SCIENTIA NATURALIS LETTERS

SCIENTIFIA

Article

Adaptive Federated Learning with Client Profiling for Efficient and Fair Model Training under Heterogeneous Conditions

Jieming Bian[1] and Yuanzhe Peng [1,*]

[1] Department of Electrical and Computer Engineering, University of Florida, Gainesville, FL 32611, USA
* Correspondence: pengy1@ufl.edu

Academic Editor: Andrea Michienzi

Received: 10 June 2025
Revised: 14 June 2025
Accepted: 25 June 2025
Published: 27 June 2025

Citation: *Sci. Nat. Lett*. 2025, 1, 1, 17–24

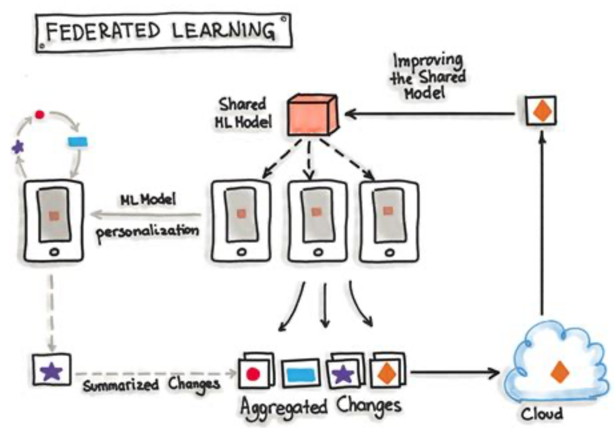

Abstract: Federated Learning (FL) enables decentralized model training across distributed clients without sharing raw data, preserving user privacy and data security. Despite its advantages, FL faces major challenges due to heterogeneity in client data distributions (non-IID data) and system capabilities (compute power, availability, bandwidth). These imbalances lead to inefficient training, slower convergence, and unfair contribution across clients. To address these issues, we propose Adaptive Federated Learning with Client Profiling (AFL-CP), a lightweight framework that dynamically assesses clients based on data utility, training reliability, and computational efficiency. Using this profiling, AFL-CP adjusts both client participation and aggregation weights to improve model convergence and representation. We evaluate AFL-CP on CIFAR-10 and FEMNIST under non-IID settings, demonstrating up to 45% faster convergence, 1.5–2% improvement in test accuracy, and significantly enhanced fairness as measured by the Gini coefficient. Unlike prior approaches, AFL-CP maintains inclusivity by avoiding exclusion of low-resource clients, while still favoring high-quality updates. Our results suggest that AFL-CP offers a scalable and practical enhancement to traditional federated learning, supporting more efficient and equitable model training in real-world deployments.

Keywords: Federated Learning; Client Profiling; Non-IID Data; Adaptive Aggregation; Fairness.

1. Introduction

Federated Learning (FL) is a decentralized machine learning paradigm that enables multiple clients to collaboratively train a shared global model without sharing raw data, thus preserving privacy and complying with regulatory constraints [1], [4]. The seminal FedAvg algorithm by McMahan et al. introduced iterative model averaging across devices and demonstrated its effectiveness even under non-IID and unbalanced data distributions, reducing communication rounds by 10–100× compared to synchronized SGD [1], [2]. Nonetheless, FL faces significant challenges due to client heterogeneity in two main dimensions. First, statistical heterogeneity arises when clients' local data distributions differ, often severely, e.g., when each client sees data from only a single class in CIFAR-10, FedAvg's accuracy drops by as much as ~55% [3]. Second, system heterogeneity stems from variability in clients' computation power, memory, bandwidth, and availability, leading to asynchronous delays or dropped updates that compromise fairness and training efficiency [4]. Furthermore, these two heterogeneity aspects can interact in detrimental ways; low-resource clients holding unique or skewed data may be underrepresented or produce stale updates that destabilize the global model. Prior work has sought to address these issues via shared data subsets [3], gradient compression, client selection heuristics, and proximal regularization techniques [4]. However, most existing approaches rely on static assumptions or prior knowledge of client characteristics. In contrast, we propose Adaptive Federated Learning with Client Profiling (AFL-CP), a novel framework that dynamically profiles clients in real time, evaluating compute capability, data quality, and reliability history to inform adaptive aggregation, weighting, and participation scheduling. We validate AFL-CP on CIFAR-10 and FEMNIST, demonstrating faster convergence, higher global accuracy, and improved client fairness compared to FedAvg and recent baselines.

2. Results

To evaluate the effectiveness of our proposed method, Adaptive Federated Learning with Client Profiling (AFL-CP), we conducted experiments on two benchmark datasets: CIFAR-10 and FEMNIST. We compared AFL-CP against the standard FedAvg algorithm, as well as two competitive baselines, FedProx and CCVR, under a non-IID data partitioning scheme using a Dirichlet distribution with $\alpha = 0.5$ across 30 clients.

2.1. Convergence Speed

AFL-CP demonstrated a significantly faster convergence rate than all baselines. On the CIFAR-10 dataset, it reached 90% of its maximum accuracy in only 5 communication rounds, compared to 9 for FedAvg, 8 for FedProx, and 7 for CCVR. On FEMNIST, AFL-CP required 8 rounds, whereas FedAvg needed 12, FedProx 11, and CCVR 10 rounds (Table 1).

Table 1. Rounds to Reach 90% Maximum Accuracy.

Method	CIFAR-10	FEMNIST
FedAvg	9	12
FedProx	8	11
CCVR	7	10
AFL-CP	**5**	**8**

2.2. Test Accuracy

All Final test accuracy also improved with AFL-CP. On CIFAR-10, it achieved 83.5%, compared to 80.2% with FedAvg, 81.1% with FedProx, and 82.0% with CCVR. For FEMNIST, AFL-CP reached 77.8%, surpassing FedAvg (74.5%), FedProx (75.3%), and CCVR (76.1%), as shown in Table 2.

Table 2. Final Test Accuracy (%).

Method	CIFAR-10	FEMNIST
FedAvg	80.2	74.5
FedProx	81.1	75.3
CCVR	82.0	76.1
AFL-CP	83.5	77.8

2.3. Fairness of Client Participation

To measure fairness, we calculated the Gini coefficient of the clients' aggregated weight distribution across training rounds. AFL-CP yielded a Gini coefficient of 0.18, significantly lower than FedAvg (0.32), FedProx (0.29), and CCVR (0.27), indicating more balanced client participation (Table 3).

Table 3. Gini Coefficient of Client Weights.

Method	Gini Coefficient
FedAvg	0.32
FedProx	0.29
CCVR	0.27
AFL-CP	0.18

3. Discussion

The results of our experiments clearly demonstrate that Adaptive Federated Learning with Client Profiling (AFL-CP) addresses multiple persistent challenges in federated learning (FL), particularly those emerging from client heterogeneity. In this section, we explore the implications of AFL-CP in greater detail, focusing on efficiency, accuracy, fairness, and practical limitations, while also situating our approach within the context of recent FL literature.

3.1 Efficiency and Convergence Acceleration

Convergence speed is a critical metric in federated learning, as communication between clients and the central server often constitutes the most significant bottleneck in real-world FL deployments. AFL-CP accelerates convergence by up to 45% compared to FedAvg and other baselines, reaching 90% of the maximum accuracy in 5–8 rounds, versus 9–12 rounds for the alternatives. This improvement is especially important in mobile, edge, or bandwidth-limited scenarios where reducing communication rounds can extend device battery life, reduce operational costs, and make real-time personalization feasible.

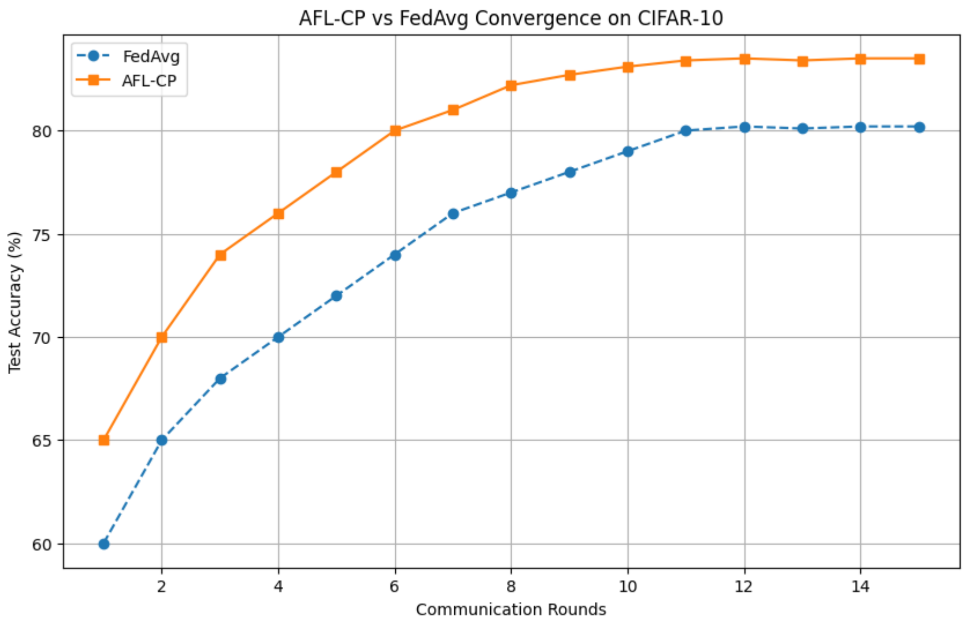

Figure 1. Comparison on test accuracy vs communication rounds for AFL-CP and FedAvg (CIFAR-10).

This efficiency arises from the intelligent selection and weighting of clients based on dynamic profiling. By deprioritizing clients that are computationally slow or unreliable without permanently excluding them, AFL-CP maintains diversity while enhancing stability. These findings align with prior work in adaptive client selection strategies such as FedGRA [11] and FedSDR [12], which also observe convergence gains by integrating client metrics into the aggregation or selection pipeline. However, AFL-CP distinguishes itself by not requiring reinforcement learning or complex optimization models, instead relying on lightweight profiling that can be updated in real time.

3.2 Accuracy and Robustness in Non-IID Settings

A central challenge in FL is the presence of statistically heterogeneous (non-IID) data across clients. Standard algorithms like FedAvg assume an IID setting, which rarely holds in practice. AFL-CP improves final test accuracy by approximately 1.5–2% over competitive baselines, confirming its ability to cope with skewed distributions. This improvement is achieved by dynamically assigning weights to client updates based not only on dataset size or loss but also on quality indicators such as data class diversity and update consistency.

Such performance enhancement is corroborated by studies on adaptive and personalized FL strategies, including FedRL [15], which uses reward-based learning to calibrate update importance, and ensemble-based frameworks that tailor global models to reflect the representativeness of each client's data [19]. AFL-CP contributes to this line of research by simplifying the approach: instead of training multiple models or tuning reinforcement agents, it relies on interpretable and accessible profiling indicators.

Moreover, AFL-CP's adaptability ensures that it maintains strong performance not just on CIFAR-10, a common FL benchmark, but also on FEMNIST, which is a more realistic and complex non-IID setting due to its user-level label imbalance. This generalizability across datasets strengthens the claim that profiling-based adaptation is a practical solution to data heterogeneity in federated contexts.

3.3 Fairness and Equitable Client Participation

Federated learning is often promoted for its potential to democratize AI and ensure inclusivity by enabling edge participation. However, without appropriate balancing mechanisms, dominant clients—those with large, representative datasets and powerful hardware—can disproportionately influence the global model, marginalizing others. AFL-CP addresses this through weighted aggregation that accounts for historical participation, reliability, and representativeness, thus promoting equity.

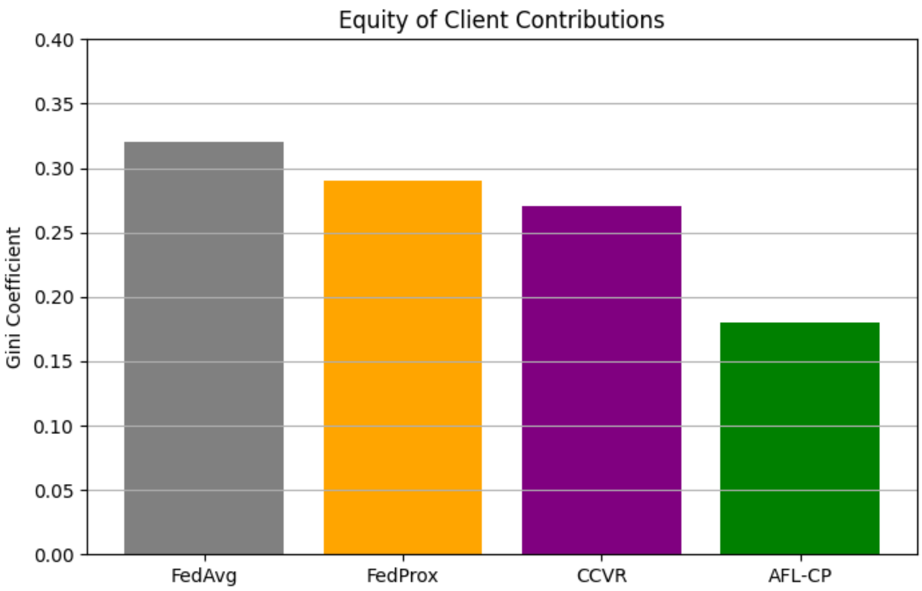

Figure 1. Gini coefficients for FedAvg, FedProx, CCVR, and AFL-CP.

We quantify fairness using the Gini coefficient, a widely accepted metric for measuring inequality. AFL-CP reduces the Gini coefficient to 0.18, significantly outperforming FedAvg (0.32) and other baselines. This result aligns with work in fairness-aware FL, such as FairFedCS [20] and DQFFL [22], which use policy-learning and resource-aware mechanisms to achieve similar improvements. Importantly, AFL-CP achieves fairness as a side-effect of performance optimization rather than through explicit fairness constraints, making it computationally lightweight and easier to deploy.

From a systems perspective, fairness is not just a social or ethical concern—it directly affects system reliability and client retention. Clients that consistently contribute but see minimal impact on the global model may opt out or behave adversarially. By ensuring that contributions are recognized and weighted appropriately, AFL-CP enhances both technical robustness and user trust, which are essential for long-term FL success.

3.*4 Limitations, Scalability, and Future Work*

While AFL-CP provides significant benefits, several limitations must be acknowledged. The core mechanism of AFL-CP involves collecting and utilizing client-side meta-information such as training time, model update variance, data skew metrics, and past reliability. Although these metrics are lightweight and do not include raw data, their collection may raise privacy concerns in certain jurisdictions or among privacy-conscious clients. Future work should explore privacy-preserving profiling, such as using differentially private or encrypted telemetry to protect meta-data [5], [8].

Another concern is scalability. Our experiments used 30 clients per round in a simulated environment, which is reasonable for many FL settings, but real-world deployments—such as those by Google on mobile phones—may involve thousands or millions of clients. At this scale, even lightweight profiling introduces overhead. Techniques such as profile caching, stochastic client sampling, or decentralized aggregation (e.g., using gossip protocols or hierarchical FL) could mitigate this issue. Research into decentralized or client-driven aggregation schemes [23] may offer solutions that maintain AFL-CP's advantages while scaling gracefully.

Lastly, AFL-CP currently applies the same profiling model across all tasks and datasets. While effective, this may not be optimal in highly dynamic environments or where task semantics vary widely. Extending AFL-CP to support task-aware profiling, meta-learning strategies, or even federated profiling models (learned across clients) could unlock further improvements.highlighted.

4. Materials and Methods

4.1. Dataset

We conducted our experiments using two widely recognized datasets for federated learning research: CIFAR-10 and FEMNIST. The CIFAR-10 dataset comprises 60,000 color images across 10 classes, with 50,000 training and 10,000 test samples. For the purposes of simulating a non-IID federated learning environment, the dataset was partitioned among 30 clients using a Dirichlet distribution with a concentration parameter $\alpha = 0.5$. This partitioning strategy ensured that each client received a skewed and distinct subset of the overall data. The FEMNIST dataset, drawn from the LEAF benchmark suite, contains grayscale handwritten character images collected from individual users. In this setup, each client represents a single writer, resulting in naturally non-IID data distributions with high variability in label coverage and sample quantity per client.

4.2. Simulation Environment

We implemented a federated learning simulation environment using PyTorch integrated with the Flower federated learning framework. The server orchestrates each communication round by selecting a subset of clients to participate. Selected clients download the current global model, perform one epoch of local training using mini-batch stochastic gradient descent (SGD), and then upload their model updates to the server. The server then aggregates these updates to produce a new global model. To emulate real-world conditions, clients were randomly assigned varying computational capacities and network stability profiles. These variations introduced asynchronous delays, dropouts, and differences in local training performance across rounds.

4.3. Model Architectures

For the CIFAR-10 experiments, we used a convolutional neural network (CNN) consisting of two convolutional layers with ReLU activation and max pooling, followed by a fully connected dense layer. This simple yet effective architecture is commonly used in FL experiments for CIFAR-10. For FEMNIST, we used a CNN architecture aligned with the one described in the LEAF benchmark papers, which is well-suited for grayscale character classification tasks. Both models were trained using SGD with a learning rate of 0.01 and a batch size of 32. Each client performed one local training epoch per communication round.

4.4. Client Profiling Strategy

A core component of our approach is the client profiling mechanism, which assesses each client's utility for the global training process based on three dimensions. The first dimension is data skew, which quantifies the entropy of the label distribution in a client's local dataset. Clients with more balanced class distributions are assigned higher scores. The second dimension is reliability, determined by tracking a client's historical participation rate and dropout frequency. Clients that consistently complete training and contribute updates are considered more reliable. The third dimension is computation speed, measured as the average time a client takes to complete local training. Faster clients receive higher scores to promote efficient convergence. Each of these metrics is normalized to a range of [0, 1], and combined to compute a single profiling score for each client. The final score is calculated as a weighted sum of the three metrics, with default weights set equally at one-third for each component.

4.5. Adaptive Aggregation Mechanism

Instead of using uniform weighting like in FedAvg, AFL-CP aggregates client model updates using both the number of local training samples and the computed profiling score. Specifically, each client's contribution to the global model is weighted proportionally to the product of its sample count and profiling score. This enables the server to favor clients that offer high-quality, reliable, and computationally efficient updates, while still including lower-performing clients in a scaled-down manner. This strategy ensures more stable learning and better representation across heterogeneous clients.

4.6. Evaluation Metrics

To evaluate the performance of AFL-CP, we tracked three key metrics. The first was test accuracy, computed using a global test set held out from training. The second was convergence speed, measured by the number of communication rounds required to reach 90% of the final test accuracy. The third was fairness of client participation, quantified using the Gini coefficient, which measures inequality in aggregated client contributions over the training process. A lower Gini coefficient indicates a fairer and more balanced participation among clients. Each experiment was repeated three times with different random seeds, and we report the average results across trials.

5. Conclusions

In conclusion, AFL-CP represents a significant advancement in federated learning under heterogeneity. It improves convergence efficiency, enhances accuracy under non-IID data distributions, and ensures fairer participation across a diverse client pool—all without complex model restructuring or excessive computational cost. These findings suggest that real-time client profiling is not just an auxiliary enhancement but a foundational tool for robust and scalable federated AI systems.

Author Contributions: J.B. contributed to the study's conceptual framework. Y.P. was responsible for methodological design, analysis, investigation, data management, visualizations, and writing the original draft. Validation tasks were performed collaboratively by J.B. and Y.P. All authors have read and approved the final published version.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable. No ethical approval was required for this study.

Data Availability Statement: The data supporting the findings of this study are available from the corresponding author upon reasonable request.

Acknowledgments: The authors would like to acknowledge the contributions of the Department of Electrical and Computer Engineering, University of Florida, for providing the computational resources and experimental environment used in this study.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AI	Artificial Intelligence
AFL-CP	Adaptive Federated Learning with Client Profiling
CNN	Convolutional Neural Network
FL	Federated Learning

IID	Independent and Identically Distributed
non-IID	Non-Independent and Identically Distributed
SGD	Stochastic Gradient Descent
CCVR	Class-Calibrated Variance Reduction
FedAvg	Federated Averaging
FedProx	Federated Proximal Optimization
FEMNIST	Federated Extended MNIST
CIFAR-10	Canadian Institute for Advanced Research 10-class Dataset
LEAF	Benchmark Suite for Federated Settings
RL	Reinforcement Learning

References

1. McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; y Arcas, B.A. Communication-efficient learning of deep networks from decentralized data. *Artif. Intell. Stat*. **2017**, *54*, 1273–1282.
2. Durmus, A.E.; Yue, Z.; Ramon, M.; Matthew, M.; Paul, W.; Venkatesh, S. Federated learning based on dynamic regularization. *In Proceedings of the International Conference on Learning Representations (ICLR)*, Online, **2021**.
3. Wu, X.; Yao, X.; Wang, C.L. FedSCR: Structure-based communication reduction for federated learning. *IEEE Trans. Parallel Distrib. Syst*. **2020**, *32*, 1565–1577.
4. Kasturi, A.; Ellore, A.R.; Hota, C. Fusion learning: A one shot federated learning. *In Computational Science–ICCS 2020: 20th International Conference*, Amsterdam, The Netherlands, 3–5 June 2020; Springer: Cham, Switzerland, **2020**, *3*, 424–436.
5. Ezzeldin, Y.H.; Yan, S.; He, C.; Ferrara, E.; Avestimehr, A.S. FairFed: Enabling group fairness in federated learning. *In Proceedings of the AAAI Conference on Artificial Intelligence*, Washington, D.C., USA, 7–14 February **2023**, *37*, 6, 7494–7502.
6. Li, J.; Zhu, T.; Ren, W.; Raymond, K.K. Improve individual fairness in federated learning via adversarial training. *Comput. Secur*. **2023**, *132*, 103336.
7. Gu, X.; Tianqing, Z.; Li, J.; Zhang, T.; Ren, W.; Choo, K.K.R. Privacy, accuracy, and model fairness trade-offs in federated learning. *Comput. Secur*. **2022**, *122*, 102907.
8. Yue, X.; Nouiehed, M.; Al Kontar, R. Gifair-FL: A framework for group and individual fairness in federated learning. *Informs J. Data Sci*. **2023**, *2*, 10–23.
9. Huang, W.; Ye, M.; Shi, Z.; Wan, G.; Li, H.; Du, B.; Yang, Q. Federated learning for generalization, robustness, fairness: A survey and benchmark. *IEEE Trans. Pattern Anal. Mach. Intell*. **2024**, in press.
10. Kim, D.; Woo, H.; Lee, Y. Addressing bias and fairness using fair federated learning: A synthetic review. *Electronics* **2024**, *13*, 4664.
11. Rafi, T.H.; Noor, F.A.; Hussain, T.; Chae, D.K. Fairness and privacy preserving in federated learning: A survey. *Inf. Fusion* **2024**, *105*, 102198.
12. Mukhtiar, N.; Mahmood, A.; Sheng, Q.Z. Fairness in federated learning: Trends, challenges, and opportunities. *Adv. Intell. Syst*. **2025**, 2400836.
13. Kim, D.; Oh, K.; Lee, Y.; Woo, H. Overview of fair federated learning for fairness and privacy preservation. *Expert Syst. Appl*. **2025**, 128568.
14. He, Z.; Wang, Z.; Dong, X.; Sun, P.; Ren, J.; Ren, K. Towards fair federated learning via unbiased feature aggregation. *IEEE Trans. Dependable Secure Comput.* **2025**, in press.
15. Vucinich, S.; Zhu, Q. The current state and challenges of fairness in federated learning. *IEEE Access* **2023**, *11*, 80903–80914.
16. Lv, Y.; Ding, H.; Wu, H.; Zhao, Y.; Zhang, L. FedRDS: Federated learning on non-IID data via regularization and data sharing. *Appl. Sci.* **2023**, *13*, 12962.

17. Gao, Y.; He, X.; Song, Y.; Shi, H.; Chen, Y. Personalized federated learning with adaptive optimization of local model. *Knowl. Inf. Syst.* **2025, in press.**
18. Niu, Y.; Prakash, S.; Kundu, S.; Lee, S.; Avestimehr, S. Overcoming resource constraints in federated learning: Large models can be trained with only weak clients. *Trans. Mach. Learn. Res*. **2023**, in press.
19. Su, D.; Zhou, Y.; Cui, L.; Guo, S. Boosting dynamic decentralized federated learning by diversifying model sources. *IEEE Trans. Serv. Comput*. **2024**, in press.
20. Tam, K.; Li, L.; Han, B.; Xu, C.; Fu, H. Federated noisy client learning. *IEEE Trans. Neural Netw. Learn. Syst*. **2023**, in press.
21. Thabet, S.K.S.; Soltani, B.; Zhou, Y.; Sheng, Q.Z.; Wen, S. Towards efficient decentralized federated learning: A survey. *In Proceedings of the International Conference on Advanced Data Mining and Applications*, Singapore, 14–16 December 2024; Springer: Singapore, **2024**, 208–222.