parameters, n_k is the number of samples of client k, and $n = \sum_k n_k$ is the total number of samples. Each client uploads the model parameters after performing several rounds of gradient descent locally, and the server performs weighted aggregation.

To address the heterogeneity of data distribution between clients, an adaptive local optimization strategy based on regularization terms is introduced to control the overfitting of local models to their own data. Specifically, client k minimizes the following objectives in local updates:

$$\min_w F_k(w_k) + \frac{\mu}{2} \| w_k - w^t \|^2 \qquad (2)$$

Where w^t is the current global model parameter, and μ is the regularization coefficient for adjusting the local offset. This method can ensure personalized learning capabilities while strengthening the coordination consistency with the global model and alleviating the oscillation convergence problem in the Non-IID scenario.

After each round of communication is completed, the server receives the model update Δw_k uploaded by each client and uses weighted average to perform global updates in the form of:

$$w^{t+1} = w^t - \eta \sum_{k=1}^{K} \frac{n_k}{n} \Delta w_k \qquad (3)$$

doi:10.20944/preprints202507.1160.v1

Where η is the learning rate. To further protect privacy information during transmission, the client uses a differential privacy mechanism to add noise perturbations to the update vector before uploading the model update. Specifically, let the local update be Δw_k, then its published version is:

$$\Delta' w_k = \Delta w_k + N(0, \sigma^2 I) \qquad (4)$$

Where $N(0, \sigma^2 I)$ is Gaussian noise with mean zero and variance σ^2, which achieves privacy protection for each round of update process.

In addition, in order to adapt to the heterogeneity of client resources and differences in computing power, this method introduces an adaptive dynamic weighting mechanism on the server side. This mechanism dynamically adjusts the aggregation weight α_k based on the stability index and local loss reduction of the client in the previous rounds of training. The final global update expression is:

$$w^{t+1} = w^t - \eta \sum_{k=1}^{K} \alpha_k \Delta' w_k \qquad (5)$$

This strategy improves the adaptability of federated optimization in resource-diverse and distributed heterogeneous environments, and improves the convergence efficiency of model aggregation. Through the above design, the overall framework realizes the collaborative modeling of heterogeneous data and resource environments under the premise of ensuring privacy, effectively taking into account security, robustness and modeling performance.

IV.Experimental Results

A Dataset

This study uses the FEMNIST dataset from the LEAF (Benchmarking Federated Learning Framework) platform. FEMNIST is a federated version of the classic EMNIST dataset. It is designed for image recognition tasks in cross-device, non-independent and identically distributed (Non-IID) environments. The data consists of handwritten digits and letters from thousands of users. It is partitioned by user, forming independent data subsets that reflect the real-world characteristics of data silos in multi-center scenarios.

FEMNIST contains over 800,000 image samples. Each image is a 28×28 grayscale image. The labels cover 62 character classes, including digits 0–9, uppercase letters A–Z, and lowercase letters a–z. The data is organized by user. Each client represents one user. The dataset is highly heterogeneous. It is suitable for evaluating the adaptability and stability of federated learning in Non-IID and highly diverse data environments.

This dataset is widely used in research on personalized model optimization, privacy-preserving modeling, and adaptive aggregation strategies. It effectively verifies the modeling capability of federated learning systems in resource-constrained and distributed data settings. It meets the experimental requirements of this study on multi-center heterogeneous collaborative modeling.

B Experimental Results

In this section, this paper first gives the comparative experimental results of the proposed algorithm and other algorithms, as shown in Table 1.

Table 1. Comparative experimental results.

Method	Acc	Auc	F1-Score
FedAvg [23]	78.4	0.852	0.765
FedProx [24]	80.1	0.867	0.781
FedDyn [25]	81.7	0.881	0.794

MOON [26]	82.3	0.889	0.802
FedPer [27]	79.5	0.861	0.773
Ours	84.6	0.914	0.827

As shown in the table, the proposed method outperforms all baselines across Accuracy, AUC, and F1-Score, achieving 84.6% Accuracy—over 6 points higher than FedAvg—and an AUC of 0.914, indicating strong generalization on heterogeneous multi-center data. Compared to FedProx, FedDyn, and MOON, the model demonstrates superior robustness and boundary recognition without compromising performance under differential privacy constraints. Its adaptive optimization mechanism effectively balances personalization and global consistency, mitigating overfitting to local data. These results confirm the method's effectiveness and scalability for real-world federated learning applications. The model's training dynamics are further illustrated in the loss curve shown in Figure 2.

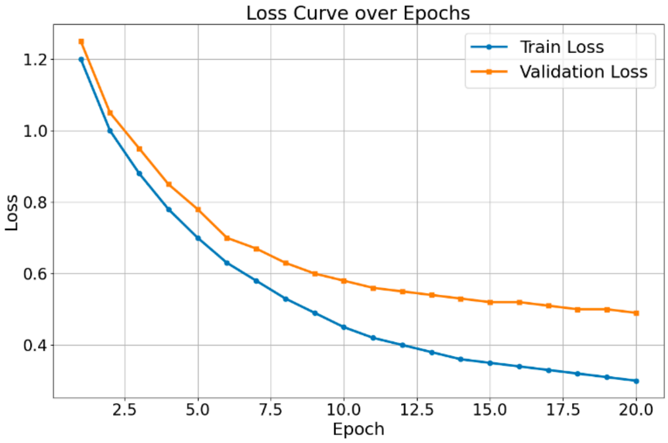

Figure 2. Loss function changes with epoch.

As shown in the figure, both the training loss and validation loss decrease steadily throughout the training process. This indicates good convergence and learning ability of the model during optimization. In the early stages, both losses drop rapidly, suggesting that the model learns data patterns quickly.

As the number of training rounds increases, the training loss continues to decline. The validation loss, however, decreases more slowly after round 10 and gradually stabilizes. This suggests that the model is approaching its optimal generalization performance while continuing to learn. Around round 20, the validation loss shows minimal change, indicating that the model may be nearing convergence.

Overall, the loss curves do not show signs of overfitting. The two curves remain at a consistent distance but trend toward each other. This suggests that the model maintains strong generalization ability while achieving high training accuracy. It also confirms the effectiveness of differential privacy and adaptive optimization in supporting model stability and performance improvement.

Finally, the experimental results of different learning rates are given, as shown in Table 2.

Table 2. Experimental results of different learning rates.

Learning Rate	**Acc**	**Auc**	**F1-Score**
0.004	81.2	0.879	0.791
0.003	82.7	0.894	0.805
0.002	83.5	0.902	0.817
0.001	84.6	0.914	0.827

As shown in the table, the model demonstrates a steady improvement across all performance metrics as the learning rate decreases. A relatively large learning rate, such as 0.004, leads to faster convergence. However, it may cause the model to overshoot during training, which negatively affects final accuracy and stability. This setting results in lower scores in Accuracy, AUC, and F1-Score.

When the learning rate is reduced to 0.002 and 0.001, the model's performance significantly improves. At a learning rate of 0.001, the model achieves its best results, with an Accuracy of 84.6%, an AUC of 0.914, and an F1-Score of 0.827. This indicates that a smaller learning rate enables more precise parameter updates. It helps the model better adapt to the Non-IID data structure and enhances overall modeling performance. These results demonstrate that in a federated learning framework, proper tuning of the learning rate significantly impacts training stability and generalization. A lower but stable learning rate facilitates more effective coordination of differences across clients and improves the overall performance of multi-center collaborative modeling.

V. Conclusions

This study proposes a heterogeneous federated learning approach that combines differential privacy and adaptive optimization for multi-center collaborative modeling. It aims to address the conflicts between inconsistent data distributions, heterogeneous resources, and privacy protection in real-world multi-institutional environments. By introducing a dynamic weighted aggregation mechanism and a local regularization-based optimization strategy, the method improves the stability and generalization ability of the global model while preserving the privacy of sensitive data. Experimental results demonstrate that the proposed method outperforms existing classical federated learning algorithms across various performance metrics. It exhibits enhanced robustness and adaptability in non-IID settings. These findings validate the practical significance of integrating differential privacy with adaptive optimization to enhance federated model performance. The method holds great promise for applications involving sensitive data, such as healthcare, finance, and public services.

In addition, the proposed framework exhibits good generalizability and scalability. It provides theoretical support and technical guidance for future multi-center collaborative intelligent modeling. In real-world deployments, this method can effectively reduce the impact of data silos between institutions. It enables efficient data value utilization under secure and controllable conditions and is expected to have significant impact in areas such as intelligent healthcare, personalized recommendation, and smart manufacturing.

VI. FUTURE Work

Future work may explore more efficient communication compression mechanisms, stronger personalized modeling methods, and the extension of cross-modal federated learning. Moreover, integrating advanced techniques such as large language models and self-supervised learning, under strict privacy guarantees, may further enhance the intelligence and adaptability of federated learning in complex real-world environments. This will support broader deployment and application of federated learning across diverse scenarios.

References

1. M. Li, R. Hao, S. Shi, Z. Yu, Q. He, and J. Zhan, “A CNN-Transformer approach for image-text multimodal classification with cross-modal feature fusion,” Proceedings of the 2025 8th International Conference on Advanced Algorithms and Control Engineering (ICAACE), pp. 1182–1186, 2025.
2. Y. Xiang, Q. He, T. Xu, R. Hao, J. Hu, and H. Zhang, “Adaptive transformer attention and multi-scale fusion for spine 3D segmentation,” arXiv preprint, arXiv:2503.12853, 2025.
3. Z. Xu, Q. Bao, Y. Wang, H. Feng, J. Du, and Q. Sha, “Reinforcement learning in finance: QTRAN for portfolio optimization,” Journal of Computer Technology and Software, vol. 4, no. 3, 2025.

4. Q. Sha, T. Tang, X. Du, J. Liu, Y. Wang, and Y. Sheng, “Detecting credit card fraud via heterogeneous graph neural networks with graph attention,” arXiv preprint, arXiv:2504.08183, 2025.
5. Q. Bao, J. Wang, H. Gong, Y. Zhang, X. Guo, and H. Feng, “A deep learning approach to anomaly detection in high-frequency trading data,” Proceedings of the 2025 4th International Symposium on Computer Applications and Information Technology (ISCAIT), pp. 287–291, 2025.
6. X. Wu, et al., "An adaptive federated learning scheme with differential privacy preserving," Future Generation Computer Systems, vol. 127, pp. 362–372, 2022.
7. A. El Ouadrhiri and A. Abdelhadi, "Differential privacy for deep and federated learning: a survey," IEEE Access, vol. 10, pp. 22359–22380, 2022.
8. J. Fu, et al., "Differentially private federated learning: a systematic review," arXiv preprint, arXiv:2405.08299, 2024.
9. M. Adnan, et al., "Federated learning and differential privacy for medical image analysis," Scientific Reports, vol. 12, no. 1, p. 1953, 2022.
10. R. Aziz, et al., "Exploring homomorphic encryption and differential privacy techniques towards secure federated learning paradigm," Future Internet, vol. 15, no. 9, p. 310, 2023.
11. K. Wei, et al., "Personalized federated learning with differential privacy and convergence guarantee," IEEE Transactions on Information Forensics and Security, vol. 18, pp. 4488–4503, 2023.
12. Y. Zhang, J. Liu, J. Wang, L. Dai, F. Guo, and G. Cai, “Federated learning for cross-domain data privacy: a distributed approach to secure collaboration,” arXiv preprint, arXiv:2504.00282, 2025.
13. X. Huang, Z. Zhang, X. Li, and Y. Li, “Reinforcement learning-based Q-learning approach for optimizing data mining in dynamic environments,” 2025.
14. J. Liu, X. Gu, H. Feng, Z. Yang, Q. Bao, and Z. Xu, “Market turbulence prediction and risk control with improved A3C reinforcement learning,” Proceedings of the 2025 8th International Conference on Advanced Algorithms and Control Engineering (ICAACE), pp. 2634–2638, 2025.
15. X. Sun, Y. Duan, Y. Deng, F. Guo, G. Cai, and Y. Peng, “Dynamic operating system scheduling using double DQN: a reinforcement learning approach to task optimization,” Proceedings of the 2025 8th International Conference on Advanced Algorithms and Control Engineering (ICAACE), pp. 1492–1497, 2025.
16. W. Huang, J. Zhan, Y. Sun, X. Han, T. An, and N. Jiang, “Context-aware adaptive sampling for intelligent data acquisition systems using DQN,” arXiv preprint, arXiv:2504.09344, 2025.
17. S. Wang, C. Wang, J. Gao, Z. Qi, H. Zheng, and X. Liao, “Feature alignment-based knowledge distillation for efficient compression of large language models,” arXiv preprint, arXiv:2412.19449, 2024.
18. H. Zheng, L. Zhu, W. Cui, R. Pan, X. Yan, and Y. Xing, “Selective knowledge injection via adapter modules in large-scale language models,” 2025.
19. H. Zheng, Y. Wang, R. Pan, G. Liu, B. Zhu, and H. Zhang, “Structured gradient guidance for few-shot adaptation in large language models,” arXiv preprint, arXiv:2506.00726, 2025.
20. W. Zhang, Z. Xu, Y. Tian, Y. Wu, M. Wang, and X. Meng, “Unified instruction encoding and gradient coordination for multi-task language models,” 2025.
21. J. Liu, Y. Zhang, Y. Sheng, Y. Lou, H. Wang, and B. Yang, “Context-aware rule mining using a dynamic transformer-based framework,” Proceedings of the 2025 8th International Conference on Advanced Algorithms and Control Engineering (ICAACE), pp. 2047–2052, 2025.
22. Y. Zhao, W. Zhang, Y. Cheng, Z. Xu, Y. Tian, and Z. Wei, “Entity boundary detection in social texts using BiLSTM-CRF with integrated social features,” 2025.
23. L. Collins, et al., "FedAvg with fine tuning: local updates lead to representation learning," Advances in Neural Information Processing Systems, vol. 35, pp. 10572–10586, 2022.

doi:10.20944/preprints202507.1160.v1

24. T. An, et al., "Consideration of FedProx in privacy protection," Electronics, vol. 12, no. 20, p. 4364, 2023.
25. C. Jin, et al., "FedDyn: a dynamic and efficient federated distillation approach on recommender system," Proceedings of the 2022 IEEE 28th International Conference on Parallel and Distributed Systems (ICPADS), 2023.
26. Q. Li, B. He, and D. Song, "Model-contrastive federated learning," Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021.
27. M. G. Arivazhagan, et al., "Federated learning with personalization layers," arXiv preprint, arXiv:1912.00818, 2019.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.