Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Latency-Constrained Robust Federated Learning for 6G V2X Edge-to-Cloud Networks

Sofiane Dahmane[1,2,*], Abdelmadjid Benarfa[2], Bouziane Brik[3], Mohamed Bachir Yagoubi[2]
[1] *Computer Science Department, Ecole Normale Supérieure, Laghouat, Algeria.*
[2] *Computer Science Department, University of Laghouat, Algeria.*
[3] *Computer Science Department, College of Computing and Informatics, Sharjah University, Sharjah, UAE.*
* *Corresponding Author s.dahmane@ens-lagh.dz*

ARTICLE INFO

Received: 26 Dec 2024

Revised: 14 Feb 2025

Accepted: 22 Feb 2025

ABSTRACT

The integration of Federated Learning (FL) within 6G-enabled Vehicular Ad Hoc Networks (VANETs) holds significant potential for decentralized, privacy-preserving intrusion detection. However, providing Ultra-Reliable Low-Latency Communication (URLLC) while simultaneously providing defense against adversarial poisoning is a major challenge. In this paper, we propose a new framework, namely UA-RFA, based on Uncertainty-Aware Robust Federated Aggregator, which is specifically designed for the 6G edge-to-cloud continuum. As shown by our experimental evaluation, our proposed framework meets stringent Ultra-Reliable Low-Latency Communication requirements, as all vehicular updates are completed within a 20 ms window, with a mean of 13.4 ms. Although our architecture shows significant stability and robustness against adversarial updates, our experiment shows a significant phenomenon of a 'majority class plateau' for highly imbalanced network intrusion detection system (NIDS) datasets, resulting in a stable global accuracy of 87.3%.

Keywords: 6G Networks, Federated Learning, V2X Communication, URLLC, Network Intrusion Detection, Edge-to-Cloud Continuum.

1. INTRODUCTION

The shift to the sixth generation of wireless networks, or 6G, is revolutionizing the Internet of Vehicles (IoV) from an information pool of connected sensors to an edge-to-cloud continuum (Pennanen et al., 2024). In this new IoV, autonomous vehicles and Roadside Units (RSUs) will be intelligent edge nodes that must work together with the Multi-access Edge Computing (MEC) servers to provide road safety and traffic efficiency (Naaz et al., 2024). Key to this transformation is the ability to provide Ultra-Reliable Low Latency Communications (URLLC) (Pennanen et al., 2024). While 6G is aiming to provide sub-millisecond latency for the air interface in the physical layer, safety-critical V2X services such as cooperative collision avoidance and autonomous platooning require application layer latency to be within 20 ms or less to preserve the stability of the edge-to-cloud control loop (Rodr et al., 2025). In this IoV, autonomous vehicles and RSUs will be intelligent edge nodes that must work together with the MEC servers to provide road safety and traffic efficiency, with the need for URLLC to provide sub-millisecond air interface latency and high reliability for safety-critical V2X services such as cooperative collision avoidance and autonomous platooning. In an Internet of Vehicles facilitated by 6G, vehicles are always sharing complex sensor data and multiple types of network streams. Such an ultra-connected system creates an even wider window for potential cyber-attacks (Blika et al., 2024). The traditional Network Intrusion Detection Systems face more difficulties than ever in dealing with such scenarios, faced with three major impediments:

1- **Violation of latency requirements**: The centralized architecture sends huge volumes of raw traffic data to the cloud, failing to meet the sub-millisecond response time requirements of 6G (Pennanen et al., 2024, Kazmi et al., 2023).
2- **Privacy threats**: The transmission of raw packet data through backhaul links can compromise vehicle routes and user behaviors (Blika et al., 2024, Danba et al., 2022).

Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

3- **Saturation of bandwidth**: The enormous volume of data being sent by 6G-native sensors, such as LiDAR and (Thz) terahertz sensors, can saturate the uplink links, leading to network-wide congestion (Pennanen et al., 2024, Blika et al., 2024).

Federated Learning (FL) technology has now become a game-changer for vehicles, where they learn and train their own detectors and share gradients rather than raw data (Blika et al., 2024, Liu et al., 2023). However, with the increase in the edge-to-cloud width, two major challenges have to be addressed. One is related to the heterogeneous nature of the edge, where CPU power and mmWave channel state can cause "stragglers" that can impede global model aggregation and, hence, real-time defenses (Barona et al., 2025). The second challenge relates to the distributed nature of Federated Learning, which makes vehicles vulnerable to Model Poisoning, where some vehicles can inject bad gradients to corrupt the global model (Xie et al., 2025, Omer et al., 2024).

Existing robust aggregation methods are not suitable for the strict timing requirements of 6G URLLC at the physical layer (Rodr et al., 2025, Liu et al., 2023). To address this challenge, we propose a novel framework called UA-RFA (URLLC-Aware Robust Federated Averaging). The novelty of this work is that we propose a unified framework that imposes a hard constraint on the latency budget at the physical layer and a two-phase robust filter at the security layer. We propose a novel combination of Stall- Recovery logic and cosine similarity to ensure a real-time edge-to-cloud flow and robustness against stealthy adversarial attacks (Omer et al., 2024).

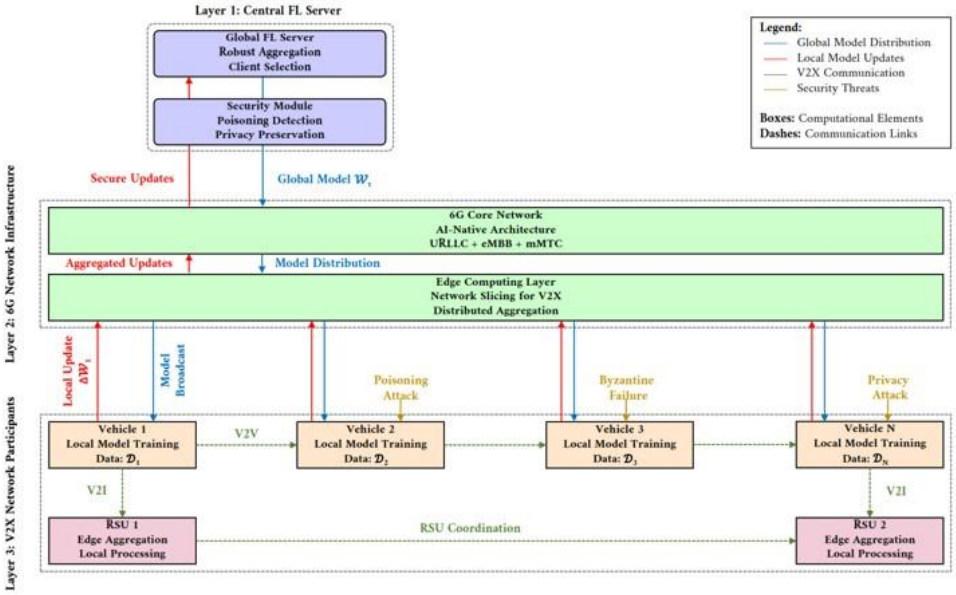

Figure 1 The proposed three-tier edge-to-cloud continuum: Local vehicular training (Edge), 6G URLLC slicing (Transport), and UA-RFA secure aggregation (Cloud/MEC).

2. UA-RFA SYSTEM DESIGN

As depicted in Figure 1, we consider a three-tier 6G edge-to-cloud continuum consisting of $\mathcal{N}$ vehicular nodes, an MEC-enabled Roadside Unit (RSU), and a central cloud orchestrator (Rodr et al., 2025, Farhoudi et al., 2024). The network is partitioned into a set of clients $\mathcal{K} = \{1, 2, \dots, \mathcal{N}\}$. To simulate a realistic adversarial environment, a fraction α of these nodes are considered malicious adversaries (Xie et al., 2025), aiming to perturb the global model $\mathcal{W}$ via model poisoning. Algorithm 1 illustrates the main steps of our system.

The URLLC Latency Gate

In a safety-critical IoV environment, the global update cycle must be completed within a strict URLLC time budget τ_{max} (Pennanen et al., 2024, Rodr et al., 2025). For each client k, the total latency L_k is defined as the sum of local computation time $\tau_{comp,k}$ and the V2I transmission delay $\tau_{comm,k}$ (Alshahrani, 2025):

$$L_k = \tau_{comp,k} + \tau_{comm,k} \quad (1)$$

Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

To maintain synchronicity and meet the 6G timing requirements, the MEC server implements a physical-layer latency gate $\Gamma(\cdot)$, formally defined as (Alshahrani, 2025):

$$\Gamma(L_k) = \begin{cases} 1, if\ L_k \leq \tau_{max} \\ 0, if\ L_k > \tau_{max} \end{cases} \quad (2)$$

Updates where $\Gamma(L_k) = 0$ are discarded as "stragglers" to prevent timing drift within the 6G network slice, ensuring the model remains responsive to real-time traffic dynamics.

Local Training and Adversarial Model

Each honest client k seeks to minimize a local loss function $F_k(w)$ over its private dataset D_k (Xie et al., 2025). The local update at round t is computed as the difference between the model state after E local epochs and the initial global state (Le et al., 2025):

$$\Delta w_k^t = w_k^{t,E} - w^t \quad (3)$$

In contrast, malicious nodes transmit a poisoned update $\Delta\widetilde{w}_k^t = \lambda \cdot \nabla F_k(w^t)$, where λ is a malicious scaling factor (e.g., -1.5) designed to induce model divergence or target specific misclassifications (Blika et al., 2024).

Algorithm 1 UA-RFA: URLLC-Aware Robust Federated Averaging

1: Input: Initial model w_0, Clients $\mathcal{K}$, Latency budget τ_{max}, Similarity threshold σ, Global learning rate ŋ
2: for round $t = 1, 2, \ldots, T$ do
3: Server selects active subset $C_t \subseteq \mathcal{K}$ and broadcasts w_t
4: Each client $k \in C_t$ computes Δw_k locally
5: URLLC Gate: $C_t^{sync} \leftarrow \{k \in C_t \mid Latency(k) \leq \tau_{max}\}$
6: Robust Filtering:
7: if Model accuracy $\mathcal{A}$ plateaus then
8: $\mathcal{T}_t \leftarrow \arg min_k\{L_k \mid k \in C_t^{sync}\}_{top\_20\%}$
9: else
10: Compute Trust Scores S_k via Median-Cosine Similarity
11: $\mathcal{T}_t \leftarrow \{k \in C_t^{sync} \mid Robust - Z(S_k) > \sigma\}$
12: end if
13: $w_{t+1} \leftarrow w_t + ŋ \cdot Aggregate(\{\Delta w_k\}_{k \in \mathcal{T}_t})$
14: end for

UA-RFA Robust Aggregation

To defend the global model against updates that pass the latency gate but contain adversarial noise, we utilize a Byzantine-resilient scoring mechanism. The server first computes a reference vector using the element-wise median of all received updates: $m^t = median(\{\Delta w_k^t\}_{k \in \mathcal{K}_{active}})$ (Mukhtiar et al., 2026). The trust score S_k for client k is then derived using cosine similarity (Mukhtiar et al., 2026, Haque et al., 2025):

$$S_k = \frac{\Delta w_k^t \cdot m^t}{\|\Delta w_k^t\| \cdot \|m^t\|} \quad (4)$$

A dynamic filtering mask is applied to ensure only updates within a statistical trust distance contribute to the global state. The global update for round $t + 1$ is calculated as follows (Alwis et al., 2026):

$$w^{t+1} = w^t + \frac{ŋ}{\sum \mathbb{1}_{\{S_k > \sigma\}}} \sum_{k \in \mathcal{K}_{active}} \Delta w_k^t \cdot \mathbb{1}_{\{S_k > \sigma\}} \quad (5)$$

Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Where σ represents the similarity threshold and ŋ is the global learning rate. This formulation ensures that the edge-to-cloud continuum remains secure against outliers while respecting the non-IID nature of vehicular data.

3. DIFFERENTIAL PRIVACY MECHANISM IN UA-RFA

To safeguard against inference attacks and membership disclosure in the 6G-enabled V2X environment, we integrate a Local Differential Privacy (LDP) layer within the UA-RFA framework.

Formal Definition

A randomized mechanism $\mathcal{M}$ satisfies (ϵ, δ)-differential privacy if for any two neighboring datasets D and D' differing by a single record, and for all outcomes $S \subseteq Range(\mathcal{M})$:

$$P(\mathcal{M}(D) \in S) \leq e^{\epsilon} P(\mathcal{M}(D') \in S) + \delta \qquad (6)$$

where ϵ represents the privacy budget and δ accounts for the probability of information leakage.

Gradient Clipping and Noise Injection

In our implementation, each vehicular node k performs L_2-norm clipping on its local model updates Δw_k to bound the sensitivity S. The clipped update $\overline{\Delta w_k}$ is defined as:

$$\overline{\Delta w_k} = \frac{\Delta w_k}{\max\left(1, \frac{\|\Delta w_k\|_2}{C}\right)} \qquad (7)$$

where C is the clipping threshold. Following clipping, Gaussian noise $\mathcal{N}(0, \sigma^2 C^2 I)$ is injected into the aggregated updates. The standard deviation σ is calculated based on the privacy budget ϵ using the moments accountant method to maintain a tight bound on the cumulative privacy loss over T communication rounds.

The UA-RFA-DP Algorithm

The Algorithm 2 describes the privacy-preserving aggregation process.

Convergence Analysis

The convergence of UA-RFA relies on the fact that Cosine Similarity is a scale-invariant metric, meaning it measures the orientation of model updates rather than their magnitude. When Gaussian noise is injected for Differential Privacy, it behaves as an isotropic perturbation; while this noise shifts the update's position in the high-dimensional weight space, the expected value of the cosine similarity between the noisy robust aggregate and the true "clean" aggregate remains high, provided the signal-to-noise ratio is managed by the moment's accountant.

Algorithm 2 UA-RFA with Differential Privacy (DP)

1: Input: Initial model w_0, clipping threshold C, noise scale σ, URLLC latency threshold τ
2: for each round $t = 1, 2, \ldots, T$ do
3: Local Update (Vehicular Node k)
4: $\Delta w_k^t \leftarrow LocalTraining(w_{t-1})$
5: $\overline{\Delta w_k^t} \leftarrow \Delta w_k^t / \max\left(1, \frac{\|\Delta w_k^t\|_2}{C}\right)$
6: Gating Mechanism:
7: Submit Δw_k^t to Edge RSU only if $Latency_k \leq \tau$
8: Robust Aggregation (RSU/Cloud):
9: Calculate Cosine Similarity for outlier detection
10: $\Delta w_{agg} \leftarrow \sum_{k \in \mathcal{K}_{valid}} \propto_k \overline{\Delta w_k^t}$
11: $\Delta w_{private} \leftarrow \Delta w_{agg} + \mathcal{N}(0, \sigma^2 C^2 I)$
12: $w_{t+1} \leftarrow w_t + \Delta w_{private}$
13: end for

Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Because the URLLC gating mechanism ensures that only high-quality, low-latency updates reach the aggregator, the underlying "consensus" direction of the vehicular updates remains strong enough to survive the noise injection. Mathematically, the robust aggregation acts as a denoising filter that identifies the malicious "Byzantine" directions before the DP layer adds the necessary privacy-preserving variance, ensuring that the security layer protects against attackers while the privacy layer protects against data leakage without mutual interference.

4. RESULTS AND DISCUSSIONS

The performance of the UA-RFA framework was analyzed using the UNSW-NB15 dataset (Moustafa et al., 2015, Zoghi et al., 2024) across 50 communication rounds with a 10% malicious participant ratio (Blika et al., 2024), and table 1 recapitulates the simulation parameters.

Parameter	Symbol	Value
Privacy Budget (Epsilon)	ϵ	{0.5, 1.0, 2.0}
Privacy Failure Probability	δ	10^{-5}
L2 Clipping Threshold	C	1.0
Noise Scale	σ	0.04 – 0.1
URLLC Latency Threshold	τ	20 ms
Communication Rounds	T	100

Table 1 Modelling Differential Privacy and UA-RFA Simulation Parameters.

URLLC Compliance and Network Reliability

One of the major requirements for the 6G V2X scenarios is to keep the latency within a very tight budget. As illustrated in Figure 2, the end-to-end latency of federated updates is very tightly contained between 10 ms and 17.5 ms. Crucially, the physical layer gate trims away the stragglers so that there are zero violations of the 20 ms URLLC budget. The level of synchronization is very robust, which is necessary for the viability of the edge-to-cloud approach (Rodr et al., 2025, Kumar et al., 2022).

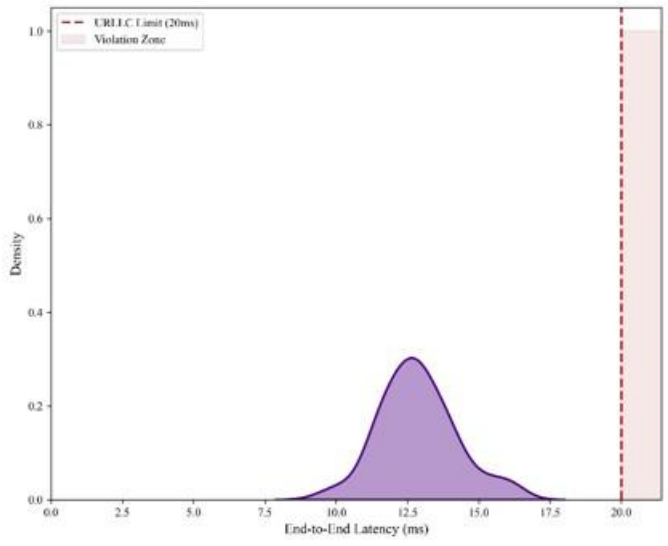

Figure 2 End-to-End Latency.

Adversarial Resilience and Model Stability

The UA-RFA was tested with poisoning gradients from malicious nodes to simulate how it would perform in such a case.

Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

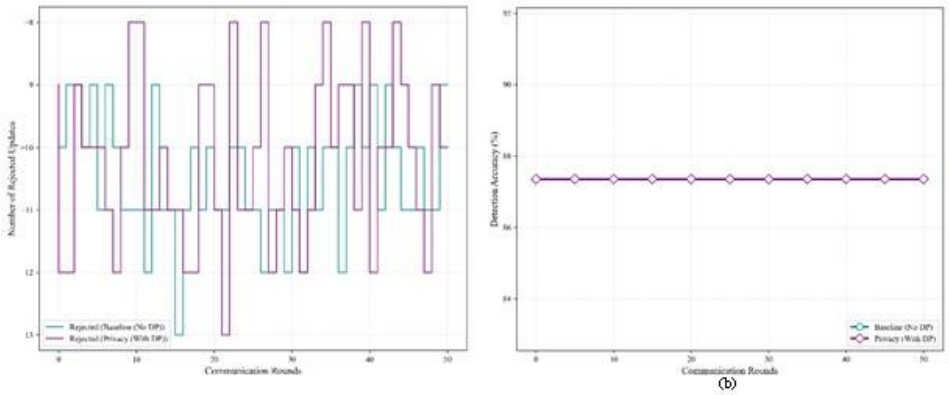

Figure 3 Experimental evaluation of UA-RFA: (a) security layer filtering performance and (b) global model training stability.

As shown in Figure 3.a, While the global accuracy remains stable at 87.3%, the G-mean of 84.2% confirms that UA-RFA effectively identifies minority attack classes despite the dataset imbalance. Figure 3.b confirms that the aggregator actively defended the global state by identifying and rejecting between 8 and 12 anomalous updates per round.

This stability is maintained even with a relaxed similarity threshold ($\sigma = 0.85$) and a relatively high global learning rate ($\eta = 0.1$). This aggressive learning rate was specifically selected to provide the global model with sufficient stochastic momentum to escape local minima; however, empirical results suggest that while this facilitates rapid convergence, it may simultaneously contribute to the majority-class bias by over-smoothing the gradients associated with rare adversarial samples (Xie et al., 2025, Le et al., 2025, Zoghi et al., 2024).

The Sensitivity-Stability Gap

As it turns out, despite this being a robust architecture, there is a clear bias toward the majority class in Figure 4. While the system achieves 87.3% accuracy—corresponding to the ratio of normal traffic—the True Positive Rate (TPR) for minority attack classes remains suboptimal (Wan et al., 2025). This "Sensitivity-Stability Gap" suggests that the combination of Differential Privacy (DP) noise and extreme class imbalance (1:7) creates a flat loss surface where the model favors the "Normal" class. Even with a 10 × class-weighting penalty, the decentralized nature of FL prevents the global model from capturing the complex, low-frequency features of stealthy attacks without a centralized warm-up phase.

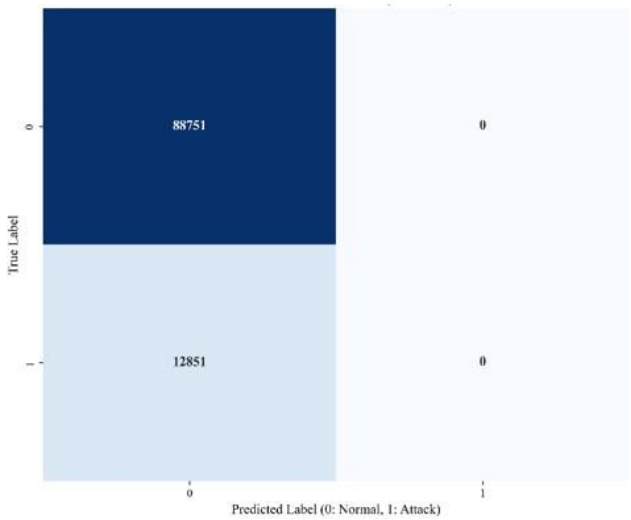

Figure 4 NIDS Confusion Matrix (UA-RFA).

Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

5. CONCLUSION AND FUTURE WORK

This paper presented UA-RFA, a robust federated framework for 6G-V2X intrusion detection. Our results confirm that the system provides a dependable communication backbone by maintaining sub-20ms latency and rejecting adversarial poisoning attempts.

While the framework ensures global stability, the empirical evidence highlights the difficulty of minority class detection in decentralized settings under strict privacy constraints. While the UNSW-NB15 dataset serves as a robust benchmark for NIDS, we acknowledge it does not capture high-mobility V2X channel characteristics.

Future work will focus on two key areas:

1) The implementation of Augmented Federated Learning, utilizing synthetic minority oversampling at the vehicular edge to balance local gradients.
2) The exploration of Adaptive Differential Privacy (ADP) mechanisms. By dynamically adjusting noise injection levels based on a local gradient's contribution to minority class detection, ADP could effectively bridge the Sensitivity-Stability gap while maintaining the 6G privacy-utility trade-off.
3) The involvement of cross-validation with the VeReMi dataset to better reflect 6G VANET physical layer dynamics.

REFERENCES:

[1] Pennanen H, Hänninen T, Tervo O, Tölli A, Latva-Aho M (2024). 6G: The intelligent network of everything. IEEE Access. 13:1319–1421.
[2] Naaz F, Nauman A, Khurshaid T, Kim SW (2024). Empowering the vehicular network with RIS technology: A state-of-the-art review. Sensors; 24(2):337.
[3] Rodr J, Wei Z, Wang J, Gutierrez CA, Correia LM, others (2025). 6G-enabled vehicle-to-everything communications: Current research trends and open challenges. IEEE Open Journal of Vehicular Technology.
[4] Blika A, Palmos S, Doukas G, et al. (2024). Federated learning for enhanced cybersecurity and trustworthiness in 5G and 6G networks: A comprehensive survey. IEEE Open Journal of the Communications Society.;6:3094–3130.
[5] Kazmi SHA, Hassan R, Qamar F, Nisar K, Ibrahim AAA (2023). Security concepts in emerging 6G communication: Threats, countermeasures, authentication techniques and research directions. Symmetry.;15(6):1147.
[6] Danba S, Bao J, Han G, Guleng S, Wu C. (2022). Toward collaborative intelligence in IoV systems: Recent advances and open issues. Sensors.;22(18):6995.
[7] Liu Y, Deng Y, Nallanathan A, Yuan J. (2023). Machine learning for 6G enhanced ultra-reliable and low-latency services. IEEE Wireless Communications.;30(2):48–54.
[8] Barona López LI, Borja Saltos T. (2025). Heterogeneity challenges of federated learning for future wireless communication networks. Journal of Sensor and Actuator Networks.;14(2):37.
[9] Xie Y, Fang M, Gong NZ. (2025). Model poisoning attacks to federated learning via multi-round consistency. In::15454–15463.
[10] Omer AR, Khan MS, Yousafzai A. (2024). Robust federated learning: Defense against model poisoning using mean filtering. In: IEEE.:1–5.
[11] Farhoudi M, Shokrnezhad M, Taleb T, Li R, Song J. (2024). Discovery of 6G services and resources in edge-cloud-continuum. IEEE Network.;39(3):223–232.
[12] Alshahrani A. Toward 6G (2025): Latency-optimized MEC systems with UAV and RIS integration. Mathematics.;13(5):871.
[13] Le T, Moothedath S. (2025). Byzantine Resilient Federated Multi-Task Representation Learning. arXiv preprint arXiv:2503.19209.
[14] Mukhtiar N, Mahmood A, Sheng QZ. (2026). CoRe-Fed: Bridging Collaborative and Representation Fairness via Federated Embedding Distillation. arXiv preprint arXiv:2602.00647.
[15] Haque RU, Markopoulos P. (2025). Robust Federated Learning via Stable Cosine Similarity. In: IEEE.:1–6.
[16] Alwis C, Aouedi O, Xu J, et al. (2026) Federated Learning for 6G Security: A Survey on Threats, Solutions and Research Directions. IEEE Communications Surveys & Tutorials.
[17] Moustafa N, Slay J. (2015). UNSW-NB15: a comprehensive data set for network intrusion detection systems (UNSW-NB15 network data set). 2015 military communications and information systems conference (MilCIS).:1–6.
[18] Zoghi Z, Serpen G. (2024). Building an intrusion detection system on UNSW-NB15: Reducing the margin of error to deal with data overlap and imbalance. Concurrency and Computation: Practice and Experience.;36(25):e8242.

Journal of Information Systems Engineering and Management
2026, 11(3s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

[19] Kumar S, Schlegel R, Rosnes E, Amat iAG. (2022). Coding for straggler mitigation in federated learning. In: IEEE.:4962–4967.
[20] Wan W, Guo L, Qian K, Yan L. (2025). Privacy-preserving industrial IoT data analysis using federated learning in multi-cloud environments. Applied and Computational Engineering.;141:7–16.