American Journal of Applied Science and Technology

VOLUME Vol.06 Issue 02 2026
PAGE NO. 107-111
DOI 10.37547/ajast/Volume06Issue02-11

Research and Development of Ai-Based Resource Allocation Methods for Aerial and Terrestrial Base Stations in Communication Networks

Tumaeva Aygul Medetbaevna

Doctoral Student of Belarusian-Uzbek Joint Interdisciplinary Institute of Applied Technical Qualifications, Uzbekistan

B.T. Kaipbergenov

Scientific Advisor, Professor of Nukus State Technical University, Uzbekistan

Received: 31 December 2025; **Accepted:** 23 January 2026; **Published:** 28 February 2026

Abstract: This study examines artificial intelligence–based approaches for efficient resource allocation in hybrid aerial and terrestrial base station networks. As wireless traffic demand and service diversity continue to increase, conventional optimization techniques become less effective in dynamic and large-scale environments. Therefore, AI methods such as deep reinforcement learning, supervised learning, and graph neural networks are employed to optimize spectrum sharing, power control, user association, and UAV trajectory planning. Furthermore, the research considers cross-layer constraints including backhaul capacity, energy consumption, and quality-of-service requirements. The results demonstrate that AI-driven strategies enhance system throughput, fairness, adaptability, and energy efficiency while maintaining reliable performance in time-varying network conditions.

Keywords: Artificial Intelligence; Resource Allocation; Aerial Base Stations; Terrestrial Base Stations; UAV Communications; Hybrid Wireless Networks; Deep Reinforcement Learning; Graph Neural Networks; Spectrum Management; Energy Efficiency; 6G Networks.

INTRODUCTION:

The rapid densification of wireless traffic, together with the emergence of new service classes such as ultra-reliable low-latency communications and massive machine-type communications, has made resource allocation one of the central bottlenecks of modern networks. At the same time, the networking paradigm is expanding beyond classical terrestrial base stations toward hybrid architectures that include aerial base stations, typically implemented via unmanned aerial vehicles (UAVs), high-altitude platforms, or other airborne nodes. Consequently, the resource allocation problem is no longer limited to power control and scheduling in relatively stable cells; instead, it becomes a coupled decision process where radio resources, backhaul capacity, mobility, energy, and three-dimensional placement interact. Therefore, artificial intelligence, and in particular learning-based optimization, is increasingly viewed as a pragmatic way to manage complexity, respond to uncertainty, and adapt to dynamic environments where accurate models are either unavailable or too expensive to use online.

In hybrid aerial–terrestrial networks, terrestrial base stations deliver stable coverage and high capacity in areas with fixed infrastructure, whereas aerial base stations are introduced to provide on-demand coverage, rapid deployment after disasters, hotspot offloading, or temporary capacity boosts during events. However, while the aerial layer adds flexibility, it also introduces nontrivial constraints. For instance, UAV base stations are limited by battery energy, payload, regulatory flight restrictions, and time-varying air-to-ground channel characteristics. Moreover, their position directly influences link quality, interference topology, and user association. Thus, resource allocation in such networks naturally becomes a multi-domain optimization problem, where spectrum, time-frequency scheduling, transmit power, beamforming, user association, caching decisions, and UAV trajectories must be co-designed [5, 16-25].

At the same time, the service heterogeneity of next-generation systems makes the objective multi-criteria. In practice, an operator may want to maximize sum-rate, yet simultaneously guarantee minimum rate for cell-edge users, maintain latency bounds for interactive services, and reduce energy cost. Hence, the “best” policy is rarely a single metric optimum; instead, it is a compromise among competing objectives under uncertainty. For this reason, AI methods are particularly attractive because they can approximate complex decision boundaries, learn from data, and adapt policies to new regimes without requiring perfect analytical channel models.

A typical hybrid resource allocation problem can be expressed as an optimization over decision variables for both layers:

- Terrestrial layer decisions: subcarrier/resource block assignment, scheduling, power allocation, beamforming, and load balancing among ground stations.
- Aerial layer decisions: UAV 3D placement, trajectory planning, hovering time, power allocation, spectrum reuse strategy, and user association.
- Cross-layer decisions: backhaul selection (e.g., mmWave, microwave, optical), partitioning between access and backhaul resources, caching or edge computing placement, and handover control [2, 98-105].

If we denote the overall system state at time as, including user positions, traffic loads, channel conditions, UAV battery levels, and interference measurements, then we can represent the control decisions as, including allocations and movement updates. The system evolves according to unknown or partially known dynamics, and the network obtains a utility reflecting throughput, latency, energy, and fairness. Consequently, the problem becomes a sequential decision-making task, which aligns well with reinforcement learning (RL) formulations.

Nevertheless, even a “single-shot” snapshot optimization is difficult because the problem is generally nonconvex due to interference coupling, integer scheduling variables, and nonlinear air-to-ground path loss. Therefore, classical methods such as mixed-integer programming, successive convex approximation, or Lagrangian relaxation often require careful modeling and still may be too slow for real-time operation. In contrast, AI can either replace the online solver (by learning a fast mapping from state to near-optimal action) or augment the solver (by providing warm starts, learned heuristics, or policy selection).

AI-based resource allocation offers three major advantages. First, it provides adaptation: a learned policy can react to changing user density, mobility patterns, and channel statistics. Second, it supports scalability: once trained, inference can be fast even in large networks. Third, it enables model-free control: RL can operate when the environment is complicated or partially observed.

However, these benefits are only realized if the design accounts for networking realities. Specifically, resource allocation is safety-critical in the sense that violations can cause coverage holes, severe interference, or failure to meet service-level agreements. Thus, AI must be constrained. In practice, this means incorporating hard constraints (e.g., maximum transmit power, minimum SINR, UAV no-fly zones) and ensuring stable performance under distribution shifts. Therefore, safe RL, constrained optimization, robust learning, and hybrid model-based/model-free schemes become essential rather than optional [1].

One practical approach is to generate training data using high-quality offline optimization (e.g., solving many instances with convex approximation or metaheuristics) and then train a neural network to imitate the solver. In this case, the model learns a mapping, where the outputs can be power levels, association probabilities, or UAV placement suggestions. The advantage is predictability and speed, and moreover the network can be deployed with low inference latency at the base station controller.

Yet, supervised learning inherits the biases and coverage limits of the generated dataset. Hence, it is effective when the operational regime is similar to the training regime, but it may degrade if user distribution, propagation conditions, or interference patterns shift. Therefore, it is often used as a component: for example, as a “policy prior” that is refined online by RL or constrained optimization.

Because UAV positions and network states evolve over time, RL is especially suitable for jointly optimizing communication and mobility. In an RL framework, the agent observes measurements (partial or full), chooses actions (movement, scheduling, power), and receives rewards related to capacity and quality-of-service. Deep RL variants can handle high-dimensional states, and multi-agent RL can distribute control across multiple UAVs and terrestrial stations.

Nevertheless, naive RL can be unstable due to sparse rewards, nonstationarity (especially in multi-agent settings), and constraint violations. Therefore, research increasingly focuses on constrained RL, where the agent optimizes performance while maintaining constraints such as outage probability, interference thresholds, or battery safety margins. In addition, the reward must be carefully shaped to avoid pathological behaviors, such as a UAV chasing throughput at the cost of abandoning underserved users [6].

Wireless networks are naturally represented as graphs, where nodes correspond to users and stations, and edges encode channel or interference relationships. Graph neural networks (GNNs) can exploit this structure to produce allocation decisions that generalize across different network sizes and layouts. For example, a GNN can learn to assign resources based on local interference neighborhoods rather than relying on fixed-size inputs.

In hybrid aerial–terrestrial systems, topology changes as UAVs move, which makes graph-based models especially meaningful. Consequently, GNNs can be integrated with RL, enabling policies that are both sequential and structure-aware. This combination is promising for scaling to large deployments where purely centralized control is too heavy.

Since user data and network measurements are sensitive, and since edge devices and base stations produce large volumes of telemetry, federated learning offers a way to train models without transferring raw data to a central server. In this setting, terrestrial base stations and aerial nodes train local models and share only gradient updates or model parameters. As a result, privacy is improved and bandwidth overhead may be reduced.

However, wireless federated learning itself faces challenges, including communication overhead, straggler effects, and non-IID data across cells. Therefore, resource allocation and learning become intertwined: the network must allocate resources to support the training process, while the training aims to improve resource allocation. This creates a feedback loop that must be engineered carefully.

User association determines which users connect to terrestrial or aerial stations. In hybrid networks, association is not just a function of signal strength; it must reflect load, backhaul constraints, and QoS demands. AI can learn association policies that predict congestion and anticipate mobility, thereby preventing oscillations and reducing handover overhead. Moreover, combining association with UAV placement yields significant gains, since moving the aerial station can be more effective than reassigning users under poor geometry.

Aerial stations can reuse spectrum with terrestrial cells, yet spectrum reuse can cause strong interference, particularly when line-of-sight links dominate. Thus, the problem is to choose reuse patterns and power levels that deliver capacity while controlling interference. AI-based controllers can learn to identify interference-sensitive regions and adjust reuse accordingly. Additionally, beamforming and directional links can be exploited, and learning methods can incorporate spatial features to decide when aggressive reuse is safe.

Trajectory design is a uniquely aerial decision variable. A UAV can reposition to improve channels, but movement costs energy and time. Consequently, trajectory planning must trade off communication gains against energy expenditure and service continuity. RL can handle this tradeoff by incorporating battery state into the system state and penalizing energy use while rewarding QoS satisfaction. Furthermore, multi-UAV coordination is crucial: without coordination, UAVs may cluster and interfere, whereas coordinated policies can cover disjoint hotspots and share backhaul resources.

Aerial base stations often rely on wireless backhaul to connect to the core network, and the backhaul capacity may be time-varying. Therefore, access scheduling should be backhaul-aware: serving more users than backhaul allows can create queues and latency violations. AI can predict backhaul fluctuations and proactively adjust access rates, caching strategies, or even UAV position to maintain stable end-to-end performance.

Energy is a dominant constraint for UAV-based stations, and it is also a cost driver for terrestrial infrastructure. Therefore, energy-aware resource allocation aims to minimize power consumption while maintaining QoS. AI can learn when to deploy UAVs, when to land or recharge, and how to balance power across links. Moreover, the system can adopt sleep modes for underutilized terrestrial cells and offload to aerial platforms temporarily, although such strategies must be optimized carefully to avoid coverage gaps.

A credible study in AI-based resource allocation requires a robust evaluation methodology. First, the simulation environment should capture realistic air-to-ground and ground-to-ground propagation, including probabilistic line-of-sight conditions and urban morphology effects. Second, mobility models for users and UAVs must reflect plausible patterns, such as commuting flows, event crowding, or disaster evacuation. Third, the evaluation should compare AI methods against strong baselines, including classical optimization heuristics, greedy association, proportional fair scheduling, and power control schemes.

In addition, the evaluation must report not only average throughput but also tail performance: outage probability, 5th percentile rates, latency distributions, and fairness indices. This is important because AI methods can sometimes improve mean metrics while harming vulnerable users. Furthermore, the study should include stress tests under distribution shifts, such as unexpected traffic surges or partial link failures, because real deployments rarely follow training assumptions.

Despite progress, several challenges remain. First, generalization is a central risk: a model trained in one city layout or traffic pattern may fail in another. Therefore, domain adaptation, meta-learning, and robust training methods are necessary. Second, constraint handling is critical: future systems require guaranteed QoS, and thus constrained RL and safe policy deployment are key. Third, multi-agent coordination remains difficult due to nonstationarity; nevertheless, communication-aware multi-agent learning and centralized training with decentralized execution offer viable pathways.

Moreover, explainability is increasingly relevant. Operators need to understand why a policy chooses a certain spectrum reuse pattern or UAV movement, especially when failures occur. Consequently, interpretable models, post-hoc explanations, and policy auditing become important for operational trust. Finally, the integration of AI into network controllers must consider standardization, interoperability, and cybersecurity, because adversarial manipulation of learning inputs could cause misallocation of resources or denial-of-service effects.

CONCLUSION

AI-based methods for resource allocation in hybrid aerial–terrestrial base station networks represent a promising direction for next-generation communication systems, especially because they can adapt to dynamics, scale to large scenarios, and operate under partial models. At the same time, the hybrid setting intensifies the complexity by introducing mobility, energy constraints, and rapidly changing interference topologies. Therefore, the most effective solutions are not purely data-driven nor purely analytical; rather, they are hybrid approaches that integrate learning with domain constraints, structured representations such as graphs, and robust evaluation against strong baselines. Consequently, future research should emphasize safe and generalizable learning, multi-agent coordination, backhaul-aware optimization, and deployment-ready architectures that can deliver performance gains while respecting the strict reliability requirements of modern networks.

REFERENCES

1. Eldar, Y. C., Goldsmith, A., Gündüz, D., & Poor, H. V. (Eds.). (2022). Machine learning and wireless communications. Cambridge University Press.
2. Jiang, C., Zhang, H., Ren, Y., Han, Z., Chen, K. C., & Hanzo, L. (2016). Machine learning paradigms for next-generation wireless networks. IEEE Wireless Communications, 24(2), 98-105.
3. Mozaffari, M., Saad, W., Bennis, M., & Debbah, M. (2016). Unmanned aerial vehicle with underlaid device-to-device communications: Performance and tradeoffs. IEEE Transactions on Wireless Communications, 15(6), 3949-3963.
4. Park, J., Samarakoon, S., Bennis, M., & Debbah, M. (2019). Wireless Network Intelligence at the Edge. Proc. Ieee, 107(11), 2204-2239.
5. Shehzad, M. K., Rose, L., Butt, M. M., Kovacs, I. Z., Assaad, M., & Guizani, M. (2022). Artificial intelligence for 6G networks: Technology advancement and standardization. IEEE Vehicular Technology Magazine, 17(3), 16-25.
6. Zhang, S., Zhang, H., & Song, L. (2020). Beyond D2D: Full dimension UAV-to-everything communications in 6G. IEEE Transactions on Vehicular Technology, 69(6), 6592-6602.