

GAN-MADRL-Driven UAV-Assisted Vehicular Edge Computing for Resilient and Smart Transportation Systems in 6G Networks

Mesfin Leranso Betalo¹, Xiaoshan Bai², Zongze Wu³, Jianqiang Li⁴

¹College of Mechatronics and Control Engineering, Shenzhen University, College of Computer Science and Software Engineering, Shenzhen University, Shenzhen 518060, China, Department of Information Technology, College of Engineering and Technology, Wachemo University, P.O. Box 667, Hosaena, Ethiopia.

²College of Mechatronics and Control Engineering, and the National Engineering Laboratory for Big Data System Computing Technology, Shenzhen University, Shenzhen 518060, China.

³College of Mechatronics and Control Engineering and the Guangdong Laboratory of Artificial Intelligence and Digital Economy (SZ), Shenzhen University, Shenzhen 518060, China.

⁴National Engineering Laboratory for Big Data System Computing Technology, and the College of Computer Science and Software Engineering, Shenzhen University, Shenzhen 518060, China Email: mesfinleranso@szu.edu.cn, baixiaoshan@szu.edu.cn, zzwu@szu.edu.cn, lijq@szu.edu.cn

Corresponding Author: mesfinleranso@szu.edu.cn

Received 01 March 2026; revised 15 April 2026; accepted 15 June 2026

Abstract

Unmanned aerial vehicles (UAVs) integrated with vehicular edge computing (VEC) in 6G-enabled Intelligent Transportation Systems (ITS) offer a resilient platform to enhance urban mobility, road safety, and public transport efficiency. Acting as agile edge nodes, UAVs enable low-latency communication, distributed intelligence, and adaptive coverage in dynamic and infrastructure-limited environments. Nevertheless, challenges such as UAV energy constraints, intermittent connectivity, and heterogeneous data distributions persist. This paper presents a generative adversarial network (GAN)-enhanced Multi-Agent Deep Reinforcement Learning (MADRL) framework for UAV-assisted VEC in 6G-enabled ITS. The proposed framework aims to minimize UAV energy consumption and network latency while ensuring reliable connectivity and accurate traffic prediction. Specifically, it jointly addresses traffic prediction, UAV path planning, and computational offloading (TPCO), formulated as a Markov Decision Process (MDP). UAVs, vehicles, and edge servers operate as decentralized agents that learn collaboratively to optimize latency, energy consumption, and connectivity. GAN modules enrich sparse and heterogeneous datasets, thereby improving training robustness and convergence. The simulation results indicate that the proposed GAN-MADRL framework achieves about 30% better traffic-prediction accuracy, 35% higher energy efficiency, and 40% lower latency compared to CDL, FL, PPO, and Greedy baselines.

Keywords: Generative Adversarial Networks, Vehicular Edge Computing, Intelligent Transportation Systems, UAV trajectory optimization, Traffic prediction

1. Introduction

The evolution of sixth-generation (6G) wireless networks is ushering in a new era of resilient and smart transportation systems, driven by the convergence of artificial intelligence (AI), the Internet of Things (IoT), and sustainable technologies. These developments are reshaping urban mobility, road safety, and public transport efficiency by enabling ultra-reliable, low-latency, and high-capacity connectivity across densely distributed and heterogeneous mobile agents (Jha et al., 2024; Mohammadisarab et al., 2024). Core

6G technologies—such as terahertz (THz) communication, massive machine-type communication (mMTC), and ultra-reliable lowlatency communication (URLLC)—provide the foundation for time-sensitive Intelligent Transportation Systems (ITS) services, including autonomous driving, vehicle-to-everything (V2X) communication, and real-time traffic management (Creß et al., 2023; Han et al., 2023).

In parallel, Vehicular Edge Computing (VEC) has emerged as a distributed paradigm that decentralizes computation by offloading data processing from the cloud to network edges. This paradigm enhances responsiveness, reduces communication latency, and enables localized, intelligent decision-making within smart cities (Rezaee et al., 2024). To further strengthen flexibility and resilience, Unmanned Aerial Vehicles (UAVs) are increasingly integrated into VEC as aerial edge nodes, providing dynamic coverage, adaptive computation offloading, and connectivity support in infrastructure-limited or emergency-prone areas (Hu et al., 2021; Yan et al., 2024). The integration of UAVs with 6G-enabled VEC forms an agile, scalable, and sustainable architecture that improves urban mobility management, supports emergency responses, and enables real-time coordination among connected vehicles (Liu et al., 2024; Wang et al., 2023).

AI-driven analytics and IoT-enabled sensing empower UAV-assisted ITS to manage traffic, predict mobility, and maintain resilience against congestion or disruptions. Federated learning (FL) enables privacy-preserving, distributed model training among UAVs and vehicles without central aggregation (Ali et al., 2024). When combined with GANs and attention mechanisms (Ali et al., 2025), these systems enhance prediction accuracy, adaptability, and robustness under dynamic, non-IID conditions. Integrated with 6G’s multi-hop and edgeintelligent infrastructure, they promote greener and more adaptive ITS operations.

A. Challenges

Despite remarkable advancements in AI-driven edge intelligence, existing UAV-assisted VEC approaches in 6G-enabled ITS still face significant challenges in scalability, communication efficiency, and learning stability under high mobility and heterogeneous data. The presence of non-independent and identically distributed (non-IID) and sparse local datasets collected by UAVs and vehicles limits model generalization and slows convergence in distributed learning. Furthermore, UAVs operate under strict energy and mobility constraints, making it difficult to balance computation offloading, communication scheduling, and trajectory planning in highly dynamic and resource-limited environments. In addition, frequent model updates and high communication overhead in traditional centralized or synchronous federated schemes introduce latency and bandwidth congestion, reducing real-time responsiveness and scalability. Ensuring robustness, privacy, and adaptability under rapidly changing traffic and environmental conditions remains a critical challenge.

B. Motivation

Despite advances in AI-driven edge intelligence, existing approaches remain insufficient for UAV-assisted VEC in large-scale, dynamic 6G-enabled ITS. Centralized Deep Learning (CDL) causes excessive communication overhead, privacy risks, and latency due to continuous data exchange (Betalo et al., 2021), while Federated Learning (FL) suffers from synchronization delays and degraded accuracy under non-IID and mobile conditions (Shinde & Tarchi, 2023). Reinforcement Learning (RL) methods such as DQN and PPO offer adaptability but converge slowly and become unstable in multi-agent real-time scenarios (Betalo et al., 2022; Fakirah et al., 2023; Guan et al., 2023). To address these issues, this work proposes a GAN-enhanced Multi-Agent Deep Reinforcement Learning (GAN-MADRL) framework in which UAVs function as collaborative edge nodes that jointly optimize traffic prediction, trajectory planning, and computation offloading under energy and mobility constraints. GAN-based data augmentation improves robustness and generalization, while decentralized coordination enhances scalability, energy efficiency, and communication resilience, enabling resilient and sustainable ITS for nextgeneration urban mobility.

C. Key Contributions

The main contributions of this work are summarized as follows:

We propose a novel GAN-based MADRL framework for UAV-assisted VEC in 6G-enabled ITS. This framework enables decentralized, privacy-preserving, and communication-efficient learning across UAVs, vehicles, and edge servers, supporting real-time decision-making under dynamic and resourceconstrained conditions.

We formulate a joint optimization problem for Traffic prediction, Path planning, and Computational offloading (TPCO) by modeling the system as a Markov Decision Process (MDP) within a MADRL paradigm. This design allows decentralized agents (UAVs, vehicles, and edge servers) to collaboratively make energy-efficient and latency-aware decisions, enhancing scalability in highly dynamic ITS environments.

To address data heterogeneity and sparsity, we integrate GAN modules at the edge to generate realistic mobility and traffic patterns. This augmentation improves local model generalization, accelerates convergence, and enhances prediction accuracy in non-IID environments.

Extensive simulations demonstrate that the proposed GAN-MADRL framework significantly outperforms CDL, FL, PPO, and greedy baselines. It achieves faster convergence, reduced communication overhead, and higher scalability under dense agent deployments and dynamic traffic patterns, validating its effectiveness for real-time, large-scale 6G-enabled ITS.

Recent advancements in UAV-assisted VEC for 6G-enabled ITS have focused on optimizing task offloading, trajectory planning, and resource management to enhance communication efficiency and energy performance (Betalo et al., 2024). However, most existing DRL- and scheduling-based UAV-assisted ITS frameworks lack scalability, privacy preservation, and robustness under non-IID data distributions and high mobility. FL approaches (Zhang et al., 2023) have emerged to mitigate data sharing and privacy concerns, yet they largely rely on synchronous updates, which are unsuitable for dynamic and heterogeneous vehicular environments. Asynchronous and incremental FL variants, such as ASYFED (Li et al., 2022), diversity-driven (Lei et al., 2022), and K-asynchronous methods (Chen et al., 2024), have improved scalability and convergence but still struggle with joint task coordination and real-time adaptability. Meanwhile, generative approaches (Betalo, Ullah, et al., 2025) demonstrate promise for mitigating data sparsity at the edge but remain underexplored for multi-agent ITS optimization. Likewise, joint optimization studies on UAV trajectory and resource allocation (Betalo et al., 2023) often employ centralized or static schemes that cannot cope with the dynamic, distributed, and large-scale nature of next-generation ITS. To bridge these gaps, our proposed GAN-based IFL framework introduces a decentralized, asynchronous learning paradigm enhanced with GAN-driven data augmentation to ensure robust, scalable, and privacy-preserving coordination across UAVs, vehicles, and edge servers. By jointly optimizing traffic prediction, UAV trajectory planning, and computational offloading, it achieves superior adaptability, energy efficiency, and latency reduction compared to conventional CDL, FL, PPO, and greedy methods, making it well-suited for real-time 6G-enabled ITS.

II. System Model

We consider a 6G-enabled UAV-assisted Vehicular Edge Computing (VEC) system comprising vehicles ((V)), UAVs ((U)), and edge servers ((E)). Vehicles collect local traffic data, UAVs act as mobile edge nodes for monitoring and computation, and edge servers coordinate global optimization. Within the proposed GANMADRL framework, these agents collaboratively perform real-time traffic prediction, trajectory planning, and computational offloading, as illustrated in Figure 1.

A. Communication Model

Three communication modes are considered: Ground-to-Air (G2A), Air-to-Ground (A2G), and Air-to-Air (A2A). Their channel capacities are modeled as:

!

$$C_{x,y}(t) = B \cdot \log_2 \left(1 + \frac{P_x \cdot G_{x,y}}{N_0 \cdot d_{x,y}^\gamma} \right) \quad (1)$$

where P_x is the transmission power of node x , $G_{x,y}$ the channel gain, N_0 noise power, $d_{x,y}$ the distance between nodes, and γ the path loss exponent. A minimum capacity constraint $C_{x,y}(t) \geq C_{\min}$ ensures reliable federated updates and low-latency communication.




Fig. 1: Proposed GAN-MADRL-Driven UAV-Assisted VEC in 6G Networks.

B. Traffic and Trajectory Models

Traffic state prediction is defined as

$$\mathcal{T}(t+1) = \mathcal{f}_{\mathcal{T}}(\mathcal{T}(t), \mathbf{u}_{\text{traffic}}(t), \mathbf{D}_{\text{local}}(t)), \quad (2)$$

where $\mathcal{f}_{\cdot}$ is implemented via ML models enhanced by GAN-generated samples. Vehicle trajectory forecasting follows accounting for propulsion, communication, and control overhead.

$$\mathbf{x}_i(t+1) = \mathcal{G}(\mathbf{x}_i(t), \mathbf{u}_i(t), \mathbf{R}_i(t)), \text{ where} \quad (3)$$

$\mathcal{G}(\cdot)$ uses vehicle dynamics and road context. For UAVs, the state evolves as

$$\mathbf{x}_u(t+1) = \mathbf{x}_u(t) + \quad (4)$$

$\Delta t \cdot \mathbf{v}_u(t)$, with energy consumption modeled by

$$E_u(t) = h(\mathbf{x}_u(t), \mathbf{v}_u(t), \mathbf{u}_u(t)), \quad (5)$$

This system model enables cooperative GAN-MADRL optimization of traffic prediction, UAV path planning, and computation offloading in dynamic ITS environments.

III. Problem Formulation

We formulate a multi-objective optimization problem for UAV-assisted VEC in 6G-enabled ITS, aiming to jointly optimize TPCO. The primary objectives are to minimize UAV energy consumption and network

latency while ensuring accurate predictions and reliable connectivity. Let $\mathcal{T}(\mathcal{t}+1)$ and $\hat{\mathcal{T}}(\mathcal{t}+1)$ denote the actual and predicted traffic states, and $\mathbf{x}(\mathcal{t}+1)$ and $\hat{\mathbf{x}}(\mathcal{t}+1)$ the actual and predicted trajectories of vehicle i . The energy consumed by UAV u at time $\mathcal{t}$ is given by $E_u(\mathcal{t})$. The joint optimization problem is expressed as:

P1:

$$\min_{T, \hat{T}, x_i, x_u} \left[\alpha \sum_t \sum_{u \in \mathcal{U}} E_u(t) + \beta \Theta(t) \right] \quad (6a)$$

subject to:

$$\text{C1: } E_u(t) \leq E_{\max}, \forall u \in \mathcal{U}, t \quad (6b)$$

$$\text{C2: } x_u(t+1) = x_u(t) + \Delta t v_u(t), \forall u \in \mathcal{U}, t \quad (6c)$$

$$\text{C3: } D(x_u(t), x_i(t)) \leq D_{\max}, \forall u \in \mathcal{U}, i \in \mathcal{V}, t \quad (6d)$$

$$\text{C4: } x_i(t+1) \in \mathcal{R}_{\text{road}}, \forall i \in \mathcal{V}, t \quad (6e)$$

$$\text{C5: } \sum_t E_u(t) \leq E_{\text{total}}, \forall u \in \mathcal{U}. \quad (6f)$$

$$\text{C6: } \sum_t \|\mathbf{T}(t+1) - \hat{\mathbf{T}}(t+1)\|^2 \leq \epsilon_{\text{traffic}}, \quad (6g)$$

$$\text{C7: } \sum_i \sum_t \|\mathbf{x}_i(t+1) - \hat{\mathbf{x}}_i(t+1)\|^2 \leq \epsilon_{\text{traj}}, \quad (6h)$$

$$\text{C8: } \eta_u(t) \leq \eta_{\text{thr}}, \forall u, t. \quad (6i)$$

In the optimization problem (P1), the constraints ensure feasibility, safety, and efficiency in UAV-assisted VEC for 6G-enabled ITS. $\mathcal{C1}$ limits the instantaneous energy consumption of UAVs to remain within their maximum power budgets, ensuring sustainable operation. $\mathcal{C2}$ enforces the UAV kinematic model, where the next position depends on velocity and time-step updates. $\mathcal{C3}$ guarantees reliable connectivity by restricting the UAV–vehicle distance to within a maximum communication range. $\mathcal{C4}$ ensures that predicted vehicle trajectories remain inside feasible road regions, maintaining realistic and safe mobility patterns. $\mathcal{C5}$ bounds the cumulative energy consumption of UAVs over the operation period, preserving long-term mission endurance. $\mathcal{C6}$ constrains the traffic prediction error between actual and predicted states, ensuring high forecasting accuracy. $\mathcal{C7}$ similarly bounds the mean-squared error of vehicle trajectory predictions, maintaining reliable trajectory forecasting. Finally, $\mathcal{C8}$ ensures that the communication and processing load on each UAV does not exceed a predefined threshold, thereby promoting energy-aware and stable operation during collaborative decision-making.

IV. Proposed Method: GAN-MADRL

To address the optimization problem in Section III, the UAV-assisted VEC system in 6G-enabled ITS is modeled as a MDP [24]. UAVs, vehicles, and edge servers act as decentralized agents that interact asynchronously with the environment, optimizing energy efficiency, latency, and trajectory accuracy. GAN modules at UAVs and edge nodes generate synthetic traffic and mobility data to overcome non-IID and sparse datasets, improving learning stability and convergence. The proposed GAN-MADRL framework integrates actor–critic learning with asynchronous coordination, enabling scalable, energy-aware, and adaptive optimization. This results in stable convergence, reduced communication overhead, and enhanced robustness for intelligent and sustainable ITS operations in dynamic 6G environments.

A. MDP Formulation

The MDP is represented by the tuple $(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma)$:

- **State ($\mathcal{S}_t \in \mathcal{S}$):** At time $\mathcal{t}$, the system state is defined as

$$s_t = (T_t, V_t, U_t, C_t), \quad (7)$$

where T_t denotes the traffic state (e.g., congestion levels, flow rates), V_t captures vehicle positions and velocities, U_t represents UAV positions and velocities, and C_t includes communication metrics such as bandwidth, SNR, and latency.

- **Action ($a_t \in A$):** The joint action space is

$$a_t = (a_V, a_U, a_E), \quad (8)$$

where a_V includes vehicle routing or speed control, a_U represents UAV trajectory and energy allocation, and a_E denotes edge server actions such as model aggregation and resource scheduling.

- **Transition Function (P):** The transition probability $P(s_{t+1}|s_t, a_t)$ models the stochastic evolution of the system based on vehicular mobility, wireless channel dynamics, and UAV decisions.
- **Reward Function (R):** The reward is designed to balance multiple objectives:

$$R(s_t, a_t) = \alpha_1 R_{\text{traffic}} + \alpha_2 R_{\text{traj}} + \alpha_3 R_{\text{energy}} + \alpha_4 R_{\text{latency}} \quad (9)$$

where R_{traffic} penalizes traffic prediction error, R_{traj} penalizes trajectory forecasting error, R_{energy} accounts for UAV energy consumption, and R_{latency} captures network delay. The coefficients α_i are tunable to prioritize different trade-offs.

- **Discount Factor (γ):** Ensures long-term optimization of cumulative reward, with $0 < \gamma \leq 1$.

B. GAN-Augmented Learning

Each agent i maintains a local dataset D_i , which is often sparse, imbalanced, and non-IID due to heterogeneous vehicular mobility and traffic variations. To overcome these challenges, a Generative Adversarial Network (GAN) is deployed at UAVs and edge nodes to synthesize realistic mobility and traffic samples. The generator G_i learns to mimic real data distributions, while the discriminator D_i distinguishes between real and synthetic data. The adversarial objective is defined as:

$$\min G_i \max D_i \mathbb{E}_{x \sim D_i} [\log D_i(x)] + \mathbb{E}_{z \sim \mathcal{N}(0,1)} [\log(1 - D_i(G_i(z)))]. \quad (10)$$

The generator produces synthetic samples $\hat{x}$, enriching the local dataset as $\hat{D}_i = G_i(z, \phi_i)$. This $\hat{D}_i = G_i(z, \phi_i)$ augmentation mitigates overfitting, accelerates convergence, and improves learning robustness under data scarcity or communication delays. Moreover, the GAN-driven augmentation enhances prediction accuracy for rare traffic states and ensures balanced participation among agents in the distributed learning process.

C. Multi-Agent DRL Optimization

In the proposed GAN-MADRL framework, each agent learns an optimal policy $\pi_i(a_t|s_t)$ to jointly optimize traffic prediction, UAV trajectory, and computational offloading. The value function (critic) is updated iteratively as:

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha [R(s_t, a_t) + \gamma \max_{a'} Q(s_{t+1}, a') - Q(s_t, a_t)], \quad (11)$$

where α is the learning rate and γ the discount factor. The reward $R(s_t, a_t)$ balances four critical objectives: energy efficiency, latency minimization, trajectory accuracy, and traffic prediction reliability, defined as

$$R(s_t, a_t) = \alpha_1 R_{\text{traffic}} + \alpha_2 R_{\text{traj}} + \alpha_3 R_{\text{energy}} + \alpha_4 R_{\text{latency}} \quad (12)$$

where α_i are adaptive weights tuned via meta-learning for dynamic environments. Through decentralized actor-critic training, UAVs, vehicles, and edge servers collaboratively learn energy-efficient and low-

latency policies. The GAN-enhanced state distribution improves value estimation, stabilizing convergence and reducing training oscillations. Figure 2 illustrates the overall architecture and cooperative learning mechanism.




Fig. 2. GAN-enhanced MADRL-Framework.

D. Training Pipeline

The end-to-end GAN-MADRL learning procedure proceeds in iterative stages: label=d)

1. **Data Collection:** UAVs and vehicles collect local data and train GANs to produce synthetic mobility and traffic samples.
2. **Local Update:** Each agent updates actor-critic networks using augmented datasets D_i^{aug} to improve learning under non-IID conditions.
3. **Interaction:** Agents interact with the environment, perform task offloading, and obtain rewards based on energy, latency, and accuracy.
4. **Coordination:** Edge servers perform asynchronous aggregation of local models to ensure scalability and reduce communication load.
5. **Convergence:** The process repeats until global policy convergence, ensuring stable performance and faster adaptation in dynamic ITS scenarios.

This hybrid framework enables efficient learning even with limited data, unstable links, and dynamic network conditions, achieving a balance between accuracy, energy, and latency performance.

E. Convergence and Complexity Analysis

GAN augmentation ensures local data diversity, preventing overfitting and accelerating convergence. The asynchronous nature of MADRL eliminates synchronization bottlenecks common in conventional FL or RL. The communication complexity is reduced to lightweight policy parameter exchanges, while computational complexity scales linearly with the number of

agents. Overall, the framework achieves communication-efficient, computation-light, and highly scalable learning suitable for real-time 6G-enabled ITS.

Algorithm 1 GAN-Enhanced MADRL (Sketch)

Require: Agents $\mathcal{A}$; replay buffers $\{B_i\}$; actor θ_i , critic ω_i ; targets θ_i^-, ω_i^- ; GAN params $\Phi_i, \gamma, \tau, \eta_\pi, \eta_Q$

- 1: **for** episode $t = 1$ to T **do**
 - 2: **for** each agent $i \in \mathcal{A}$ **in parallel do**
 - 3: Observe s_t^i ; update GAN with real data; generate synthetic samples
 - 4: Take $a_t^i \sim \pi_{\theta_i}(s_t^i)$; get r_t^i, s_{t+1}^i ; store transition
 - 5: Update critic: $y = r + \gamma \max_{a'} Q\omega^-(s', a')$; $\omega_i \leftarrow \omega_i - \eta_Q \nabla (Q\omega_i - y)$
 - 6: Update actor: $\theta_i \leftarrow \theta_i + \eta_\pi \nabla_{\theta} Q\omega_i(i, s, \pi_{\theta_i}(s))$
 - 7: Soft targets: $\{\theta_i^-, \omega_i^-\} \leftarrow \tau \{\theta_i, \omega_i\} + (1 - \tau) \{\theta_i^-, \omega_i^-\}$
 - 8: **end for**
 - 9: Periodic async aggregation and broadcast
 - 10: **end for**
-

V. Results and Discussion

A. Simulation Setup

Simulations were performed on a workstation with an Intel Core i7 (2.4 GHz) processor and 16 GB RAM using Python 3.9 with TensorFlow 2.12 and PyTorch 1.13. The 6G-enabled ITS scenario covers a (500×500) m urban area where multiple UAVs fly at altitudes of 100–200 m to coordinate with vehicles and roadside edge servers. The network consists of 50–100 moving vehicles and five RSU-based edge units connected through mmWave and sub-6 GHz V2X links. Each UAV has limited onboard computation and storage, performing traffic prediction, trajectory planning, and offloading tasks. GAN modules embedded at the edge generate synthetic traffic and mobility data to augment local datasets and enhance learning performance during federated training.

B. Comparative Evaluation Baselines

To demonstrate the effectiveness of the proposed GAN-based IFL framework, we evaluate it against four state-of-the-art baseline approaches widely used in intelligent VEC systems. The proposed GAN-IFL is a hybrid learning framework that integrates GANs with IFL. This architecture enables UAVs, vehicles, and edge servers to collaboratively learn with privacy-preserving, communication-efficient, and adaptive policies for traffic prediction, UAV coordination, and trajectory forecasting in real-time 6G ITS environments.

- **CDL (Centralized Deep Learning):** Involves centralized model training where all raw data is transmitted to a central server for processing. It achieves high accuracy but suffers from heavy communication overhead, privacy leakage, and scalability issues (Nawaz et al., 2024).
- **FL (Federated Learning):** Enables decentralized model training where devices collaboratively learn without sharing raw data, thus improving privacy. However, it faces challenges such as non-IID data, slow convergence, and communication asynchrony (Li et al., 2024).
- **PPO (Proximal Policy Optimization):** A policy-gradient reinforcement learning algorithm known for training stability and sample efficiency. Nonetheless, it exhibits limited scalability in federated and multi-agent environments due to coordination complexity (Li et al., 2025).
- **Greedy:** A heuristic approach where each agent selects the immediate best action to maximize short-term rewards without considering long-term effects. It is computationally simple but leads to suboptimal global performance and poor adaptability in dynamic ITS scenarios (Betalo et al., 2021).

C. Performance Parameters

The proposed GAN-based MADRL framework is evaluated against CDL, FL, DQN, and PPO across four key metrics—energy efficiency, latency, prediction accuracy, and convergence. The wireless channel model includes both large-scale path loss and Rayleigh fading, while UAVs and edge nodes follow asynchronous




Fig. 3. Performance comparison of the proposed GAN-MADRL under different conditions: (a) Average reward, (b) Normalized return, and (c) validation accuracy.

GAN-based IFL updates with periodic aggregation. Training spans 3500 episodes under dynamic vehicle densities and task variations to emulate real-time ITS conditions.

- **Average reward:** Mean cumulative return per episode; higher values indicate better learning and policy efficiency.
- **Normalized return:** Scaled cumulative reward (0–1); higher values show faster convergence and stronger policy performance.
- **Validation accuracy:** Measures model performance on unseen data; higher accuracy reflects better generalization and robustness.
- **Trajectory NRMSE:** Normalized error between predicted and actual trajectories; lower values indicate higher prediction precision.
- **Trajectory weight (α_2):** Controls the importance of trajectory accuracy; higher α_2 yields smoother and precise paths.
- **Collision rate:** Frequency of UAV or vehicle collisions; lower rates imply safer and more coordinated navigation.
- **Vehicle density:** Number of vehicles per road segment; higher density increases traffic complexity and coordination demands.

D. Convergence and Stability Analysis

The proposed GAN-MADRL framework achieves stable and fast convergence through decentralized cooperative learning, GAN-based data augmentation, and adaptive reward shaping. These mechanisms enhance exploration, reduce variance, and prevent policy oscillations, ensuring consistent learning performance across dynamic 6G-enabled ITS scenarios, as illustrated in Figure 3.

The average reward increases faster in our GAN-MADRL method because GAN-based data augmentation and decentralized learning enhance exploration, stability, and policy convergence compared to other baselines as shown in Figure 3a.

In Figure 3b, the normalized return increases with episodes as agents progressively learn optimal policies through repeated exploration and reward feedback. The learning rate of $2e-5$ yields better stability and convergence because smaller steps prevent overshooting and oscillations, enabling smoother policy updates and more consistent long-term performance.

The validation accuracy increases with episodes because the GAN-MADRL agents continuously refine their policies and predictions through repeated training and feedback. Over time, GAN-based data

augmentation enhances generalization, while multi-agent collaboration reduces overfitting—leading to higher and more stable validation accuracy compared to other methods as shown in Figure 3c.

E. Trajectory Planning and Prediction Accuracy Analysis

The proposed GAN-MADRL framework enhances trajectory planning and prediction accuracy through cooperative learning among UAVs, vehicles, and edge nodes, as shown in Figure 4. GAN-based data




Fig. 4. Performance comparison of the proposed GAN-MADRL under different conditions: (a) Vehicle per minute vs. traffic weight α_1 , (b) Traffic RMSE vs. traffic weight α_1 , (c) Final position error vs. trajectory weight α_2 , and (d) Collision rate vs. vehicle density.

Augmentation enables the model to capture realistic traffic dynamics, improving trajectory forecasting under varying mobility conditions. The multi-agent reinforcement learning design ensures adaptive, energy-efficient UAV path optimization with low latency and accurate task allocation. As illustrated in Figure 5, simulation results confirm that the proposed framework consistently outperforms CDL, FL, PPO, and Greedy baselines in trajectory smoothness, prediction accuracy, and convergence stability within real-time 6G-enabled ITS environments.

The vehicles per minute increase with traffic weight because a higher traffic weight (α_1) represents denser traffic conditions and greater data generation rates, as shown in Figure 4a. As α_1 increases, the proposed GAN-MADRL model efficiently allocates UAV and edge resources, maintaining smooth traffic flow prediction and adaptive offloading, leading to a proportional rise in vehicle throughput and better network utilization.

The trajectory NRMSE decreases as traffic weight (α_1) increases because higher α_1 provides richer traffic information, enabling the proposed GAN-MADRL framework to learn more accurate motion patterns. GAN-based data augmentation and multi-agent coordination enhance trajectory prediction accuracy, reducing estimation errors and ensuring smoother, more reliable UAV path planning compared to baseline methods, as shown in Figure 4b.

The final position error decreases with increasing trajectory weight (α_2) because a higher α_2 emphasizes precise path-following during training, as shown in Figure 4c. In the proposed GAN-MADRL framework,

this strengthens trajectory optimization, leading to more accurate UAV endpoint alignment and reduced deviation compared to CDL, FL, and PPO baselines.

In Figure 4d, the collision rate increases with vehicle density because higher density leads to more frequent interactions and limited maneuvering space. However, the proposed GAN-MADRL framework significantly reduces collisions by enabling cooperative trajectory planning, real-time coordination, and adaptive decision-making among UAVs and vehicles under dense traffic conditions.

The Traffic RMSE decreases as the Non-IID level (α) increases because higher α values indicate more homogeneous data distribution across agents, as shown in Figure 5a. In the proposed GAN-MADRL framework, this improves collaborative learning consistency and stabilizes model updates, leading to lower prediction errors compared to CDL, FL, and PPO baselines under heterogeneous traffic conditions.




Fig. 5. Performance comparison of the proposed GAN-MADRL under different conditions: (a) Traffic RMSE vs. Non-IID level α (b) Traffic RMSE vs. Traffic weight α_1 (c) Traffic RMSE vs. sensors, and (d) Trajectory NRMSE vs. prediction horizon.

In Figure 5b, the Traffic RMSE decreases with increasing traffic weight (α_1) because higher (α_1) emphasizes traffic-related features during learning. In the proposed GAN-MADRL framework, this prioritization enhances traffic pattern recognition and prediction accuracy, resulting in lower RMSE compared to CDL, FL, and PPO baselines.

The Traffic RMSE decreases as the number of sensors increases because more sensors provide richer and more diverse traffic data. In the proposed GAN-MADRL framework, this enhanced data availability improves learning accuracy and environmental awareness, leading to more precise traffic prediction and lower RMSE than CDL, FL, and PPO baselines as shown in Figure 5c.

The trajectory NRMSE increases with a longer prediction horizon because predicting farther into the future introduces greater uncertainty and accumulated errors as shown in Figure 5d. However, the proposed GAN-MADRL framework maintains lower NRMSE than baseline methods by leveraging GAN-based data augmentation and multi-agent coordination, which enhance long-term trajectory prediction stability and accuracy.

F. UAV Energy-Efficiency Analysis

The proposed GAN-MADRL framework enhances UAV energy efficiency by jointly optimizing trajectory planning, computational offloading, and traffic prediction. Through decentralized learning and GAN-based data augmentation, UAVs develop energy-aware policies that reduce flight distance, idle time, and redundant communication. The adaptive reward design balances energy usage and performance, enabling sustained UAV operation under dynamic 6G ITS environments, as illustrated in Figure 6.

The average energy consumption decreases as computational offloading increases because more tasks are transferred from UAVs to nearby edge servers, reducing onboard processing load and flight energy usage as shown in Figure 6a. In the proposed GAN-MADRL framework, adaptive offloading and energy-aware learning enable optimal task distribution, achieving lower energy consumption compared to CDL, FL, PPO, and Greedy methods.

In Figure 6b, the end-to-end latency decreases with increasing computational offloading because more processing tasks are executed at powerful edge servers instead of resource-limited UAVs. In the proposed GAN-MADRL framework, intelligent offloading and cooperative scheduling minimize transmission delay and processing time, achieving lower latency than CDL, FL, PPO, and Greedy baselines.




Fig. 6. Performance comparison of the proposed GAN-MADRL under different conditions: (a) Average energy consumption vs. Computational offloading b) End-to-end latency. Computational offloading c) Average energy consumption vs. number of agents, and d) Energy efficiency vs. offloading capacity.

The average energy consumption decreases as the number of agents increases because cooperative task sharing and distributed learning reduce the workload per UAV. In the proposed GAN-MADRL framework, multi-agent coordination enables balanced energy usage and efficient resource allocation, leading to lower overall energy consumption compared to CDL, FL, PPO, and Greedy approaches as shown in Figure 6c.

The energy efficiency increases with higher offloading capacity because more computation is transferred to edge servers, reducing UAV processing and transmission energy as shown in Figure 6d. In the proposed GANMADRL framework, adaptive offloading and cooperative learning maximize energy utilization, achieving superior efficiency compared to CDL, FL, PPO, and Greedy baselines.

VI. Conclusion

This paper proposed a GAN-MADRL framework for UAV-assisted VEC in 6G-enabled ITS. The framework jointly optimizes TPCO through a multi-agent MDP. By integrating GAN-based data

augmentation and asynchronous, decentralized learning, the proposed approach enhances model generalization, convergence stability, and scalability under non-IID and dynamic traffic conditions. Extensive simulations demonstrate that the GAN-MADRL framework significantly outperforms CDL, FL, PPO, and Greedy baselines in terms of energy efficiency, latency reduction, trajectory accuracy, and traffic-prediction performance. In future work, we aim to extend this framework toward privacy-preserving and trust-aware collaborative intelligence by incorporating blockchain-based secure aggregation and edge-intelligent digital twins. Furthermore, expanding to UAV-swarm coordination with proactive mobility prediction will further improve adaptability, real-time decision-making, and resilience in large-scale 6G-enabled ITS environments.

References

- Ali, A., Ullah, I., Ahmad, S., Wu, Z., Li, J., & Bai, X. (2025). An attention-driven spatio-temporal deep hybrid neural network for traffic flow prediction in transportation systems. *IEEE Transactions on Intelligent Transportation Systems*.
- Ali, W., Din, I. U., Almogren, A., & Rodrigues, J. J. (2024). Federated learning-based privacy-aware location prediction model for Internet of Vehicular Things. *IEEE Transactions on Vehicular Technology*.
- Betalo, M. L., Leng, S., Abishu, H. N., Dharejo, F. A., Seid, A. M., Erbad, A., Naqvi, R. A., Zhou, L., & Guizani, M. (2023). Multi-agent deep reinforcement learning-based task scheduling and resource sharing for O-RAN-empowered multi-UAV-assisted wireless sensor networks. *IEEE Transactions on Vehicular Technology*.
- Betalo, M. L., Leng, S., Abishu, H. N., Seid, A. M., Fakirah, M., Erbad, A., & Guizani, M. (2024). Multi-agent DRL-based energy harvesting for freshness of data in UAV-assisted wireless sensor networks. *IEEE Transactions on Network and Service Management*.
- Betalo, M. L., Leng, S., Chen, X., & Zhou, L. (2021). Joint optimization for cluster head selection in UAV-assisted WSN. In *2021 International Conference on UK-China Emerging Technologies (UCET)* (pp. 31–36). IEEE.
- Betalo, M. L., Leng, S., Seid, A. M., Abishu, H. N., Erbad, A., & Bai, X. (2025). Dynamic charging and path planning for UAV-powered rechargeable WSNs using multi-agent deep reinforcement learning. *IEEE Transactions on Automation Science and Engineering*.
- Betalo, M. L., Leng, S., Zhou, L., & Fakirah, M. (2022). Multi-UAV data collection optimization for sink node and trajectory planning in WSN. In *2022 IEEE 2nd International Conference on Computer Communication and Artificial Intelligence (CCAI)* (pp. 1–7). IEEE.
- Betalo, M. L., Ullah, I., Tesema, F. B., Wu, Z., Li, J., & Bai, X. (2025). Generative AI-driven multi-agent DRL for task allocation in UAV-assisted EMPD within 6G-enabled SAGIN networks. *IEEE Internet of Things Journal*.
- Chen, Z., Yi, W., Shin, H., & Nallanathan, A. (2024). Adaptive semi-asynchronous federated learning over wireless networks. *IEEE Transactions on Communications*.
- Creß, C., Bing, Z., & Knoll, A. C. (2023). Intelligent transportation systems using roadside infrastructure: A literature survey. *IEEE Transactions on Intelligent Transportation Systems*, 25(7), 6309–6327.
- Fakirah, M., Leng, S., Mohammad, M., & Betalo, M. L. (2023). Hybrid VLC/RF misalignment system for autonomous vehicles at roundabouts. In *2023 6th International Conference on Electronics Technology (ICET)* (pp. 581–586). IEEE.
- Guan, Y., Zou, S., Peng, H., Ni, W., Sun, Y., & Gao, H. (2023). Cooperative UAV trajectory design for disaster area emergency communications: A multiagent PPO method. *IEEE Internet of Things Journal*, 11(5), 8848–8859.

- Han, Z., Wu, Y., Li, T., Zhang, L., Pei, L., Xu, L., Li, C., Ma, C., Xu, C., Shen, S., et al. (2023). An efficient spatial-temporal trajectory planner for autonomous vehicles in unstructured environments. *IEEE Transactions on Intelligent Transportation Systems*, 25(2), 1797–1814.
- Hu, J., Chen, C., Cai, L., Khosravi, M. R., Pei, Q., & Wan, S. (2021). UAV-assisted vehicular edge computing for the 6G Internet of Vehicles: Architecture, intelligence, and challenges. *IEEE Communications Standards Magazine*, 5(2), 12–18.
- Jha, A. V., Appasani, B., Khan, M. S., Zeadally, S., & Katib, I. (2024). 6G for intelligent transportation systems: Standards, technologies, and challenges. *Telecommunication Systems*, 86(2), 241–268.
- Lei, Y., Wang, S. L., Zhong, M., Wang, M., & Ng, T. F. (2022). A federated learning framework based on incremental weighting and diversity selection for Internet of Vehicles. *Electronics*, 11(22), Article 3668.
- Li, C., Wu, J., Zhang, Y., & Wan, S. (2024). Energy-latency tradeoffs for joint optimization of vehicle selection and resource allocation in UAV-assisted vehicular edge computing. *IEEE Transactions on Green Communications and Networking*.
- Li, P., Xiao, Z., Gao, H., Wang, X., & Wang, Y. (2025). Reinforcement learning-based edge-end collaboration for multi-task scheduling in 6G-enabled intelligent autonomous transport systems. *IEEE Transactions on Intelligent Transportation Systems*.
- Li, Z., Huang, C., Gai, K., Lu, Z., Wu, J., Chen, L., Xu, Y., & Choo, K.-K. R. (2022). AsyFed: Accelerated federated learning with asynchronous communication mechanism. *IEEE Internet of Things Journal*, 10(10), 8670–8683.
- Liu, R., Hua, M., Guan, K., Wang, X., Zhang, L., Mao, T., Zhang, D., Wu, Q., & Jamalipour, A. (2024). 6G-enabled advanced transportation systems. *IEEE Transactions on Intelligent Transportation Systems*.
- Mohammadisarab, A., Nouruzi, A., Khalili, A., Mokari, N., Arand, B. A., & Jorswieck, E. A. (2024). Resilient disaster relief in industrial IoT: UAV trajectory design and resource allocation in 6G non-terrestrial networks. *IEEE Open Journal of the Communications Society*, 5, 1827–1845.
- Nawaz, M. W., Zhang, W., Flynn, D., Zhang, L., Swash, R., Abbasi, Q. H., Imran, M. A., & Popoola, O. (2024). 6G edge-networks and multi-UAV knowledge fusion for urban autonomous vehicles. *Physical Communication*, 67, Article 102479.
- Rezaee, M. R., Hamid, N. A. W. A., Hussin, M., & Zukarnain, Z. A. (2024). Comprehensive review of drones collision avoidance schemes: Challenges and open issues. *IEEE Transactions on Intelligent Transportation Systems*.
- Shinde, S. S., & Tarchi, D. (2023). Joint air-ground distributed federated learning for intelligent transportation systems. *IEEE Transactions on Intelligent Transportation Systems*, 24(9), 9996–10011.
- Wang, X., Li, K., & Chehri, A. (2023). Multi-sensor fusion technology for 3D object detection in autonomous driving: A review. *IEEE Transactions on Intelligent Transportation Systems*, 25(2), 1148–1165.
- Yan, J., Zhao, X., & Li, Z. (2024). Deep reinforcement learning-based computation offloading in UAV-assisted vehicular edge computing networks. *IEEE Internet of Things Journal*.
- Zhang, X., Chang, Z., Hu, T., Chen, W., Zhang, X., & Min, G. (2023). Vehicle selection and resource allocation for federated learning-assisted vehicular network. *IEEE Transactions on Mobile Computing*, 23(5), 3817–3829.