

Clustering and Aggregation for Wireless Federated Learning Under Label-Flipping Attacks

Mohamed Ads¹, Student Member, IEEE, Ruslan Zhagypar², Student Member, IEEE, Nour Kouzayha³, Member, IEEE, Hesham ElSawy⁴, Senior Member, IEEE, Tareq Y. Al-Naffouri⁵, Fellow, IEEE, and Hossam S. Hassanein⁶, Fellow, IEEE

Abstract—Federated Learning (FL) enables decentralized model training across distributed clients while preserving data privacy. However, its effectiveness can degrade in the presence of non-independent and identically distributed (IID) data distributions, unreliable wireless communication, and adversarial behavior such as label-flipping attacks. In this letter, we propose RANGE-CFL, a robust and network-aware framework for Clustered Federated Learning that addresses these challenges. RANGE-CFL integrates a network-aware aggregation scheme to mitigate communication unreliability and introduces a dynamic clustering mechanism with participation memory to maintain group coherence under intermittent connectivity. To defend against label-flipping attacks, we develop a gradient-based client filtering method that detects and removes malicious clients using PCA-assisted clustering over output neuron gradients. Experimental results on the EMNIST and CIFAR10 datasets show that RANGE-CFL enhances global model accuracy while simultaneously reducing inter-user accuracy variance. This demonstrates its robustness against label-flipping attacks and its effectiveness in promoting fairness and stability across heterogeneous user data distributions, outperforming the baseline FL algorithm.

Index Terms—Federated learning, wireless networks, stochastic geometry, trustworthiness, security.

I. INTRODUCTION

FEDERATED Learning (FL) has emerged as a transformative paradigm in distributed machine learning, enabling multiple clients to collaboratively train a global model without sharing their raw data [1]. This approach addresses critical concerns related to data privacy, security, and regulatory compliance, making it particularly suitable for sensitive applications such as healthcare, finance, and mobile services [2].

At the core of FL is the aggregation algorithm, which combines locally trained models from participating clients to update the global model. The original baseline method,

Federated Averaging (FedAvg), forms the foundation of the federated learning approaches due to its simplicity and effectiveness in aggregating local model updates, which performs a weighted average of client updates [3]. While FedAvg and its variants have demonstrated effectiveness in various scenarios, they often assume that all clients are trustworthy and the communication environment is ideal. These assumptions, however, do not hold in many real-world applications.

Deploying FL over wireless networks presents a set of intertwined challenges that must be jointly addressed for real-world applicability [4]. On one hand, FL systems are vulnerable to adversarial threats—particularly data poisoning and model manipulation attacks—where malicious clients aim to degrade the global model’s performance [5]. To counter these threats, several robust aggregation methods such as Krum [6], and Bulyan [7] have been developed. These approaches rely on the statistical properties of received gradients to detect and filter out malicious updates. On the other hand, deploying FL over wireless networks introduces its own set of challenges, including bandwidth limitations, variable latency, and signal interference, all of which can disrupt the efficiency and consistency of model aggregation [8], [9].

Despite their significance, these two dimensions—security and communication robustness—are often tackled in isolation. While [10], [11], [12] presents a relevant step toward simultaneously addressing security and wireless communication challenges in FL, their focus remains on model poisoning attacks, with data poisoning attacks receiving limited attention. In this letter, we bridge this gap by proposing RANGE-CFL, a clustered federated learning (CFL) framework that simultaneously enhances robustness against data poisoning and adapts to wireless communication constraints. RANGE-CFL integrates three complementary modules: a SINR-aware aggregation to mitigate communication bias, a memory-based clustering mechanism for stable grouping under intermittent participation, and a gradient-driven filtering module to detect and exclude label-flipping adversaries. Collectively, these modules improve the accuracy of all clients, ensuring that their performances remain closely aligned, while ensuring fairness among heterogeneous wireless clients.

The remainder of this letter is organized as follows: Section II shows the system model. Section III details the proposed aggregation framework. Section IV presents the experimental setup and evaluation metrics and discusses the results and insights gained from our experiments. Finally, Section V concludes this letter and outlines directions for future research.

Received 4 November 2025; revised 13 December 2025; accepted 9 January 2026. Date of publication 15 January 2026; date of current version 3 February 2026. This work was supported by King Abdullah University of Science and Technology (KAUST) under Award ORFS-CRG12-2024-6478. The associate editor coordinating the review of this article and approving it for publication was W. Mei. (Corresponding author: Mohamed Ads.)

Mohamed Ads, Hesham ElSawy, and Hossam S. Hassanein are with the School of Computing, Queen’s University, Kingston, ON K7L 2N8, Canada (e-mail: m.ads@queensu.ca; hesham.elsawy@queensu.ca; hossam@cs.queensu.ca).

Ruslan Zhagypar, Nour Kouzayha, and Tareq Y. Al-Naffouri are with the Department of Computer, Electrical and Mathematical Sciences and Engineering, King Abdullah University of Science and Technology, Thuwal 23955, Makkah, Saudi Arabia (e-mail: ruslan.zhagypar@kaust.edu.sa; nour.kouzayha@kaust.edu.sa; tareq.alnaffouri@kaust.edu.sa).

Digital Object Identifier 10.1109/LWC.2026.3654024

II. SYSTEM MODEL AND PROBLEM FORMULATION

A. Federated Learning

In an FL setting with U clients, a central server (CS), represented by a terrestrial base station (BS), coordinates the training process. Each client c holds a private local dataset D_c , performs training on it, and transmits the resulting model updates to the CS. Each BS aims to minimize the aggregate loss of the clients associated with it within its Voronoi cell, formulating the following distributed optimization problem:

$$\min_{\mathbf{g}^t} f(\mathbf{g}^t) = \sum_{c=1}^U \Omega_c F_c(\mathbf{g}^t), \quad (1)$$

where Ω_c is a weighting coefficient used in the aggregation of local models, which takes into account different factors (i.e., the size of each client's dataset). FL typically proceeds over multiple global rounds $t \in \{0, \dots, T-1\}$, where each client's local loss function $F_c(\cdot)$ is evaluated using the shared global model parameters \mathbf{g}^t . Thus, (1) seeks the optimal global model \mathbf{g}^t that minimizes the weighted average loss across all clients to capture the collective knowledge of the distributed network.

During each global round of the optimization process defined in (1), each participating device performs one local update to refine its model parameters. This is achieved by applying the following iterative update rule:

$$\mathbf{w}_c^{(t+1)} = \mathbf{w}_c^{(t)} - \lambda \nabla F_c(\mathbf{w}_c^{(t)}). \quad (2)$$

In this context, $\mathbf{w}_c^{(t+1)}$ captures the model parameters of client c after it has undergone one epoch of local training during the t -th global round. The parameter λ is the learning rate. Furthermore, $\nabla F_c(\cdot)$ is the gradient of the loss function, calculated with respect to the model parameters. After completion of all local training epochs, Each client transmits its updated weights to the BS. Since the learning rate is fixed across all clients, the BS can infer the equivalent gradient as the scaled difference between the received model and the previous global model.

B. Clustering in Federated Learning

An effective approach to handle the heterogeneity of users and non-IID data in FL is through client clustering [13], [14]. Specifically, to address these challenges, CFL employs a dynamic clustering mechanism that groups clients based on their pairwise cosine similarity [13].

The pairwise cosine similarity refers to the computation of similarity between every pair of clients update vectors $\Delta \mathbf{w}_c^{(t+1)} = \mathbf{w}_c^{(t+1)} - \mathbf{g}^{(t)}$, expressed as $\text{sim}_{i,j} = \frac{\langle \Delta \mathbf{w}_i, \Delta \mathbf{w}_j \rangle}{\|\Delta \mathbf{w}_i\| \|\Delta \mathbf{w}_j\|}$, where $\langle \cdot, \cdot \rangle$ is the inner product and $\|\cdot\|$ is the L_2 norm, forming a symmetric similarity matrix \mathbf{S} . This similarity matrix is converted to a distance matrix ($\mathbf{D} = -\mathbf{S}$) and used by the Hierarchical Agglomerative Clustering (HAC) algorithm with Complete Linkage in [13] to group all clients. Each client is assigned to exactly one cluster based on these similarity relations, ensuring that clients are not duplicated across clusters. After clustering, each BS aggregates the received models following the update rule $\mathbf{g}^{t+1} = \mathbf{g}^t + \frac{1}{U} \sum_{c=1}^U (\mathbf{w}_c^t - \mathbf{g}^t)$, where the global model is updated as the




Fig. 1. System model.

average deviation of local models from the previous global state. Each cluster ζ_i is then characterized by the average and maximum update magnitudes, μ_{avg}^t and μ_{max}^t , which measure the internal consistency and extremity of updates within the cluster. A cluster is split into two smaller sub-clusters following [13] when the *Max-to-Mean Ratio (MMR)* of update magnitudes exceeds a predefined threshold ϵ_1^t , i.e., $\frac{\mu_{\text{max}}^t}{\mu_{\text{avg}}^t} > \epsilon_1^t$, and the current cluster size is sufficiently large to allow meaningful division. As illustrated in Fig. 1, the clustering process depends solely on the behavioral similarity of model updates, reflecting local data distributions rather than geographical proximity.

C. Network Model

The considered network adopts a universal frequency reuse scheme for a set of orthogonal channels, with each client c occupying one resource block (RB). This approach, with the number of RBs greater than or equal to the number of devices in each cell, allows clients to transmit their data simultaneously without causing interference to one another within the same cell. In other words, *intra-cell interference is entirely avoided* due to the orthogonal allocation of RBs among clients in a given cell. However, despite the orthogonality within each cell, interference is still present in the form of *inter-cell interference*, which arises when devices in neighboring cells reuse the same RB. This leads to overlapping transmissions on the BS, which we denote as aggregate interference $I_c = \sum_{i \in \phi_c} P h_i r_i^{-\eta}$, where ϕ_c is the set of interfering clients affecting the uplink communication of the client c .

To simplify the analysis, we assume perfect downlink communication, meaning that the transmission from the BS to clients is error-free. For the uplink, all clients transmit with a constant power level P . The noise at the BS is modeled as additive white Gaussian noise with power N_0 and variance σ^2 , and small-scale fading is represented by a Rayleigh distribution denoted by h_c . Furthermore, the signal experiences path loss modeled by a power-law function, where the received signal power decays with distance r_c raised to the power $-\eta$, with η being the path-loss exponent. This framework captures both large-scale and small-scale effects that influence

the uplink communication performance in the considered FL system.

Building upon the described wireless communication environment, we now define the condition for successful uplink transmission. A client's local update is successfully received by the BS if the signal-to-interference-plus-noise ratio (SINR) exceeds a predefined threshold γ_t , expressed as:

$$\text{SINR}_c^{(t)} = \frac{Ph_c r_c^{-\eta}}{I_c + N_0} > \gamma_t, \quad (3)$$

The SINR threshold γ_t defines the minimum reception quality for successful model transmission. Clients failing to satisfy ($\text{SINR}_c^{(t)} > \gamma_t$) are excluded from aggregation. For controlled comparison, γ_t is fixed during all experiments.

D. Adversarial Threat Model

In the clustered FL setup, we consider adversarial clients within certain clusters that perform label-flipping attacks. Each adversary $a \in \mathcal{A} \subset U$ modifies its dataset D_a such that for all $(x, y) \in D_a$, the label is deterministically flipped from one class y_i to another y_j' ($i \neq j$), i.e., $y_i \rightarrow y_j'$. This manipulation introduces systematic bias into the global model, degrading its performance. The attack may target the most frequent classes or be randomly distributed across the label space, exploiting the dataset's label diversity. This behavior is especially harmful in CFL, as malicious clients within a cluster can collectively skew the cluster model and distort the aggregation outcome. In smaller clusters, even a few adversaries can dominate local updates and drive the model away from the true data distribution.

III. RANGE-CFL: ROBUST AGGREGATION UNDER LABEL-FLIPPING ATTACKS

In this section, we introduce our proposed framework, RANGE-CFL, which addresses the dual challenges of communication unreliability and adversarial behavior in clustered federated learning. RANGE-CFL incorporates memory-based clustering for client organization, SINR-aware weighting for fair aggregation, and gradient-based analysis for adversarial client detection.

A. Clustering With Participation Memory

Unreliable wireless communication means clients may fail to transmit their model updates in a given round. Since dynamic clustering depends on these updates to measure client behavioral similarity, RANGE-CFL introduces a participation memory mechanism. This mechanism ensures continuity by retaining and utilizing each client's most recent successfully received update vector for computing pairwise cosine similarities. This allows the clustering process to remain informed by a client's historical behavior even during participation failures.

B. Global Model Update With SINR-Aware Aggregation

Once clustering is established, RANGE-CFL incorporates communication reliability into the global model aggregation process. Specifically, each client's contribution is scaled by a SINR-aware weight $\kappa_c^{(t)}$, as introduced in (5). This weight

reflects the probability of successful transmission based on the client's uplink conditions.

The global model is updated using the following expression:

$$\mathbf{g}^{t+1} = \mathbf{g}^t + \frac{1}{U} \sum_{c=1}^U \mathbb{1}_{\{\text{SINR}_c^{(t)} > \gamma_t\}} \cdot (\mathbf{w}_c^t - \mathbf{g}^t), \quad (4)$$

where $\mathbb{1}_{\{\cdot\}}$ is the indicator function that takes the value one if $\{\cdot\}$ is true and zero otherwise, and $\mathbf{w}_{c,t}$ denotes the trained model update from client c .

C. Mitigating Label Flipping via Gradient Aggregation

To mitigate the influence of label flipping attacks, RANGE-CFL utilizes a robust gradient analysis technique introduced in [13]. The core idea is to monitor the gradients associated with the output neurons across multiple training rounds. Let K denote the total number of output classes. For each output neuron $k \in \{1, \dots, K\}$, we maintain a cumulative gradient score array $\mathbf{G} = [G_1, G_2, \dots, G_K]$, where each score G_k aggregates the influence of client-provided gradients over time.

In each round, the BS calculates the local output layer gradient vector $\nabla \mathcal{L}_c = [g_1^{(c)}, g_2^{(c)}, \dots, g_K^{(c)}] = \frac{\Delta w_c^{(c)}}{\lambda}$. However, to account for unequal participation, each gradient component is scaled inversely by the probability that the SINR condition in (3) is satisfied, i.e., $\mathbb{P}\{\text{SINR}_c^{(t)} > \gamma_t\}$. The SINR-aware weighting factor κ is therefore defined as:

$$\kappa_c^{(t)} = \frac{1}{\mathbb{P}\{\text{SINR}_c^{(t)} > \gamma_t\}} = \frac{\exp\left(\frac{\gamma_t N_0}{Pr_c^{-\eta}}\right)}{\mathcal{L}\left(\frac{\gamma_t}{Pr_c^{-\eta}}\right)}, \quad (5)$$

where $\mathcal{L}(s)$ represents the Laplace Transform of the uplink interference [15], given by:

$$\mathcal{L}(s) = \exp\left(-2\pi\lambda \int_0^\infty \frac{(1 - \exp(-\pi\lambda r^2))}{1 + \frac{r^\eta}{sP}} r dr\right). \quad (6)$$

This modeling approach ensures that devices with low SINR, often due to increased path loss or severe interference, are considered equitably in the gradient calculation process. By adjusting their weight based on their likelihood of successful transmission. The cumulative array is updated as follows:

$$G_k \leftarrow G_k + \kappa_c^{(t)} \cdot \nabla \mathcal{L}_c[k], \quad \forall k \in \{1, \dots, K\}. \quad (7)$$

After a sufficient number of training rounds, the accumulated gradient magnitudes begin to reveal patterns indicative of adversarial behavior. In particular, RANGE-CFL observes that two specific output neurons—corresponding to the *source* class (i.e., the true label) and the *target* class (i.e., the flipped label)—consistently exhibit higher gradient magnitudes compared to the rest. This disparity arises because benign clients pull the gradients toward the source neuron, while malicious clients systematically pull them toward the target neuron.

To exploit this observation, RANGE-CFL initiates a targeted clustering of the clients based on the gradients associated with these two neurons. Under the assumption that the number of malicious clients is smaller than that of benign clients, the clustering naturally separates the two groups. RANGE-CFL then identifies the smaller cluster as the likely group of adversaries and removes their updates from the aggregation process.

Algorithm 1: RANGE-CFL: Label Flipping Defense via Gradient Filtering

Input: K, U, C, T // classes, users, clusters, rounds
Output: g^T , global model after T rounds

- 1 Initialize model g^0 and gradient accumulator $G = 0$;
- 2 **for** each round $t = 0$ to $T - 1$ **do**
- 3 Server broadcasts g^t to all $c \in U$;
- 4 **for** each client $c \in U$ **in parallel do**
- 5 Client computes update w_c^{t+1} via local SGD then transmit it to their associated server;
- 6 **for** each received w_c^{t+1} by the server **do**
- 7 **for** each class $k = 1$ to K **do**
- 8 $G_k \leftarrow G_k + \kappa_c^t \cdot \nabla \mathcal{L}_c[k]$;
- 9 $g^{t+1} = g^t + \frac{1}{U} \sum_{c=1}^U \mathbb{1}_{\{\text{SINR}_c^{(t)} > \gamma_t\}} (w_c^t - g^t)$;
- 10 **if** *FilterClients* **then**
- 11 Identify top-2 gradient indices src, tgt from G_K ;
- 12 $\text{bad_clients} \leftarrow \text{Filter}((w_c^{t+1})_{c \in U}, \text{src}, \text{tgt})$;

Algorithm 2: Filter(\mathcal{D}, \mathcal{U})

Input: \mathcal{D} : matrix of weight vectors for selected clients
 \mathcal{U} : list of user IDs corresponding to rows in \mathcal{D}
Output: Set of user IDs detected as bad peers

- 1 Standardize \mathcal{D} using z-score normalization;
- 2 Apply K-Means clustering to \mathcal{D} with $K = 2$;
- 3 Assign cluster labels to each user in U based on K-Means output;
- 4 Count the number of users in each cluster: n_0, n_1 ;
- 5 **if** $n_0 < n_1$ **then**
- 6 $\text{bad_peers} \leftarrow \{c_i \in U \mid \text{label}(c_i) = 0\}$;
- 7 **else**
- 8 $\text{bad_peers} \leftarrow \{c_i \in U \mid \text{label}(c_i) = 1\}$;
- 9 **return** bad_peers

This clustering-based filtering strategy enables precise localization of malicious behavior, even under non-IID and partially observed communication conditions. It ensures that the global model remains resilient to manipulation while continuing to learn effectively from the honest subset of the network. This mechanism is summarized in Algorithm 1, and the adversary detection process is detailed in Algorithm 2. By integrating this filtering into the learning pipeline, RANGE-CFL effectively mitigates label flipping threats while maintaining robust and inclusive training across the network.

IV. EXPERIMENTAL SETUP

We evaluate the performance of RANGE-CFL using the EMNIST and CIFAR-10 dataset, which contains a large number of output classes, making it well-suited for assessing resilience against targeted label-flipping attacks in FL. To simulate statistical heterogeneity, we partition the dataset across $U = 30$ clients using a Dirichlet distribution with




Fig. 2. 2D PCA projection of clustered clients.

concentration parameter $\alpha = 0.5$, inducing non-IID label distributions. Clients are further grouped into three clusters to emulate the CFL environment. Visual heterogeneity is introduced by rotating input images: one-third of the clients receive images rotated 90° clockwise, another third 90° counter-clockwise, and the remaining third retain the original orientation. This setup encourages cluster-specific learning behavior and enables the evaluation of personalized aggregation strategies.

To model adversarial behavior, we designate 30% of clients across the clusters as malicious. These adversaries perform consistent label-flipping attacks by converting class-0 labels to class-9 during local training, simulating a targeted manipulation scenario. The global model's robustness is then evaluated against this structured threat. Clustering dynamics follow a gradient-based criterion, where a cluster is split once the Max-to-Mean Ratio (MMR) of update magnitudes surpasses the threshold $\epsilon_1^t = 2$, in accordance with [13]. Training runs for $T = 70$ and $T = 200$ rounds on EMNIST and CIFAR-10, respectively. In each round, participating clients perform $E = 1$ local epoch of SGD with a learning rate of 0.01 and a batch size of 32. The performance of RANGE-CFL is assessed based on the classification accuracy over a globally held-out test set.

Fig. 2 presents a 2D PCA projection of client neurons, where each neuron is constructed based on connections to the two most activated output neurons during local training. This projection reveals a clear separation between adversarial and benign clients after applying our filtering mechanism. Notably, adversarial clients form a distinct cluster away from the majority of benign clients, confirming the effectiveness of our gradient-based clustering strategy in isolating anomalous behavior, particularly under label-flipping attacks.

Fig. 3 illustrate the global model accuracy over 70 and 200 communication rounds across three configurations: (i) Without Filtering, (ii) Without Gradient Debiasing, and (iii) our proposed method RANGE-CFL, which integrates both adversarial filtering and gradient debiasing. The vertical line indicates the communication round at which cluster splitting is initiated, based on the threshold ϵ_1^t . Notably, the elimination of malicious clients is performed only after the split, allowing for more informed and reliable identification through accumulated gradient behavior.

As observed in Fig. 3, omitting the filtering step significantly degrades the global accuracy and introduces high variance among clients, as indicated by the wide red-shaded region (denoting standard deviation). Additionally, excluding

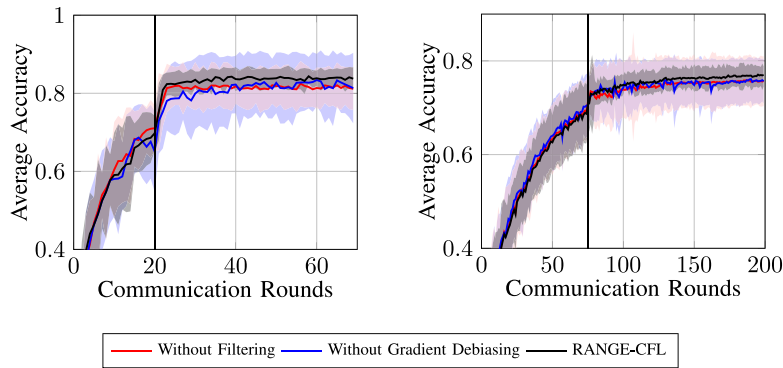


Fig. 3. Average test accuracy across communication rounds on EMNIST and CIFAR-10 datasets (left and right). Shaded regions denote \pm one standard deviation.

gradient debiasing—as defined in (7)—can lead to unstable learning behavior and performance oscillations. This degradation stems from inaccurate identification of the source and target output neurons, especially when malicious clients participate infrequently due to poor wireless conditions or intentional evasion. Without compensating for these participation disparities, the aggregated gradients become biased, potentially misguiding the filtering mechanism. In such cases, some benign clients may be incorrectly identified as malicious and excluded from training, while actual adversaries remain in the system. This misclassification is particularly harmful, as it can amplify the model’s vulnerability and degrade performance more severely than omitting filtering altogether, resulting in increased variance, as reflected in the figures.

In contrast, RANGE-CFL demonstrates superior and more consistent performance by jointly applying filtering and SINR-aware gradient debiasing. This combination ensures accurate adversary detection and robust learning, even under non-IID data distributions and unreliable communication scenarios. This confirms that our RANGE-CFL effectively adapts to adversarial behavior while preserving model performance under communication constraints and data heterogeneity.

V. CONCLUSION

This letter presented RANGE-CFL, a robust and communication-aware clustered federated learning framework that tackles three key challenges in decentralized learning: adversarial label-flipping attacks, unreliable wireless communication, and statistical heterogeneity. RANGE-CFL integrates three core components—SINR-aware aggregation for the communication, memory-based dynamic clustering for intermittent participation, and gradient-driven filtering for adversary detection. Experiments on EMNIST and CIFAR-10 dataset under realistic non-IID and adversarial conditions show that RANGE-CFL consistently outperforms baseline methods in both accuracy and robustness. The proposed clustering and filtering pipeline effectively isolates malicious clients while maintaining stable performance among benign participants. Moreover, the SINR-weighted aggregation ensures fairness and resilience for clients with poor connectivity. Future work

will explore asynchronous and decentralized FL extensions, as well as defenses against more advanced attacks such as multi-label flipping and backdoor injection, and evaluate RANGE-CFL’s scalability in large-scale deployments.

REFERENCES

- [1] A. Bakambekova et al., “On the interplay of artificial intelligence and space-air-ground integrated networks: A survey,” *IEEE Open J. Commun. Soc.*, vol. 5, pp. 4613–4673, 2024.
- [2] J. Wen et al., “A survey on federated learning: Challenges and applications,” *Int. J. Mach. Learn. Cybern.*, vol. 14, no. 2, pp. 513–535, 2023.
- [3] B. McMahan et al., “Communication-efficient learning of deep networks from decentralized data,” in *Proc. 20th Int. Conf. Artif. Intell. Statist.*, 2017, pp. 1273–1282.
- [4] R. Zhagypar et al., “Characterization of the global bias problem in aerial federated learning,” *IEEE Wireless Commun. Lett.*, vol. 12, no. 8, pp. 1339–1343, Aug. 2023.
- [5] Y. Chen et al., “Federated learning attacks and defenses: A survey,” in *Proc. IEEE Int. Conf. Big Data (Big Data)*, 2022, pp. 4256–4265.
- [6] P. Blanchard et al., “Machine learning with adversaries: Byzantine tolerant gradient descent,” in *Proc. 31st Adv. Neural Inf. Process. Syst.*, 2017, pp. 1–11.
- [7] E. M. El Mhamdi et al., “The hidden vulnerability of distributed learning in byzantium,” in *Proc. 35th Int. Conf. Mach. Learn.*, 2018, pp. 3521–3530.
- [8] M. Ads et al., “Risk-aware accelerated wireless federated learning with heterogeneous clients,” in *Proc. IEEE Int. Conf. Commun. (ICC)*, 2024, pp. 2950–2955.
- [9] R. Zhagypar et al., “UAV-assisted unbiased hierarchical federated learning: Performance and convergence analysis,” *IEEE Trans. Mach. Learn. Commun. Netw.*, vol. 3, pp. 420–447, 2025.
- [10] M. Ads et al., “RARE-FL: Resilient accelerated and risk-aware edge federated learning in scarce data scenario,” *IEEE Wireless Commun. Lett.*, vol. 13, no. 12, pp. 3424–3428, Dec. 2024.
- [11] H. Wen et al., “A unified federated learning framework for wireless communications: Towards privacy, efficiency, and security,” in *Proc. IEEE Conf. Comput. Commun. Workshops (INFOCOM WKSHPS)*, 2020, pp. 653–658.
- [12] Y. Liu et al., “A secure federated learning framework for 5G networks,” *IEEE Wireless Commun.*, vol. 27, no. 4, pp. 24–31, Aug. 2020.
- [13] F. Sattler et al., “Clustered federated learning: Model-agnostic distributed multitask optimization under privacy constraints,” *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 32, no. 8, pp. 3710–3722, Aug. 2021.
- [14] P. Sun et al., “Dual-segment clustering strategy for hierarchical federated learning in heterogeneous wireless environments,” *IEEE Wireless Commun. Lett.*, vol. 14, no. 6, pp. 1777–1781, Jun. 2025.
- [15] H. ElSawy et al., “Stochastic geometry for modeling, analysis, and design of multi-tier and cognitive cellular wireless networks: A survey,” *IEEE Commun. Surveys Tut.*, vol. 15, no. 3, pp. 996–1019, 3rd Quart., 2013.