

Article

FedGAT: Federated Graph Attention for User Association and Interference Mitigation in C-RAN

Hussein Ali Taleb 

School of Arts & Sciences, Lebanese International University (LIU), P.O. Box 146404, Beirut 1107 2020, Lebanon; hussein.taleb@liu.edu.lb

Abstract

This paper presents FedGAT, a federated graph attention network for joint user association and inter-remote radio head (RRH) interference mitigation in cloud radio access networks (C-RAN). Centralized optimization requires global channel state information (CSI) at the baseband unit (BBU) pool, which adds fronthaul overhead and exposes user data. Existing federated learning schemes use flat local models that do not represent the interference graph, which lowers performance in heterogeneous multi-cell settings. In FedGAT, each RRH builds a local interference graph from its own CSI and trains a GATv2 model by local Adam gradient steps. Only the parameter increments are sent to the BBU pool, which combines them by weighted FedAvg without exchanging raw CSI. We formulate the joint user association and power allocation problem as a mixed-integer non-convex program and derive a graph neural network (GNN)-based continuous relaxation suitable for distributed training. A non-asymptotic convergence bound is obtained for non-identically distributed local interference graphs, showing that the optimality gap grows with the local step count and a measure of graph heterogeneity. Under the 3GPP Urban Macrocell channel model, FedGAT increases the steady-state weighted sum-rate by 18.7% over a federated multilayer perceptron baseline and by 49.2% over a centralized Graph Attention Network version 2 (GATv2) model at equal training budgets, while keeping raw CSI local. These margins are averaged over ten Monte Carlo channel realizations and reported with their standard deviations, and the fronthaul benefit of FedGAT refers to privacy and CSI-free inference after training rather than training-phase traffic.

Keywords: cloud radio access network; federated learning; graph attention network; interference mitigation; resource allocation; user association; 5G heterogeneous networks; GATv2; convergence analysis



Academic Editors: Dejan Drajic,
Zoran Cica and Philipp Svoboda

Received: 22 June 2026

Revised: 17 July 2026

Accepted: 18 July 2026

Published: 7 August 2026

Copyright: © 2026 by the author.
Licensee MDPI, Basel, Switzerland.
This article is an open access article
distributed under the terms and
conditions of the [Creative Commons
Attribution \(CC BY\) license](https://creativecommons.org/licenses/by/4.0/).

1. Introduction

The cloud radio access network (C-RAN) is a common architecture for fifth-generation (5G) network densification [1]. It centralizes baseband processing at a baseband unit (BBU) pool that connects to distributed remote radio heads (RRHs) through high-capacity fronthaul links, which supports network-wide coordinated multipoint (CoMP) transmission and centralized interference management. As deployment density grows toward beyond-5G (B5G) and sixth-generation (6G) systems, this architecture faces two main bottlenecks.

The first bottleneck is fronthaul capacity. Centralized resource allocation, such as the weighted minimum mean-square error (WMMSE) [2], requires each RRH to send instantaneous full channel state information (CSI) to the BBU pool at every scheduling slot. The signaling volume scales as $N \times K \times S \times M$ complex coefficients per slot, where

N , K , S , and M denote the numbers of RRHs, user equipment units (UEs), sub-bands, and antennas per RRH, respectively. In dense deployments, this transfer consumes a large part of the fronthaul capacity that the optimization is meant to use efficiently [3]. Sending raw CSI also raises privacy concerns, because channel measurements can reveal user location and mobility.

The second bottleneck is computational complexity. The joint user association and power allocation problem is NP-hard in its mixed-integer form. Iterative solvers such as WMMSE have $\mathcal{O}(K^3)$ per-iteration complexity, which makes real-time use difficult in large deployments. Graph neural networks (GNNs) are a lower-complexity option for interference-limited resource allocation [4,5]. When the interference channel is modeled as a weighted graph whose nodes are UEs and whose edge weights encode pairwise interference, a GNN can reach near-WMMSE performance at polynomial inference complexity. Most existing GNN-based methods are still centralized, so they keep the CSI aggregation burden at the BBU pool.

Federated learning (FL) [6] trains a shared model across distributed nodes without disclosing raw measurements. Each node trains a local model on its own data and sends only parameter updates to a central server. Three properties make FL attractive for C-RAN resource allocation. First, channel measurements stay at the RRHs, which protects user location information and follows the data minimization principle. Second, the exchanged parameter updates have a size that depends only on the model, not on the number of UEs, sub-bands, or antennas, so the signaling does not grow with the network load. Third, the aggregation step lets every RRH benefit from the interference patterns observed by the other RRHs without building a central dataset. FL has been applied to RAN tasks such as channel estimation [7] and distributed power control [8]. The local models in these works are usually multilayer perceptrons (MLPs) or convolutional neural networks (CNNs), which do not represent the interference graph of the radio environment. As the simulations in this paper show, this leads to policy degeneracy in heterogeneous multi-cell scenarios.

In this paper, we propose FedGAT, which combines GNN-based resource allocation with FL. Each RRH constructs a local interference graph from its own channel measurements and trains a graph attention network version 2 (GATv2) model [9] over τ local gradient steps. Only the parameter increments are sent to the BBU pool, which performs weighted FedAvg aggregation. Raw CSI never crosses the fronthaul.

The main contributions of this work are as follows:

- We formulate the joint user association and power allocation problem as a mixed-integer non-convex program (P1) and derive a GNN-based continuous relaxation that is differentiable and suitable for distributed stochastic optimization.
- We specify the FedGAT algorithm, in which GATv2 local models are deployed at the RRHs and aggregated at the BBU pool by weighted FedAvg, with fronthaul traffic limited to model parameter increments.
- We establish a non-asymptotic convergence bound under non-i.i.d. local interference graphs by extending the FedProx analysis [10] to GATv2 local objectives, and we quantify the trade-off between local computation τ and inter-RRH graph heterogeneity Γ .
- We evaluate FedGAT under 3GPP urban macrocell (UMa) channel conditions against WMMSE, a centralized GATv2 baseline (C-GAT), FedAvg-MLP, and a greedy reference over multiple Monte Carlo channel realizations.

The rest of the paper is organized as follows. Section 2 reviews related work. Section 3 presents the system model, the FedGAT algorithm, and the convergence analysis. Section 4 reports the simulation results. Section 5 discusses the main findings and the practical scope of the method. Section 6 concludes the paper.

2. Related Work

2.1. GNN-Based Radio Resource Allocation

Graph neural networks have emerged as a powerful framework for radio resource allocation, offering permutation-equivariant policies that scale naturally to varying network sizes. The original graph attention network (GAT) [11] introduced attention-weighted neighborhood aggregation as a general-purpose graph learning primitive, enabling each node to assign learned attention coefficients to its neighbors. Building on this foundation, Eisen and Ribeiro [4] represented the power allocation policy with a random edge graph neural network (REGNN) and learned it without labeled data, obtaining a sum-rate close to that of the WMMSE algorithm at a lower computational cost. Sun et al. [12] demonstrated that deep neural networks can directly approximate the input-output mapping of iterative interference management algorithms, establishing the viability of learned resource allocation for real-time operation. Shen et al. [5] proved that message-passing GNNs are permutation equivariant when applied to interference graphs, which explains their scalability to large and varying network topologies. Chowdhury et al. [13] combined algorithm unrolling with GNNs by unfolding WMMSE iterations into a trainable graph network, achieving near-optimal power allocation with a compact, interpretable architecture. The GATv2 model [9] extended the original GAT [11] by replacing its static, query-independent attention mechanism with a dynamic, fully expressive attention function, yielding strictly superior representational capacity. All of these methods are trained and deployed at a central node, requiring global CSI, which creates prohibitive fronthaul overhead and privacy exposure in large-scale C-RAN deployments.

2.2. C-RAN Architecture and User Association

The C-RAN paradigm and its resource management challenges have been extensively studied. Peng et al. [14] provided a comprehensive survey of C-RAN system architectures and fronthaul compression techniques, highlighting the scalability limits imposed by fronthaul capacity constraints as deployment density grows toward beyond-5G and 6G. The joint user association and interference management problem is a classical combinatorial challenge in heterogeneous networks. Ye et al. [15] formulated user association as a network-wide utility maximization problem and showed that biased-SINR association with a dual decomposition solver approaches the NP-hard globally optimal solution, providing the theoretical foundation that motivates the mixed-integer formulation of FedGAT. The need for global channel knowledge in such centralized solvers underscores the value of a distributed, learning-based approach.

2.3. Federated Learning for Wireless Networks

Federated learning [6] trains a shared model across distributed nodes without disclosing raw measurements by exchanging only model parameter updates. Li et al. [10] studied FL convergence under non-i.i.d. data and established an $\mathcal{O}(1/\sqrt{T})$ rate for the FedProx algorithm, which we adopt as the analytical basis of our convergence result. Niknam et al. [16] surveyed the motivations, opportunities, and key challenges of applying FL to 5G wireless systems, noting that privacy preservation and reduced fronthaul signaling are among its most compelling advantages for RAN applications. Lim et al. [17] provided a comprehensive treatment of FL for mobile edge networks, identifying communication overhead and device heterogeneity as the principal implementation barriers. In wireless systems, Elbir and Coleri [7] applied FL to channel estimation in RIS-assisted massive MIMO, and Tehrani et al. [8] applied federated deep reinforcement learning to distributed power control in multi-cell networks. Both works use MLP or CNN local models that do

not represent the interference graph structure and therefore cannot leverage the topology of the radio environment.

2.4. DRL and Multi-Agent Learning for Resource Allocation

Beyond supervised and unsupervised GNN training, deep reinforcement learning (DRL) is a major line of work in radio resource management. Luong et al. [18] surveyed DRL applications across communications and networking, including dynamic spectrum access, rate control, caching, and offloading. Nasir and Guo [19] developed a multi-agent DRL scheme in which each transmitter selects its power level from delayed local observations and approaches the WMMSE sum-rate without instantaneous global CSI. Liang et al. [20] applied multi-agent reinforcement learning to spectrum sharing in vehicular networks, where vehicle-to-vehicle links learn to reuse the spectrum of vehicle-to-infrastructure links in a distributed way. Wang et al. [21] proposed a hierarchical DRL method for resource allocation in the cognitive Internet of Vehicles under unknown channel state information. In air-ground networks, Mozaffari et al. [22] provided a tutorial on UAV communications that covers deployment, channel modeling, and resource management, and Ding et al. [23] proposed an optimization-driven DRL framework for resource allocation in licensed and unlicensed UAV spectrum sharing networks under uncertain jamming, where a model-based optimization step reduces the training complexity of the learning agent. Recent works push this line further. Eldeeb and Alves [24] proposed an offline multi-agent reinforcement learning framework for radio resource management. The agents are trained on fixed datasets with conservative Q-learning. Their centralized training with decentralized execution variant, denoted CTDE-MARL-CQL, gives the best trade-off between complexity and performance, and it improves the weighted combination of average and tail rates by more than 15% over the baselines. Wang et al. [25] presented IRL-D3QN, a multi-agent framework for dynamic spectrum management in vehicular networks. The framework combines inverse reinforcement learning with a dueling double deep Q-network, so the reward function is recovered from expert demonstrations instead of being fixed by hand. Zhang and Guo [26] formulated multi-cell spectrum and power allocation as a distributed learning problem. They solved it with a multi-agent proximal policy optimization (MAPPO) algorithm with recurrent networks and queueing dynamics. Each agent controls one cell from local information and minimizes packet delays under stochastic traffic arrivals. Wang et al. [27] surveyed multi-agent deep reinforcement learning (MADRL) for wireless spectrum management and organized the existing schemes over power control, spectrum sensing, spectrum allocation, and channel access. These methods learn control policies by interaction and typically operate on flat state vectors or per-agent observations. FedGAT differs in two respects: it encodes the interference topology explicitly through graph attention, and it trains by federated aggregation of parameter increments, so neither raw observations nor experience samples leave the RRHs.

2.5. Federated GNNs and the Gap Addressed by FedGAT

The combination of FL with GNNs has been explored outside the wireless domain. Wu et al. [28] proposed FedPerGNN, a federated graph neural network framework for privacy-preserving personalization, where the clients train GNN models on local user item graphs and the server aggregates the gradients under local differential privacy, and Liu et al. [29] provided a comprehensive survey of federated graph learning techniques and challenges across multiple application domains. In the wireless domain, Feys et al. [30] applied a graph neural network to quantized precoding optimization in massive MIMO, but training was carried out centrally with full global CSI. To our knowledge, FedGAT is

the first method that trains GATv2 models locally at the RRHs and aggregates them with FedAvg for joint user association and interference coordination in C-RAN.

3. Materials and Methods

3.1. System Model

We consider the downlink of a C-RAN comprising N RRHs $\mathcal{N}=\{1, \dots, N\}$ and K single-antenna UEs $\mathcal{K}=\{1, \dots, K\}$. The total system bandwidth B (Hz) is partitioned into S orthogonal sub-bands. Each RRH $n \in \mathcal{N}$ is equipped with M transmit antennas and is connected to the BBU pool through a dedicated fronthaul link. The system architecture is depicted in Figure 1.

The received signal at UE k on sub-band s is:

$$y_k^s = \sum_n \mathbf{h}_{nk}^s \mathbf{w}_{nk}^s x_k^s + \sum_{j \neq k} \sum_n \mathbf{h}_{nj}^s \mathbf{w}_{nj}^s x_j^s + z_k^s, \quad (1)$$

where $\mathbf{h}_{nk}^s \in \mathbb{C}^{1 \times M}$ is the channel vector from RRH n to UE k , $\mathbf{w}_{nk}^s \in \mathbb{C}^M$ is the beamforming vector, x_k^s is the transmitted data symbol intended for UE k on sub-band s , normalized to unit power ($\mathbb{E}[|x_k^s|^2] = 1$), and $z_k^s \sim \mathcal{CN}(0, \sigma^2)$. Channels are generated according to the 3GPP UMa path-loss model (TR 38.901) with superimposed Rayleigh small-scale fading. The signal-to-interference-plus-noise ratio (SINR) at UE k on sub-band s under maximum-ratio transmission (MRT) is:

$$\text{SINR}_k^s = \frac{\sum_n \beta_{nk} p_{nk}^s g_{nk}^s}{\sum_{j \neq k} \sum_n \beta_{nj} p_{nj}^s g_{nk}^s + \sigma^2}, \quad (2)$$

where $g_{nk}^s = M^{-1} \|\mathbf{h}_{nk}^s\|^2$ is the effective MRT gain, $\beta_{nk} \in \{0, 1\}$ is the binary user association indicator, and p_{nk}^s is the transmit power. The achievable rate of UE k is $R_k = \sum_s \frac{B}{S} \log_2(1 + \text{SINR}_k^s)$.




Figure 1. C-RAN system architecture for FedGAT ($N = 3$ RRHs, $K = 15$ UEs, BBU pool as FL aggregation server).

At each scheduling slot, RRH n constructs a local interference graph $\mathcal{G}_n = (\mathcal{V}_n, \mathcal{E}_n)$ using only locally available CSI. The vertex set $\mathcal{V}_n$ comprises all K UEs. The feature vector of node k is:

$$\mathbf{f}_k = [g_{nk}, \overline{\text{SINR}}_k, \beta_k, d_{nk}/d_{\max}, Q_k/Q_{\max}]^T \in \mathbb{R}^5, \quad (3)$$

where g_{nk} is the sub-band-averaged effective gain, $\overline{\text{SINR}}_k$ the sub-band-averaged SINR at the previous slot, d_{nk} the Euclidean distance between UE k and RRH n , and Q_k the queue occupancy. A directed interference edge $(k, j) \in \mathcal{E}_n$ is instantiated whenever the normalized interference metric $a_{kj} = g_{nk}g_{nj}/\sqrt{g_{nk}^2 + g_{nj}^2}$ satisfies $a_{kj} \geq \gamma_{\text{th}}$.

The joint optimization problem is:

$$(P1) \quad \underset{\{\beta_{nk}\}, \{p_{nk}^s\}}{\text{maximize}} \quad \sum_k \mu_k R_k \quad (4)$$

$$\text{s.t.} \quad \sum_k \sum_s p_{nk}^s \beta_{nk} \leq P_n^{\max}, \quad \forall n, \quad (5)$$

$$\sum_n \beta_{nk} = 1, \quad \forall k, \quad (6)$$

$$\beta_{nk} \in \{0, 1\}, \quad p_{nk}^s \geq 0,$$

where μ_k is the quality of service (QoS) weight and $P_n^{\max}$ is the per-RRH power budget. Problem (P1) is NP-hard owing to the binary association variables and the non-convex SINR coupling in the objective. FedGAT learns a distributed continuous relaxation of this problem, as described in Section 3.2.

3.2. The FedGAT Framework

Each RRH n operates a two-layer GATv2 network [9] with trainable parameters θ_n . At layer ℓ , the embedding of node k is updated as:

$$\mathbf{h}_k^{(\ell+1)} = \phi \left(\sum_{j \in \mathcal{N}_k \cup \{k\}} \alpha_{kj}^{(\ell)} \mathbf{W}^{(\ell)} \mathbf{h}_j^{(\ell)} \right), \quad (7)$$

where $\mathcal{N}_k$ is the interference neighborhood of k in $\mathcal{G}_n$, $\mathbf{W}^{(\ell)}$ is a trainable weight matrix, ϕ denotes the ReLU activation, and $\alpha_{kj}^{(\ell)}$ is the normalized GATv2 attention coefficient computed via:

$$e_{kj}^{(\ell)} = \mathbf{a}^{(\ell)\top} \text{LeakyReLU}(\mathbf{W}_1^{(\ell)} \mathbf{h}_k^{(\ell)} + \mathbf{W}_2^{(\ell)} \mathbf{h}_j^{(\ell)}). \quad (8)$$

The score $e_{kj}^{(\ell)}$ is the unnormalized attention value and is distinct from $\alpha_{kj}^{(\ell)}$: the coefficient used in Equation (7) is obtained by the softmax normalization $\alpha_{kj}^{(\ell)} = \exp(e_{kj}^{(\ell)}) / \sum_{j'} \exp(e_{kj'}^{(\ell)})$, where the sum runs over $j' \in \mathcal{N}_k \cup \{k\}$.

After two message-passing iterations, two output heads produce a soft association score $\hat{\beta}_{nk} \in (0, 1)$ and power fractions $\hat{p}_{nk}^s \in [0, P_n^{\max}]$ for each UE.

The local training objective at RRH n is:

$$\begin{aligned} \mathcal{L}_n(\theta_n) = & - \sum_k \mu_k \hat{R}_k + \lambda_1 [\max(0, \sum_k \hat{p}_k - P_n^{\max})]^2 \\ & + \lambda_2 \frac{1}{K} \sum_k \hat{\beta}_{nk} (1 - \hat{\beta}_{nk}), \end{aligned} \quad (9)$$

where $\hat{R}_k = \sum_s \log_2(1 + \hat{\beta}_{nk} \hat{p}_{nk}^s \gamma_k^s)$ with $\gamma_k^s \triangleq \text{SINR}_k^s$. The second term penalizes violation of the per-RRH power budget (5); the third term drives $\hat{\beta}_{nk}$ toward binary values.

Algorithm 1 and Figure 2 describe the FedGAT training protocol. In each global round t , the BBU pool broadcasts the global parameter vector θ^t to all N RRHs. Each RRH n initializes its local parameters at θ^t and executes τ Adam gradient steps on $\mathcal{L}_n$. Only the

increment $\Delta\theta_n^t = \theta_n^{t,\tau} - \theta^t$ is transmitted to the BBU pool over the uplink fronthaul; raw CSI is never disclosed. The BBU pool aggregates the received increments via weighted FedAvg:

$$\theta^{t+1} = \theta^t + \sum_n \frac{|\mathcal{D}_n|}{|\mathcal{D}|} \Delta\theta_n^t, \quad (10)$$

where $|\mathcal{D}_n|$ is the local dataset cardinality at RRH n and $|\mathcal{D}| = \sum_m |\mathcal{D}_m|$.

Algorithm 1 FedGAT Training Protocol

Require: N RRHs, T rounds, τ local steps, lr η , λ_1 , λ_2

```

1: Initialize  $\theta^0$  at BBU pool
2: for  $t = 0, 1, \dots, T - 1$  do
3:   BBU broadcasts  $\theta^t$  to all  $N$  RRHs
4:   for each RRH  $n \in \mathcal{N}$  in parallel do
5:     Construct  $\mathcal{G}_n^t$  from local CSI
6:      $\theta_n^{t,0} \leftarrow \theta^t$ 
7:     for  $i = 0, \dots, \tau - 1$  do
8:        $\theta_n^{t,i+1} \leftarrow \theta_n^{t,i} - \eta \nabla \mathcal{L}_n(\theta_n^{t,i})$ 
9:     end for
10:    Upload  $\Delta\theta_n^t = \theta_n^{t,\tau} - \theta^t$  ▷ no raw CSI transmitted
11:  end for
12:   $\theta^{t+1} \leftarrow \theta^t + \sum_n \frac{|\mathcal{D}_n|}{|\mathcal{D}|} \Delta\theta_n^t$ 
13: end for
14: return  $\theta^T$ 
  
```




Figure 2. FedGAT algorithm flowchart.

The fronthaul traffic per round comprises $|\theta|$ floating-point values per RRH in the uplink and one broadcast of $|\theta|$ values in the downlink, independent of K , S , or M . Once training converges, inference is executed locally at each RRH with no further fronthaul signaling.

3.3. Convergence Analysis

We establish the convergence of FedGAT under the following standard conditions.

Assumption 1. Each local objective $\mathcal{L}_n$ is L -smooth.

Assumption 2. The stochastic gradient variance satisfies $\mathbb{E}[\|\nabla \mathcal{L}_n(\theta; \xi) - \nabla \mathcal{L}_n(\theta)\|^2] \leq \sigma_g^2$.

Assumption 3. The expected squared norm of each local gradient is bounded, $\mathbb{E}[\|\nabla \mathcal{L}_n(\theta)\|^2] \leq G^2$.

Assumption 4. The inter-RRH graph heterogeneity measure $\Gamma = \mathcal{L}(\theta^*) - \sum_n p_n \mathcal{L}_n(\theta_n^*) \geq 0$ is bounded, where θ^* is the global minimizer and $p_n = |\mathcal{D}_n|/|\mathcal{D}|$. The condition $\Gamma=0$ holds if and only if all local interference graphs are identically distributed.

Theorem 1. Under Assumptions 1–4, with step size $\eta \leq 1/(4L\tau)$, the iterates of Algorithm 1 satisfy

$$\begin{aligned} & \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|\nabla \mathcal{L}(\theta^t)\|^2] \\ & \leq \frac{2[\mathcal{L}(\theta^0) - \mathcal{L}^*]}{\eta\tau T} + \frac{4\eta L\sigma_g^2}{N} + 4\eta^2 L^2 \tau^2 (2\Gamma G^2 + \sigma_g^2). \end{aligned} \quad (11)$$

Sketch of the proof of Theorem 1. Applying the L -smoothness descent lemma to the FedAvg update $\theta^{t+1} = \theta^t + \sum_n p_n \Delta \theta_n^t$ and bounding the client drift over τ local steps by recursively unrolling the gradient iterations under Assumptions 2–4 yields $\mathbb{E}[\|\Delta \theta_n^t\|^2] \leq \eta^2 \tau^2 (2\Gamma G^2 + \sigma_g^2)$. Telescoping over T rounds and dividing by T under the step-size constraint $\eta \leq 1/(4L\tau)$ produces (11). The derivation follows Theorem 4 of [10] adapted to GATv2 local objectives. The complete proof is given in Appendix A. $\square$

Corollary 1. The second term in (11) decays as $1/N$: adding more participating RRHs reduces the stochastic gradient variance and accelerates convergence. The third term scales as $\tau^2 \Gamma$: under conditions of high inter-RRH graph heterogeneity, the local step count τ must be kept small to contain the client drift. The empirically suitable range is $\tau^* \in \{5, 10\}$ for the configurations evaluated in Section 4.

Two clarifications delimit the scope of this result. First, Theorem 1 bounds the averaged squared gradient norm of the relaxed objective $\mathcal{L}$, so it guarantees convergence to a stationary point of the continuous surrogate problem, not optimality of the original mixed-integer program (P1). The binary association is recovered by rounding the soft scores, and the third loss term drives these scores toward binary values, but no optimality certificate for (P1) is claimed. Second, the bound describes the training dynamics; it does not by itself quantify the sum-rate gap to the WMMSE reference reported in Section 4.

4. Simulation Results

All simulations are implemented in PyTorch 2.2 and PyTorch Geometric 2.5 (Python 3.10) and trained using the Adam optimizer with learning rate $\eta = 5 \times 10^{-4}$ (Simulation code is available upon request from the corresponding author). Each configuration is evaluated

over ten independent Monte Carlo channel realizations and the reported results are averaged. The main parameters are summarized in Table 1, including the number of RRHs $N \in \{3, 6, 9\}$, the number of UEs K , the number of sub-bands $S = 4$, the per-RRH power budget $P_n^{\max} = 33$ dBm, and the 3GPP UMa channel model. We report the standard deviation of every scheme in Table 2 and show error bars in Figure 3, and we treat differences whose error bars overlap as indicative rather than conclusive. A larger campaign with confidence intervals and formal significance tests is left for future work. The number of global rounds differs across experiments on purpose: in each experiment, T is set to the smallest budget after which the learning curves flatten in the corresponding configuration, which keeps the total training cost of four learning schemes over all configurations manageable. Within any single figure, all schemes use the same T , so every comparison is made at equal budget.




Figure 3. Weighted sum-rate with $\pm$ one standard deviation ($N = 3$, $K = 15$, $\tau = 5$).

Table 1. Simulation parameters.

Parameter	Value
Number of RRHs N	3, 6, 9
Number of UEs K	15 (baseline), $K = 5N$
Antennas per RRH M	4
Sub-bands S	4
Per-RRH power budget $P_n^{\max}$	33 dBm
System bandwidth B	10 MHz
Channel model	3GPP UMa (TR 38.901)
Monte Carlo realizations	10
GATv2 layers	2
Embedding dimension	64
Attention heads	4
Local Adam steps τ	5
Global FL rounds T	25–40

FedGAT is compared with four reference schemes. The first is WMMSE, which performs centralized water-filling power allocation with nearest-RRH association and full CSI at the BBU pool, and serves as an upper bound. The second is C-GAT, a centralized GATv2 model with the same architecture as FedGAT trained at the BBU pool using global CSI. The third is FedAvg-MLP, a federated multilayer perceptron with the same communication overhead as FedGAT but without graph structure. The fourth is a greedy scheme with nearest-RRH association and uniform power allocation, which requires no training and serves as a lower-complexity reference.

4.1. Convergence Behavior

Figure 3 shows the evolution of the weighted sum-rate versus the global FL round for $N = 3$, $K = 15$, and $\tau = 5$, with $\pm$ one standard deviation error bars over the ten Monte Carlo seeds. WMMSE remains approximately constant around 11.75 b/s/Hz, as it recomputes the solution at every slot with full CSI and does not involve learning. FedAvg-MLP converges rapidly to about 5.99 b/s/Hz by round 10 and then remains essentially unchanged, which indicates that the flat local model settles to a nearly uniform association policy insensitive to channel realizations. C-GAT declines from the common starting point over the first 25 rounds and stabilizes around 4.76 b/s/Hz with pronounced cross-seed variance, owing to the difficulty of training a single centralized graph model under heterogeneous interference conditions from multiple RRHs. We validated this baseline before drawing any comparison. C-GAT uses the same two-layer GATv2 architecture, embedding dimension, attention heads, optimizer, and learning rate as FedGAT, and it trains on the pooled global interference graph, so it receives strictly more information than any single RRH. Its weakness is structural rather than an implementation artifact: the global loss sums terms from RRH neighborhoods with different link densities and gain statistics, so gradient directions that improve one neighborhood degrade another, and the single shared model keeps moving between these regimes instead of settling. FedGAT avoids this conflict because each RRH fits its own graph and the BBU pool only averages parameter increments. FedGAT starts from the same initialization as the learning-based baselines, moves through an exploration phase during the first half of the horizon, and then improves steadily, reaching 7.10 b/s/Hz at steady state. Throughout the paper, a steady-state value is the average over the last ten global rounds. The FedGAT curve is still rising at the end of the horizon, so this average is a conservative summary of the level reached within the fixed training budget, and the residual round-to-round variation comes from the stochastic local gradients. This corresponds to gains of 18.7% over FedAvg-MLP and 49.2% over C-GAT at equal training budgets. Given the overlap of the error bars, these margins are indicative; Table 2 reports the corresponding standard deviations.

The confidence intervals around the WMMSE and greedy curves remain narrow, confirming stable behavior across realizations. For FedAvg-MLP, the variance is negligible, consistent with its channel-independent association policy. C-GAT displays high variance across rounds, which is consistent with its unstable behavior across channel realizations. FedGAT exhibits a large standard deviation of about ± 4.11 b/s/Hz, which shows that it adapts to channel realizations and does not collapse to a single, realization-independent policy. The mean FedGAT curve lies below the FedAvg-MLP curve during the early rounds and crosses above it in the last third of the horizon; the overlap of the error bars means this margin should be read as indicative rather than statistically conclusive.

4.2. Scalability with the Number of RRHs

Figure 4 depicts the steady-state sum-rate as a function of $N \in \{3, 6, 9\}$ with $K = 5N$ and a fixed training budget of $T = 25$ rounds. WMMSE and greedy exhibit approximately linear scaling with N , which is consistent with the increased spatial reuse. The absolute sum-rate of FedGAT and C-GAT still increases with N , but their growth is clearly slower than the linear scaling of WMMSE and greedy, so both learning schemes lose ground at $N = 9$. The reason is that the fixed number of rounds becomes insufficient when the interference graph gets denser; this behavior is aligned with the $\tau^2 T$ drift term in the convergence analysis. The C-GAT mean lies above the FedGAT mean for $N \geq 6$, but its standard deviation spans several b/s/Hz at these network sizes, so this ordering is not stable across channel realizations. FedAvg-MLP remains essentially flat around 5.8 b/s/Hz

for all N : since the offered traffic grows linearly with N , a flat sum-rate means that its uniform association policy fails to exploit the additional spatial reuse as the network grows.




Figure 4. Steady-state sum-rate versus N ($K = 5N$, $T = 25$ rounds).

4.3. Spectral Efficiency Versus Number of UEs

Figure 5 shows the per-UE spectral efficiency as a function of K for $N = 3$ and $T = 20$ rounds. All schemes exhibit a decreasing trend as K grows from 6 to 21, due to the increased inter-cell interference and resource contention. At the training budget of this experiment, $T = 20$ rounds, the shortest of the paper, FedGAT sits slightly below FedAvg-MLP and close to C-GAT over the range of K ; as the convergence experiment shows, FedGAT is still improving at this point of training, and the full-budget ranking is the one reported in Table 2. The purpose of this figure is the trend with density rather than the ranking at a fixed short budget: the FedGAT policy trained at the baseline density $K = 15$ decays smoothly as K grows and does not collapse at high density, which indicates that encoding the interference relationships through the GATv2 attention mechanism yields policies that generalize across different UE densities. When reading this figure, the WMMSE and greedy curves should be treated as references rather than competitors, since both rely on information that FedGAT deliberately avoids collecting: full CSI at the BBU pool for WMMSE and network-wide geometry for the greedy rule. The gap to WMMSE widens at higher K , reflecting that additional training would be required to approach the upper bound in denser scenarios.




Figure 5. Per-UE spectral efficiency versus K ($N = 3$, $T = 20$ rounds).

4.4. Rate–Fronthaul Overhead Trade-Off

Figure 6 presents the steady-state sum-rate as a function of the per-round fronthaul overhead on a logarithmic scale for $N = 3$ and $K = 15$. The centralized schemes WMMSE and C-GAT require the exchange of full CSI, which amounts to $NKSM \times 2 \times 32$ bits per round, about 46.1 kbits under the considered parameters. They achieve approximately 11.75 and 4.76 b/s/Hz, respectively. FedGAT and FedAvg-MLP transmit only model parameter increments. With $|\theta| = 14,053$ parameters, this corresponds to about 1349.1 kbits per round at 32-bit precision, which is the aggregate uplink traffic from the N RRHs. Under this overhead, FedGAT achieves 7.10 b/s/Hz, whereas FedAvg-MLP attains 5.99 b/s/Hz despite transmitting more bits due to its larger model size. The greedy scheme operates without any fronthaul signaling. Since a zero-overhead point has no true position on a logarithmic axis, the greedy marker is drawn at the left edge of the axis as a reference level and its placement carries no quantitative meaning. FedGAT dominates FedAvg-MLP in this trade-off, matching or beating its rate at lower communication cost. All points in this figure are the steady-state values of the baseline configuration after the full training budget, so the figure introduces no new operating point: it rearranges the results of Figure 3 and Table 2 along a communication cost axis. FedGAT stands above FedAvg-MLP and C-GAT and below WMMSE, exactly as in those results, and no earlier trend is reversed. Relative to CSI sharing, FedGAT's training-phase overhead is higher at the evaluated network sizes and only falls below the CSI payload beyond $K \approx 439$ UEs. Its main advantage is therefore privacy and CSI-free inference rather than raw training-phase efficiency. After training, inference is performed locally at each RRH without further fronthaul usage.




Figure 6. Steady-state sum-rate versus per-round fronthaul overhead ($N = 3$, $K = 15$).

4.5. Summary of Steady-State Performance

Table 2 summarizes the steady-state weighted sum-rate for all schemes at the baseline configuration ($N = 3$, $K = 15$, $\tau = 5$), including mean value, standard deviation, and percentage of the WMMSE upper bound. FedGAT achieves about 60.5% of the upper bound, compared with 51.0% for FedAvg-MLP and 40.5% for C-GAT, confirming the benefit of graph-structured local models within the federated learning framework.

The greedy result merits a brief comment. It reaches 11.48 b/s/Hz at the baseline load ($N = 3$, $K = 15$, $S = 4$), where the interference is moderate. With orthogonal subbands and only three RRHs, full-power transmission under nearest-RRH association is already near the centralized water-filling solution. Coordinated association thus adds little extra gain. This holds only at the evaluated operating point. The greedy rule still needs global geometric channel knowledge, uses a fixed full-power policy, and does not learn or adapt as the network grows denser. FedGAT is therefore assessed among learning-based

methods that train on local data without exchanging raw CSI. In that class it improves on FedAvg-MLP by 18.7% and on C-GAT by 49.2%. WMMSE and the greedy rule remain non-learning references that require global channel or geometric knowledge. The strength of the greedy rule at this operating point comes from the orthogonal sub-band allocation and the small number of RRHs, which keep interference moderate. The rule depends on network-wide geometric knowledge for the association step, applies full power regardless of the interference state, and offers no mechanism to follow time-varying QoS weights or CSI restrictions. The practical case for FedGAT is therefore not to replace simple heuristics at light load, but to serve deployments where CSI cannot leave the RRHs and the policy must adapt to heterogeneous interference conditions.

Table 2. Steady-state performance ($N = 3$, $K = 15$, $\tau = 5$).

Scheme	Mean (b/s/Hz)	Std	% UB
WMMSE	11.75	± 1.52	100.0
C-GAT (centr. GNN)	4.76	± 2.20	40.5
FedGAT (proposed)	7.10	± 4.11	60.5
FedAvg-MLP	5.99	± 0.04	51.0
Greedy	11.48	± 1.47	97.7

5. Discussion

The simulation results confirm that graph-structured local models provide a clear advantage over flat MLP baselines in federated radio resource allocation. The GATv2 attention mechanism allows each RRH to weight interfering neighbors differently depending on the instantaneous channel state, which is a capability unavailable to architectures that treat all nodes uniformly. The convergence bound in Theorem 1 provides a theoretical explanation: the optimality gap grows with $\tau^2\Gamma$, so heterogeneous interference topologies require small local step counts to limit client drift.

The comparison with C-GAT reveals an important practical insight. Training a single centralized GATv2 model on data from all RRHs simultaneously leads to oscillatory behavior, because the global loss aggregates interference graphs with very different statistical properties. By contrast, FedGAT keeps local training within each RRH's own interference graph and uses FedAvg only to consolidate the learned parameters. This separation of local adaptation from global aggregation is the key design principle behind FedGAT's performance gain. This gap should be read at equal training budget and identical hyperparameters: with longer training or per-deployment tuning the centralized model may improve, but such tuning would also raise the CSI and computation load at the BBU pool.

The fronthaul analysis shows that FedGAT's training-phase overhead exceeds raw CSI transfer for networks with fewer than 439 UEs at 32-bit precision. This trade-off becomes favorable at larger scales or when 8-bit quantization is applied, reducing the crossover to 110 UEs. The primary benefit of FedGAT at current network sizes is therefore privacy preservation and fully CSI-free inference after training, rather than training-phase communication efficiency.

The comparison with the greedy rule delimits the practical scope of FedGAT further. At the evaluated operating points the greedy scheme stays close to the WMMSE upper bound, so the value of FedGAT does not lie in raw spectral efficiency against simple heuristics at light load. It lies in privacy-preserving operation, CSI-free inference, and adaptation to heterogeneous interference topologies, and the reported gains are accordingly restricted to the class of learning-based schemes that keep CSI local.

6. Conclusions

In this paper, we proposed FedGAT, a federated graph attention network for joint user association and inter-RRH interference mitigation in C-RAN. In FedGAT, each RRH constructs a local interference graph from its own channel measurements, trains a GATv2 model through τ local gradient steps, and forwards only the resulting parameter increment to the BBU pool. Global model consolidation proceeds via weighted FedAvg, and raw CSI is not disclosed at any stage.

A non-asymptotic convergence bound shows that the optimality gap scales as $\mathcal{O}(\tau^2\Gamma)$ with respect to the inter-RRH graph heterogeneity, which gives a practical guideline for selecting the local step count τ . Simulation results under 3GPP UMa conditions show that encoding the interference topology through GATv2 attention prevents the policy degeneracy that affects flat MLP local models. FedGAT surpasses FedAvg-MLP by 18.7% at lower fronthaul cost and exceeds C-GAT by 49.2% at the same training budget. These margins are averaged over ten Monte Carlo channel realizations and are reported with their standard deviations, so they should be read as indicative trends rather than statistically confirmed gains. The training-phase fronthaul of FedGAT exceeds the CSI payload of centralized schemes for $K < 439$ UEs, and 8-bit gradient quantization reduces this crossover to $K < 110$ UEs, while post-convergence inference is fully CSI-free.

Future work will further enlarge the Monte Carlo campaign with confidence intervals and statistical significance tests. Future work will also reduce the remaining gap to WMMSE through extended training with learning-rate annealing and by adding an explicit precoding vector optimization head. Model compression via structured pruning will further reduce the fronthaul breakeven. The framework will also be extended to RIS-assisted C-RAN architectures for B5G and 6G deployments.

Funding: This research received no external funding.

Data Availability Statement: The raw data supporting the conclusions of this article will be made available by the authors on request.

Conflicts of Interest: The author declares no conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

B5G	Beyond-5G
BBU	Baseband Unit
C-RAN	Cloud Radio Access Network
CoMP	Coordinated Multipoint
CSI	Channel State Information
FL	Federated Learning
GAT	Graph Attention Network
GNN	Graph Neural Network
IEEE	Institute of Electrical and Electronics Engineers
MIMO	Multiple-Input Multiple-Output
MLP	Multilayer Perceptron
MRT	Maximum-Ratio Transmission
RAN	Radio Access Network
REGNN	Random Edge Graph Neural Network
RRH	Remote Radio Head
SINR	Signal-to-Interference-Plus-Noise Ratio

UE	User Equipment
WMMSE	Weighted Minimum Mean-Square Error
5G	Fifth Generation
6G	Sixth Generation

Appendix A. Proof of Theorem 1

Lemma A1. By the L -smoothness of $\mathcal{L}$ and the FedAvg update:

$$\mathcal{L}(\theta^{t+1}) \leq \mathcal{L}(\theta^t) + \langle \nabla \mathcal{L}(\theta^t), \theta^{t+1} - \theta^t \rangle + \frac{L}{2} \|\theta^{t+1} - \theta^t\|^2. \quad (\text{A1})$$

Lemma A2. Fix a round t and an RRH n , and denote by $\mathbf{g}_n^{t,i} = \nabla \mathcal{L}_n(\theta_n^{t,i}; \zeta_n^{t,i})$ the stochastic gradient at local step i . Unrolling the τ local gradient steps gives

$$\theta_n^{t,\tau} - \theta^t = -\eta \sum_{i=0}^{\tau-1} \mathbf{g}_n^{t,i}. \quad (\text{A2})$$

Taking the squared norm, applying the Cauchy–Schwarz inequality to the sum of τ terms, and taking expectations:

$$\mathbb{E}[\|\theta_n^{t,\tau} - \theta^t\|^2] \leq \eta^2 \tau \sum_{i=0}^{\tau-1} \mathbb{E}[\|\mathbf{g}_n^{t,i}\|^2]. \quad (\text{A3})$$

Each term splits into a variance part and a mean part by Assumption 2:

$$\mathbb{E}[\|\mathbf{g}_n^{t,i}\|^2] \leq \sigma_g^2 + \mathbb{E}[\|\nabla \mathcal{L}_n(\theta_n^{t,i})\|^2]. \quad (\text{A4})$$

Following the bounded dissimilarity argument of Theorem 4 in [10], Assumptions 3 and 4 bound the mean part by $2\Gamma G^2$: when $\Gamma = 0$ every local objective shares its stationary points with the global objective and the local gradients carry no excess drift, while for $\Gamma > 0$ the deviation of each local gradient from the global direction grows with the heterogeneity measure, and the product ΓG^2 captures the worst case under the gradient bound G^2 . Combining the three displays yields the client-drift bound

$$\mathbb{E}[\|\theta_n^{t,\tau} - \theta^t\|^2] \leq \eta^2 \tau^2 (2\Gamma G^2 + \sigma_g^2). \quad (\text{A5})$$

To conclude, substitute the server update $\theta^{t+1} - \theta^t = \sum_n p_n (\theta_n^{t,\tau} - \theta^t)$ into Lemma A1. The inner product term contributes $-\eta \tau \mathbb{E}[\|\nabla \mathcal{L}(\theta^t)\|^2]$ plus a drift correction controlled by Lemma A2, and the quadratic term is bounded directly by Lemma A2 together with the convexity of the squared norm over the aggregation weights p_n . Averaging the independent stochastic gradients of the N RRHs reduces the variance contribution by the factor $1/N$, which produces the middle term of the bound. The step size condition $\eta \leq 1/(4L\tau)$ makes the negative descent term dominate the drift correction. Telescoping the resulting inequality over $t = 0, \dots, T-1$, dividing by $\eta \tau T$, and rearranging yields (11).

References

1. Checko, A.; Christiansen, H.L.; Yan, Y.; Scolari, L.; Kardaras, G.; Berger, M.S.; Dittmann, L. Cloud ran for mobile networks: A technology overview. *IEEE Commun. Surv. Tutor.* **2015**, *17*, 405–426. [CrossRef]
2. Shi, Q.; Razaviyayn, M.; Luo, Z.-Q.; He, C. An iteratively weighted mmse approach to distributed sum-utility maximization for a mimo interfering broadcast channel. *IEEE Trans. Signal Process.* **2011**, *59*, 4331–4340. [CrossRef]
3. Wubben, D.; Rost, P.; Bartelt, J.S.; Lalam, M.; Savin, V.; Gorgoglione, M.; Dekorsy, A.; Fettweis, G. Benefits and impact of cloud computing on 5g signal processing: Flexible centralization through cloud-ran. *IEEE Signal Process. Mag.* **2014**, *31*, 35–44. [CrossRef]
4. Eisen, M.; Ribeiro, A. Optimal wireless resource allocation with random edge graph neural networks. *IEEE Trans. Signal Process.* **2020**, *68*, 2977–2991. [CrossRef]

5. Shen, Y.; Shi, Y.; Zhang, J.; Letaief, K.B. Graph neural networks for scalable radio resource management: Architecture design and theoretical analysis. *IEEE J. Sel. Areas Commun.* **2021**, *39*, 101–115. [\[CrossRef\]](#)
6. McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; Agüera y Arcas, B. Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*; PMLR: Fort Lauderdale, FL, USA, 2017; Volume 54, pp. 1273–1282.
7. Elbir, A.M.; Coleri, S. Federated learning for channel estimation in conventional and RIS-assisted massive MIMO. *IEEE Trans. Wirel. Commun.* **2022**, *21*, 4255–4268. [\[CrossRef\]](#)
8. Tehrani, P.; Restuccia, F.; Levorato, M. Federated deep reinforcement learning for the distributed control of NextG wireless networks. In *Proceedings of the IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*; IEEE: Piscataway, NJ, USA, 2021; pp. 248–253.
9. Brody, S.; Alon, U.; Yahav, E. How attentive are graph attention networks? In *Proceedings of the International Conference on Learning Representations (ICLR)*, Virtual Event, 25–29 April 2022.
10. Li, T.; Sahu, A.K.; Zaheer, M.; Sanjabi, M.; Talwalkar, A.; Smith, V. Federated optimization in heterogeneous networks. In *Proceedings of the Machine Learning and Systems (MLSys)*, Austin, TX, USA, 2–4 March 2020; Volume 2, pp. 429–450.
11. Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Liò, P.; Bengio, Y. Graph attention networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*, Vancouver, BC, Canada, 30 April–3 May 2018.
12. Sun, H.; Chen, X.; Shi, Q.; Hong, M.; Fu, X.; Sidiropoulos, N.D. Learning to optimize: Training deep neural networks for interference management. *IEEE Trans. Signal Process.* **2018**, *66*, 5438–5453. [\[CrossRef\]](#)
13. Chowdhury, A.; Verma, G.; Rao, C.; Swami, A.; Segarra, S. Unfolding WMMSE using graph neural networks for efficient power allocation. *IEEE Trans. Wirel. Commun.* **2021**, *20*, 6004–6017. [\[CrossRef\]](#)
14. Peng, M.; Sun, Y.; Li, X.; Mao, Z.; Wang, C. Recent advances in cloud radio access networks: System architectures, key techniques, and open issues. *IEEE Commun. Surv. Tutor.* **2016**, *18*, 2282–2308. [\[CrossRef\]](#)
15. Ye, Q.; Rong, B.; Chen, Y.; Al-Shalash, M.; Caramanis, C.; Andrews, J.G. User association for load balancing in heterogeneous cellular networks. *IEEE Trans. Wirel. Commun.* **2013**, *12*, 2706–2716. [\[CrossRef\]](#)
16. Niknam, S.; Dhillon, H.S.; Reed, J.H. Federated learning for wireless communications: Motivation, opportunities and challenges. *IEEE Commun. Mag.* **2020**, *58*, 46–51. [\[CrossRef\]](#)
17. Lim, W.Y.B.; Luong, N.C.; Hoang, D.T.; Jiao, Y.; Liang, Y.C.; Yang, Q.; Niyato, D.; Miao, C. Federated learning in mobile edge networks: A comprehensive survey. *IEEE Commun. Surv. Tutor.* **2020**, *22*, 2031–2063. [\[CrossRef\]](#)
18. Luong, N.C.; Hoang, D.T.; Gong, S.; Niyato, D.; Wang, P.; Liang, Y.-C.; Kim, D.I. Applications of deep reinforcement learning in communications and networking: A survey. *IEEE Commun. Surv. Tutor.* **2019**, *21*, 3133–3174. [\[CrossRef\]](#)
19. Nasir, Y.S.; Guo, D. Multi-agent deep reinforcement learning for dynamic power allocation in wireless networks. *IEEE J. Sel. Areas Commun.* **2019**, *37*, 2239–2250. [\[CrossRef\]](#)
20. Liang, L.; Ye, H.; Li, G.Y. Spectrum sharing in vehicular networks based on multi-agent reinforcement learning. *IEEE J. Sel. Areas Commun.* **2019**, *37*, 2282–2292. [\[CrossRef\]](#)
21. Wang, J.; Jiang, W.; Xu, H.; Hu, J.; Wu, L.; Shu, F.; Fang, Z. A novel resource allocation method based on hierarchical deep reinforcement learning for cognitive internet of vehicles with unknown channel state information. *Comput. Netw.* **2025**, *264*, 111270. [\[CrossRef\]](#)
22. Mozaffari, M.; Saad, W.; Bennis, M.; Nam, Y.-H.; Debbah, M. A tutorial on uavs for wireless networks: Applications, challenges, and open problems. *IEEE Commun. Surv. Tutor.* **2019**, *21*, 2334–2360. [\[CrossRef\]](#)
23. Ding, R.; Zhou, F.; Wu, Q.; Wong, K.-K.; Al-Dhahir, N. Optimization-driven drl for resource allocation under licensed and unlicensed uav spectrum sharing networks against uncertain jamming. *IEEE Trans. Mob. Comput.* **2026**, *early access*. [\[CrossRef\]](#)
24. Eldeeb, E.; Alves, H. An offline multi-agent reinforcement learning framework for radio resource management. *IEEE Trans. Mob. Comput.* **2026**, *25*, 1137–1150. [\[CrossRef\]](#)
25. Wang, J.; Dan, W.; Yang, K.; Tang, X.; Yan, L. IRL-D3QN: An intelligent multi-agent learning framework for dynamic spectrum management in vehicular networks. *Future Gener. Comput. Syst.* **2026**, *180*, 108371. [\[CrossRef\]](#)
26. Zhang, Y.; Guo, D. Multi-agent reinforcement learning for multi-cell spectrum and power allocation. *IEEE Trans. Commun.* **2025**, *73*, 5980–5992. [\[CrossRef\]](#)
27. Wang, Y.; Lei, J.; Shang, F.; Li, Y. A comprehensive survey of multi-agent deep reinforcement learning for wireless spectrum management. *Neurocomputing* **2025**, *653*, 131236. [\[CrossRef\]](#)
28. Wu, C.; Wu, F.; Lyu, L.; Qi, T.; Huang, Y.; Xie, X. A federated graph neural network framework for privacy-preserving personalization. *Nat. Commun.* **2022**, *13*, 3091. [\[CrossRef\]](#) [\[PubMed\]](#)

29. Liu, R.; Xing, P.; Deng, Z.; Li, A.; Guan, C.; Yu, H. Federated graph neural networks: Overview, techniques, and challenges. *IEEE Trans. Neural Netw. Learn. Syst.* **2025**, *36*, 4279–4295. [[CrossRef](#)] [[PubMed](#)]
30. Feys, T.; Van der Perre, L.; Rottenberg, F. Learning to quantize and precode in massive MIMO systems for energy reduction: A graph neural network approach. *IEEE J. Sel. Top. Signal Process.* **2025**, *19*, 983–998. [[CrossRef](#)]

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