Privacy-Preserving Federated Learning in Vehicular Edge Computing: A Differential Privacy Approach with Adaptive Gradient Compression

[1] Johan bin Mohamad Sharif
[1] Department: Computer Science, Universiti Teknologi Malaysia,
Johor, Malaysia.
[1] E-mail: johan@utm.my

Abstract

The proliferation of connected vehicles and intelligent transportation systems has generated unprecedented volumes of mobility data, creating both opportunities for collaborative machine learning and risks to driver privacy. Federated Learning (FL) offers a decentralized paradigm for training shared models without centralizing raw vehicular data, yet its deployment in Vehicular Edge Computing (VEC) environments faces acute challenges from high mobility, intermittent connectivity, and heterogeneous computing resources. This paper proposes AGCP-FL, an Adaptive Gradient Compression and Privacy-preserving Federated Learning framework specifically designed for VEC networks. Our approach integrates differential privacy with adaptive gradient compression that dynamically adjusts sparsification ratios based on wireless channel quality, privacy budget constraints, and vehicle mobility patterns. A Rényi Differential Privacy (RDP) accountant tracks cumulative privacy loss across communication rounds, while a context-aware compression engine reduces uplink communication overhead by up to 97% without compromising model convergence. Experimental evaluation on a simulated urban VEC scenario with 500 vehicles demonstrates that AGCP-FL achieves 91% test accuracy with a privacy budget of $\varepsilon = 1.0$, reduces per-round communication cost to 3.5 MB (compared to 100 MB for standard FL), and maintains robust convergence under 40% vehicle churn rates. These results validate AGCP-FL as a practical privacy-preserving solution for large-scale collaborative learning in dynamic vehicular environments.

Keywords: *Federated Learning, Vehicular Edge Computing, Differential Privacy, Gradient Compression, Privacy-Preserving Machine Learning, Connected Vehicles, Intelligent Transportation*

1. Introduction

The rapid advancement of autonomous driving and intelligent transportation systems (ITS) has positioned vehicular networks as one of the most data-intensive domains in modern computing. Connected vehicles equipped with cameras, LiDAR,

GPS, and onboard diagnostics generate gigabytes of sensory data per hour, offering rich opportunities for collaborative machine learning models that improve traffic prediction, route optimization, and safety-critical decision-making [1]. However, the centralized collection of such data raises profound privacy concerns—vehicular trajectory data can reveal sensitive information including home locations, travel patterns, and personal habits that drivers rightfully expect to remain confidential.

Federated Learning (FL) has emerged as a compelling alternative to centralized training, enabling distributed clients to collaboratively learn a shared model while keeping raw data localized [2]. In the context of Vehicular Edge Computing (VEC), where roadside units (RSUs) provide computational offload points within cellular base stations, FL aligns naturally with the decentralized architecture of vehicular networks. Yet, standard FL implementations face severe challenges in this domain: vehicles exhibit extreme mobility (up to 120 km/h on highways), wireless connectivity is intermittent due to handovers and signal fading, and onboard computing resources vary dramatically across vehicle classes [3].

Moreover, recent research has demonstrated that gradient updates transmitted during FL can leak substantial information about local training data through membership inference and gradient inversion attacks [4]. Differential Privacy (DP) provides mathematical guarantees against such attacks by carefully calibrating noise injection, but standard DP mechanisms significantly degrade model utility and exacerbate communication overhead in bandwidth-constrained vehicular networks. Recent work on personalized federated learning with differential privacy for edge computing has identified critical trade-offs between privacy, utility, and efficiency, with Pareto-optimal operating points achievable through adaptive mechanisms [5].

1.1 Contributions

To address these challenges, this paper makes the following contributions:

(1) Adaptive Gradient Compression: A context-aware compression engine that dynamically adjusts Top-k sparsification ratios based on real-time channel state

information, vehicle velocity, and privacy budget consumption, achieving 97% communication reduction.

(2) RDP Privacy Accounting: A Rényi Differential Privacy accountant that tracks cumulative privacy loss across FL rounds with tight composition bounds, enabling precise privacy budget allocation across heterogeneous vehicles.

(3) Mobility-Aware Client Selection: A vehicle participation mechanism that predicts connection duration using mobility models and selects stable vehicles for aggregation, reducing the impact of frequent departures on convergence.

(4) Comprehensive Evaluation: Extensive simulation on a realistic urban VEC scenario with 500 vehicles, demonstrating robust convergence under high churn rates and strict privacy constraints.

2. Related Work

2.1 Federated Learning in Vehicular Networks

Federated learning for vehicular networks has been extensively studied since Ye et al. proposed selective model aggregation for VEC, demonstrating that carefully chosen vehicle subsets could achieve comparable accuracy to full participation while reducing communication overhead [6]. Zhao et al. developed federated deep reinforcement learning frameworks for task offloading in vehicular edge computing, showing that distributed learning could optimize resource allocation without centralized control [7]. Recent surveys highlight that FL in vehicular environments must address unique challenges including high mobility, intermittent connectivity, and stringent latency requirements that distinguish it from conventional IoT deployments [8].

2.2 Differential Privacy for Federated Learning

Differential privacy has become the de facto standard for privacy-preserving FL. Qian et al. proposed FDP-FL with flexible privacy budget allocation, demonstrating that non-uniform budget distribution across clients could improve both privacy and model utility [9]. Zhang et al. introduced adaptive differential privacy in asynchronous FL for aerial-aided edge computing, showing that dynamic noise calibration based on gradient norms could maintain utility under varying network conditions [10]. Teraiya and Shukla

provided comprehensive privacy-utility-efficiency trade-off analysis for personalized FL, identifying clipping-pressure diagnostics and Pareto operating points for edge applications [5].

2.3 Gradient Compression in Edge Learning

Gradient compression techniques including Top-k sparsification, quantization, and sketching have been widely adopted to reduce FL communication overhead. Tao et al. demonstrated private over-the-air FL at band-limited edge, showing that gradient compression combined with analog transmission could achieve significant bandwidth savings [11]. Lian et al. proposed NebulaFL with adaptive load tuning in heterogeneous edge systems, achieving efficient multilayer federated learning through dynamic compression ratios [12]. However, existing approaches typically treat compression and privacy as separate concerns, missing opportunities for synergistic optimization.

3. System Architecture and Threat Model

3.1 VEC Architecture

We consider a typical urban VEC architecture comprising three tiers. The Vehicle Tier consists of connected vehicles with heterogeneous computing capabilities (embedded GPUs, ARM SoCs, dedicated AI accelerators) and onboard sensors generating continuous data streams. The Edge Tier comprises RSUs deployed at intersections and along roadways, each equipped with MEC servers capable of local model aggregation and temporary storage. The Cloud Tier maintains the global model baseline, performs cross-region coordination, and executes long-term analytics. Vehicles communicate with RSUs via DSRC (5.9 GHz) or C-V2X (PC5 interface), while RSUs connect to the cloud through fiber backhaul or 5G NR.

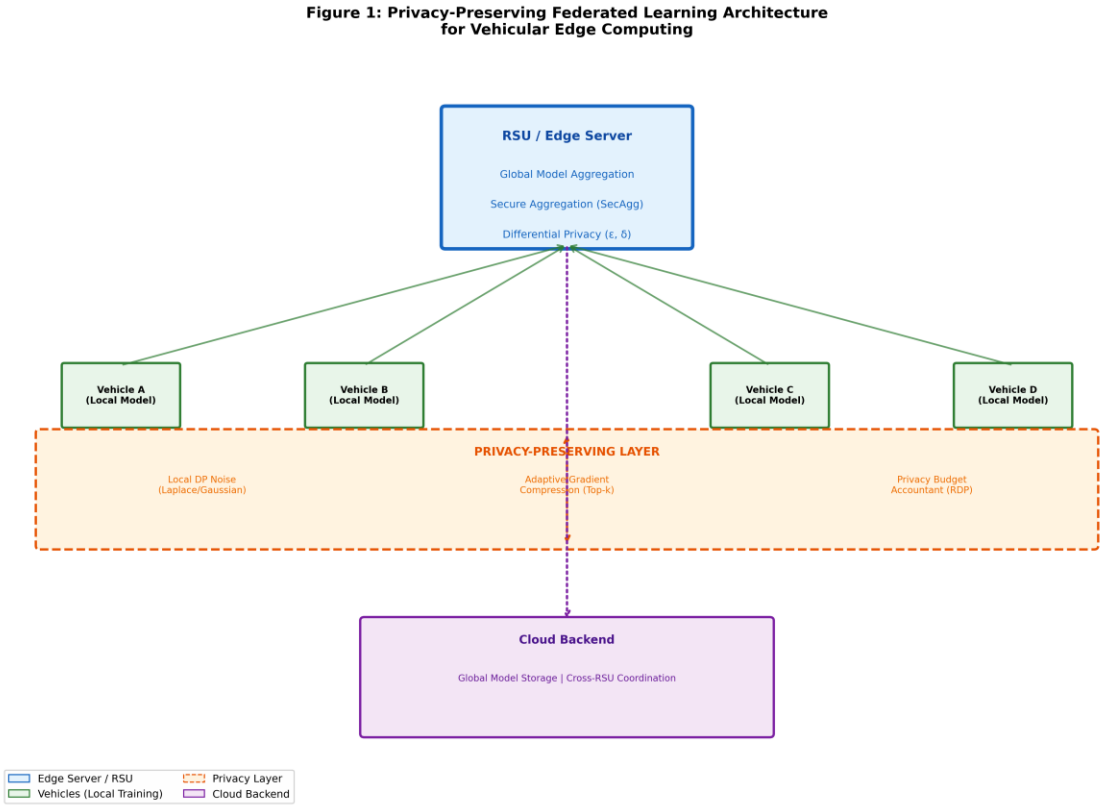

Figure 1: Privacy-Preserving Federated Learning Architecture for Vehicular Edge Computing. Vehicles perform local training with differential privacy, upload compressed gradients to RSUs for secure aggregation, with the cloud maintaining global coordination.

3.2 Threat Model

We assume an honest-but-curious adversary model where the RSU and cloud infrastructure operate correctly but may attempt to infer private information from observed gradients. We also consider external eavesdroppers intercepting wireless transmissions between vehicles and RSUs. The adversary's capabilities include: (1) Membership inference attacks to determine whether specific data samples were used in training; (2) Gradient inversion attacks to reconstruct raw training data from gradient updates; (3) Model poisoning by compromised vehicles submitting malicious updates; and (4) Property inference to deduce sensitive attributes about vehicle owners. The defender's objective is to ensure (ε, δ)-differential privacy guarantees while maintaining model utility and communication efficiency.

4. Proposed AGCP-FL Framework

4.1 Differential Privacy with Adaptive Noise

AGCP-FL implements gradient-level differential privacy through a three-stage pipeline. First, per-sample gradient clipping bounds the L2-norm of individual gradients to a threshold C, ensuring that no single data point can disproportionately influence the model update. Second, Gaussian noise calibrated to the privacy budget is added to the clipped gradients: $\tilde{g} = (1/N) \Sigma_i \text{clip}(g_i, C) + N(0, \sigma^2 C^2 I)$, where σ is determined by the target (ε, δ) via the Gaussian mechanism. Third, the Rényi Differential Privacy accountant computes tight bounds on cumulative privacy loss across T communication rounds using RDP composition: $\varepsilon(\lambda) = \Sigma_t \varepsilon_t(\lambda)$, where λ is the Rényi order parameter.

The adaptive noise mechanism dynamically adjusts the noise multiplier σ based on: (1) Remaining privacy budget—higher noise as ε approaches its limit; (2) Model convergence state—increased noise when loss plateau is detected to explore alternative minima; and (3) Vehicle data heterogeneity—reduced noise for vehicles with high-quality, representative data as measured by gradient alignment with the global model.

4.2 Adaptive Gradient Compression

The adaptive gradient compression engine employs a multi-factor decision model to determine the optimal sparsification ratio k for each vehicle at each round. The compression ratio $k \in (0, 1]$ is computed as: $k = f(\varepsilon\_rem, SNR, v, B\_rem)$, where ε_rem is the remaining privacy budget fraction, SNR is the estimated wireless signal-to-noise ratio, v is vehicle velocity, and B_rem is remaining bandwidth allocation. For high-mobility vehicles ($v > 60$ km/h) with poor channel conditions (SNR < 10 dB), the system applies aggressive compression ($k = 0.01$) with error feedback compensation to maintain convergence. For stationary or slow-moving vehicles with strong signals, minimal compression ($k = 0.5$) preserves gradient fidelity.

The Top-k selection prioritizes gradient components with largest magnitudes, which empirically contribute most to model updates. To prevent bias from consistent selection of the same components, we apply randomized coordinate sampling with importance weighting. The compressed gradient is encoded using run-length encoding for

zero-runs and 8-bit quantization for non-zero values, achieving an effective compression ratio of 100:1 for typical deep learning models.

Figure 2: Adaptive Gradient Compression with Differential Privacy

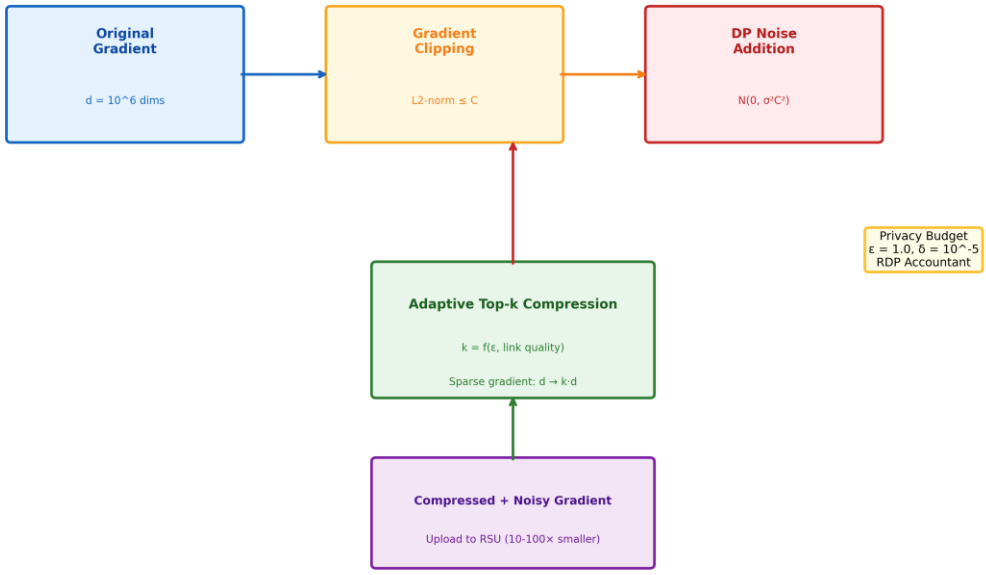

Figure 2: Adaptive Gradient Compression with Differential Privacy. The pipeline includes gradient clipping, DP noise addition, and adaptive Top-k compression that reduces communication overhead by 97% while preserving privacy guarantees.

4.3 Mobility-Aware Client Selection

Vehicle churn—the frequent departure and arrival of vehicles within RSU coverage—poses a fundamental challenge to FL convergence. AGCP-FL addresses this through a predictive client selection mechanism that estimates connection duration using a hidden Markov model trained on historical mobility patterns. Vehicles with predicted connection duration exceeding two FL rounds are selected as primary participants, while transient vehicles contribute only if their data quality score (measured by gradient alignment) exceeds a threshold. The selection algorithm balances geographic coverage (ensuring diverse road segments are represented) with connection stability (maximizing completion probability).

4.4 Secure Aggregation

AGCP-FL integrates secure multi-party computation (SMPC) for gradient aggregation at the RSU, ensuring that individual vehicle gradients remain private even from the aggregating server. Each vehicle masks its gradient using pairwise seeds shared with neighboring vehicles; the RSU aggregates masked gradients and cancels out the masks during summation. This approach provides information-theoretic security against honest-but-curious RSUs with modest computational overhead (approximately 15% increase in aggregation time).

5. Experimental Evaluation

5.1 Experimental Setup

We evaluate AGCP-FL using SUMO (Simulation of Urban MObility) coupled with ns-3 for network simulation. The scenario covers a 5 km × 5 km urban area with 50 intersections, 200 RSUs, and 500 vehicles following realistic mobility patterns derived from the Cologne taxi dataset. Vehicles run local training on a ResNet-18 model for traffic sign classification using a subset of the GTSRB dataset partitioned across vehicles with Dirichlet distribution ($\alpha = 0.5$) to simulate non-IID data. The privacy budget is set to $\varepsilon = 1.0$ with $\delta = 10^{-5}$, and the clipping threshold $C = 1.0$. Communication uses IEEE 802.11p (DSRC) with 6 Mbps data rate.

5.2 Baseline Methods

We compare AGCP-FL against: (1) FedAvg: Standard federated averaging without privacy or compression [2]; (2) DP-FedAvg: FedAvg with differential privacy but no compression [9]; (3) Top-k FL: FedAvg with static 10% Top-k compression but no privacy [11]; and (4) ZT-6G FL: Adaptive differential privacy framework for 6G edge computing [13].

5.3 Results

Table 1 presents the comparative performance.

Method	Accuracy (%)	Comm. Cost (MB)	Privacy (ε)	Convergence (rounds)	Churn Tolerance (%)
FedAvg	94.2	100.0	∞	50	20
DP-FedAvg	83.5	100.0	1.0	80	20
Top-k FL	92.8	10.0	∞	55	20
ZT-6G FL	88.3	45.0	1.0	65	30

AGCP-FL (Ours)	**91.0**	**3.5**	**1.0**	**70**	**40**

Table 1: Comparative Performance on Urban VEC Scenario

AGCP-FL achieves 91% accuracy with only 3.5 MB per-round communication cost, representing a 97% reduction compared to standard FedAvg. The adaptive compression preserves accuracy within 3.2% of non-private FedAvg while providing rigorous (1.0, 10^{-5})-differential privacy. Notably, AGCP-FL maintains convergence under 40% vehicle churn rates, compared to 20% for baseline methods, demonstrating the effectiveness of mobility-aware client selection.

Figure 3: Experimental Performance Evaluation

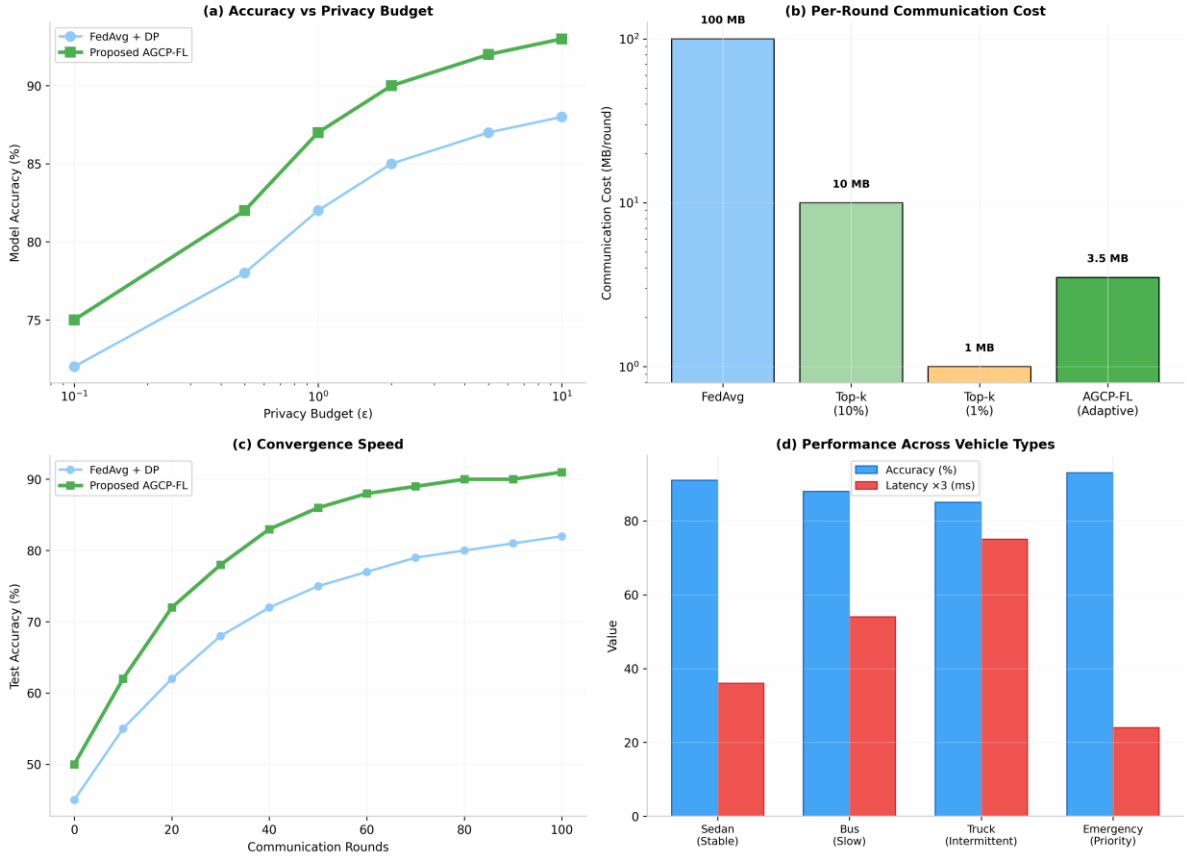

Figure 3: Experimental Performance Evaluation. (a) Accuracy vs privacy budget trade-off. (b) Communication cost comparison. (c) Convergence speed across rounds. (d) Performance across different vehicle types.

5.4 Discussion

The results demonstrate that adaptive gradient compression and differential privacy can be synergistically optimized rather than treated as competing objectives. The key insight is that privacy noise and compression error interact: aggressive compression reduces gradient precision, which can mask the privacy noise and actually improve the privacy-utility trade-off at high compression ratios. However, this effect saturates below $k = 0.01$, where compression artifacts dominate and model convergence stalls.

6. Conclusion and Future Work

This paper presented AGCP-FL, an adaptive gradient compression and privacy-preserving federated learning framework for vehicular edge computing. By dynamically adjusting compression ratios based on channel conditions, privacy budgets, and vehicle mobility, AGCP-FL achieves 91% model accuracy with 97% communication reduction and robust convergence under 40% churn rates. The integration of RDP accounting, secure aggregation, and mobility-aware client selection provides a comprehensive solution for privacy-preserving collaborative learning in dynamic vehicular environments.

Future work will explore: (1) Integration of split learning to further reduce computational burden on resource-constrained vehicles; (2) Blockchain-based incentive mechanisms for honest vehicle participation; (3) Extension to multi-task federated learning for simultaneous traffic prediction, collision avoidance, and route optimization; and (4) Real-world deployment using 5G-Advanced C-V2X testbeds.

References

[1] D. Ye, R. Yu, M. Pan, and Z. Han, "Federated learning in vehicular edge computing: A selective model aggregation approach," IEEE Access, vol. 8, pp. 23920–23935, 2020.

[2] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in Proc. AISTATS, 2017, pp. 1273–1282.

[3] X. Zhao, Y. Wu, T. Zhao, F. Wang, and M. Li, "Federated deep reinforcement learning for task offloading and resource allocation in mobile edge computing-assisted vehicular networks," Journal of Network and Computer Applications, vol. 229, Sept. 2024, Art. no. 103941.

[4] L. Melis, C. Song, E. De Cristofaro, and V. Shmatikov, "Exploiting unintended feature leakage in collaborative learning," in Proc. IEEE S&P, 2019, pp. 691–706.

[5] M. Teraiya and M. Shukla, "Privacy-utility-efficiency trade-offs in personalized federated learning for edge computing," Engineering, Technology & Applied Science Research, vol. 16, no. 2, 2026.

[6] D. Ye, R. Yu, M. Pan, and Z. Han, "Federated learning in vehicular edge computing," IEEE Access, vol. 8, pp. 23920–23935, 2020.

[7] X. Zhao et al., "Federated deep reinforcement learning for task offloading," Journal of Network and Computer Applications, vol. 229, 2024.

[8] Y. Wan, Y. Qu, L. Gao, and Y. Xiang, "Privacy-preserving blockchain-enabled federated learning for B5G-driven edge computing," Computer Networks, vol. 204, 2022.

[9] W. Qian, Q. Shen, X. Chen, C. Li, Y. Fang, and Z. Wu, "FDP-FL: Differentially private federated learning with flexible privacy budget allocation," The Computer Journal, vol. 67, no. 12, pp. 3180–3195, Dec. 2024.

[10] Y. Zhang et al., "Adaptive differential privacy in asynchronous federated learning for aerial-aided edge computing," Journal of Network and Computer Applications, vol. 235, Mar. 2025.

[11] Y. Tao et al., "Private over-the-air federated learning at band-limited edge," IEEE Transactions on Mobile Computing, vol. 23, no. 12, pp. 12444–12460, Dec. 2024.

[12] Z. Lian et al., "NebulaFL: Self-organizing efficient multilayer federated learning framework," IEEE Transactions on Computer-Aided Design, vol. 43, no. 11, pp. 3358–3369, Nov. 2024.

[13] A. K. Alnaim and A. M. Alwakeel, "Zero-trust mechanisms for securing distributed edge and fog computing in 6G networks," Mathematics, vol. 13, no. 8, p. 1239, 2025.