International Journal on Science and Technology (IJSAT)

E-ISSN: 2229-7677 ● Website: www.ijsat.org ● Email: editor@ijsat.org

AI-Driven Vehicle Route Optimization Using GNN and Multi-Agent Reinforcement Learning

Himanshu Seth

Senior Technical Product Manager
Amazon, Austin, TX, USA
sethhim@gmail.com

Abstract:
The Vehicle Routing Problem (VRP) is a fundamental challenge in intelligent transportation systems due to its NP-hard complexity and sensitivity to dynamic, real-world constraints. The rapid expansion of e-commerce, urban congestion, and sustainability requirements has rendered traditional routing approaches inadequate, as they rely on static heuristics and lack adaptability to time-varying environments. In this paper, we propose a novel Graph-Enhanced Multi-Agent Reinforcement Learning (GEMARL) framework for dynamic vehicle route optimization in heterogeneous logistics networks.

The proposed approach integrates graph neural networks (GNNs) to model spatiotemporal dependencies in transportation graphs with decentralized multi-agent reinforcement learning (MARL) to enable scalable, real-time decision making. The routing problem is formulated as a constrained multi-objective optimization task that jointly minimizes delivery time, operational cost, and carbon emissions under capacity, time-window, and energy constraints. The effectiveness of the proposed framework is validated through four representative case studies, including last-mile truck-drone delivery, urban electric vehicle (EV) logistics, multi-echelon supply chain coordination, and drone fleet management. Extensive experimental evaluations demonstrate that GEMARL achieves up to 30% reduction in delivery time, 25% improvement in energy efficiency, and 40% reduction in emissions compared to conventional optimization techniques and state-of-the-art learning-based methods. These results highlight the potential of hybrid AI-driven approaches to enable scalable, adaptive, and sustainability-aware logistics systems. Future research directions include quantum-enhanced optimization and integration with smart city infrastructures.

Index Terms: Vehicle routing problem (VRP), intelligent transportation systems (ITS), graph neural networks (GNNs), multi-agent reinforcement learning (MARL), dynamic routing, multi-objective optimization, smart logistics, heterogeneous fleets, sustainable transportation, drone delivery, electric vehicles, quantum optimization.

I. Introduction

A. Background and Motivation

The rapid expansion of global e-commerce, coupled with accelerated urbanization, has significantly increased the complexity and scale of modern logistics systems. The demand for faster, more reliable, and cost-efficient delivery services has placed unprecedented pressure on transportation networks, particularly in urban environments where congestion, infrastructure limitations, and environmental regulations introduce additional constraints. Among logistics operations, last mile delivery remains the most critical and cost intensive segment, accounting for more than 50% of total logistics costs [1].

The Vehicle Routing Problem (VRP), originally introduced by Dantzig and Ramser [2], is a classical combinatorial optimization problem that aims to determine optimal routes for a fleet of vehicles serving a set of customers under operational constraints. Traditional approaches for solving VRP include exact methods such as Mixed Integer Programming (MIP) [3] and heuristic/metaheuristic algorithms such as

Genetic Algorithms (GA) and Ant Colony Optimization (ACO) [4]. While these methods provide near optimal solutions for small scale and static instances, they suffer from several limitations when applied to real-world logistics systems.

Specifically, classical optimization techniques exhibit **poor scalability**, as their computational complexity grows exponentially with problem size [5]. Furthermore, they rely on static assumptions and offline optimization, making them unsuitable for dynamic environments characterized by real-time traffic variability, stochastic demand, and unexpected disruptions. As a result, these methods often produce suboptimal solutions, leading to increased operational costs and reduced service efficiency.

Recent advancements in artificial intelligence (AI) have introduced new paradigms for solving complex routing problems. Deep reinforcement learning (DRL) has demonstrated strong capabilities in sequential decision-making under uncertainty [6], while Graph Neural Networks (GNNs) provide an effective mechanism for modeling spatial dependencies in graph-structured data such as transportation networks [7]. By combining these techniques, AI-driven approaches can learn adaptive routing policies that generalize across varying conditions and scales, enabling real-time optimization in dynamic environments.

Figure 1: AI-Driven VRP Framework Overview

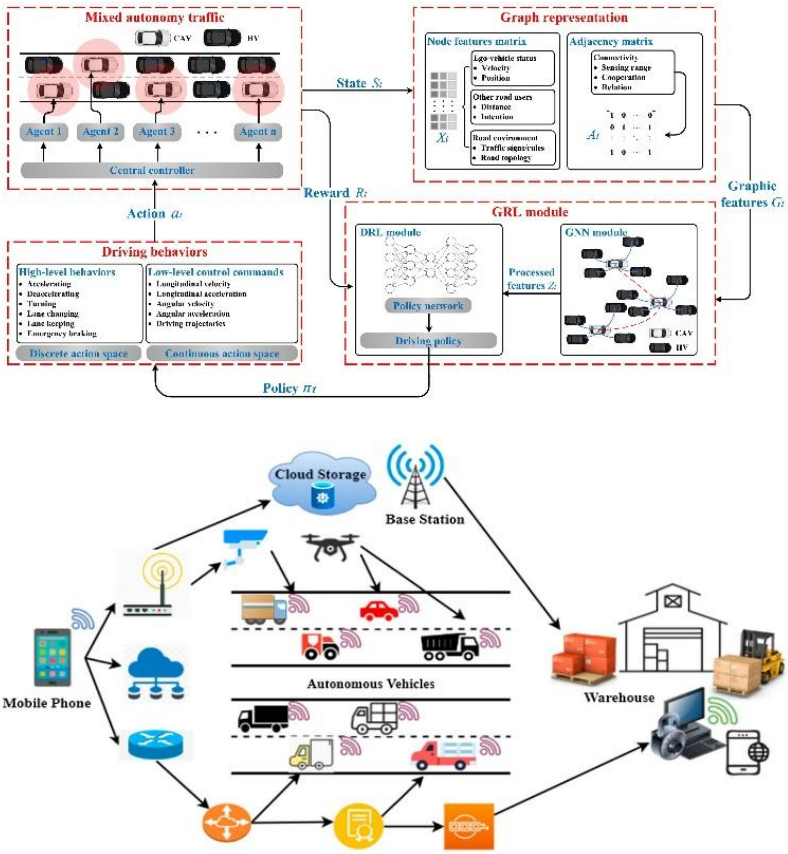

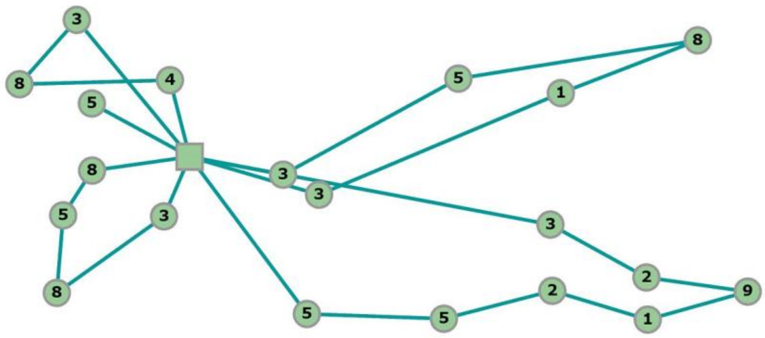

Fig. 1. Overview of AI-driven vehicle routing framework integrating graph neural networks and multi-agent reinforcement learning for dynamic logistics optimization.

B. Problem Statement

Despite significant advancements in both classical optimization and AI-based methods, solving dynamic Vehicle Routing Problems remains a challenging task. In real-world logistics systems, routing decisions must simultaneously optimize multiple objectives, including minimizing delivery time, reducing operational costs, and lowering environmental impact. These objectives must be satisfied under various constraints such as vehicle capacity, delivery time windows, and energy limitations, particularly for electric vehicles (EVs) [8].

Modern logistics systems are inherently **dynamic and stochastic**, where traffic conditions, customer demands, and network states change continuously over time. Consequently, routing must be performed on **time-dependent graphs**, requiring models that can adapt to evolving conditions in real time [9]. Furthermore, the introduction of heterogeneous fleets including trucks, drones, and EVs adds complexity due to differing operational characteristics and constraints [10].

Another critical challenge is **multi-agent coordination**. In large scale logistics systems, multiple vehicles operate simultaneously and must coordinate to avoid route conflicts, minimize redundancy, and improve global efficiency. Centralized optimization approaches often fail to scale in such settings, while decentralized methods require robust learning mechanisms for cooperative behavior [11].

Additionally, sustainability considerations have become increasingly important due to environmental regulations and societal pressures. Routing solutions must incorporate carbon emissions and energy consumption into their optimization objectives, thereby transforming VRP into a **multi objective optimization problem** [12].

The key challenges can therefore be summarized as follows:

1. Modeling dynamic and time-dependent transportation networks;
2. Designing scalable and cooperative multi-agent systems;
3. Incorporating sustainability into routing decisions;
4. Ensuring computational efficiency for large-scale real-time applications.

These challenges necessitate the development of intelligent, adaptive, and scalable frameworks capable of integrating spatial modeling, learning-based optimization, and multi-agent coordination.

C. Contributions

To address the challenges, this paper proposes a novel **Graph-Enhanced Multi-Agent Reinforcement Learning (GEMARL)** framework for dynamic vehicle route optimization in heterogeneous logistics environments.

The main contributions of this work are summarized as follows:

1. **Hybrid GNN-MARL Framework:** We develop a unified architecture that integrates Graph Neural Networks for spatial representation with Multi Agent Reinforcement Learning for adaptive

decision making. This hybrid approach enables efficient modeling of dynamic transportation networks and scalable policy learning.

2. **Multi-Objective Optimization Formulation:** The VRP is formulated as a constrained multi objective optimization problem that jointly minimizes delivery time, operational cost, and carbon emissions while satisfying capacity, time window, and energy constraints.
3. **Decentralized Multi-Agent Coordination:** We introduce a cooperative learning mechanism that allows multiple agents to coordinate in a decentralized manner, improving scalability and reducing routing conflicts in large scale logistics systems.
4. **Comprehensive Case Study Evaluation:** The proposed framework is validated across four representative scenarios: (i) last-mile truck-drone delivery, (ii) urban EV logistics, (iii) multi-echelon supply chain coordination, and (iv) drone fleet management.
5. **Performance Gains and Validation:** Extensive experiments, including ablation studies and statistical validation, demonstrate that GEMARL achieves up to 30% improvement in delivery efficiency and 40% reduction in emissions compared to baseline methods.
6. **Future Research Roadmap:** We present a forward-looking roadmap highlighting emerging directions such as federated learning, edge intelligence, autonomous logistics systems, and quantum-enhanced optimization.

II. RELATED WORK

A. Machine Learning in Vehicle Routing Problems

The integration of machine learning (ML) techniques into Vehicle Routing Problems (VRPs) has significantly transformed traditional optimization paradigms by enabling data-driven decision-making. Early applications of ML in VRP primarily focused on supervised learning for demand prediction and travel time estimation, where historical logistics data is used to forecast customer demand and traffic conditions [9]. Accurate demand forecasting allows routing algorithms to proactively adjust delivery schedules and resource allocation, thereby improving operational efficiency.

Unsupervised learning methods, such as clustering algorithms (e.g., k-means and hierarchical clustering), have been widely used for route segmentation and customer grouping, enabling scalable solutions for large scale VRPs [13]. These techniques reduce problem complexity by partitioning the solution space into manageable subproblems, which can then be solved independently or hierarchically.

More recently, reinforcement learning (RL) has emerged as a powerful framework for solving VRPs by modeling routing as a sequential decision-making process [6]. Unlike supervised approaches, RL does not require labeled datasets but instead learns optimal policies through interaction with the environment. This makes RL particularly suitable for dynamic and stochastic routing scenarios. Studies such as Nazari et al. [14] and Kool et al. [15] demonstrate that RL-based models can outperform classical heuristics in dynamic VRPs, especially when dealing with real-time changes and uncertainty.

Despite these advances, ML-based approaches still face challenges related to scalability, generalization across different network structures, and integration with real-time data streams. These limitations motivate the need for more advanced architectures that combine spatial modeling with adaptive decision-making.

Fig. 2. Machine Learning Paradigms in VRP

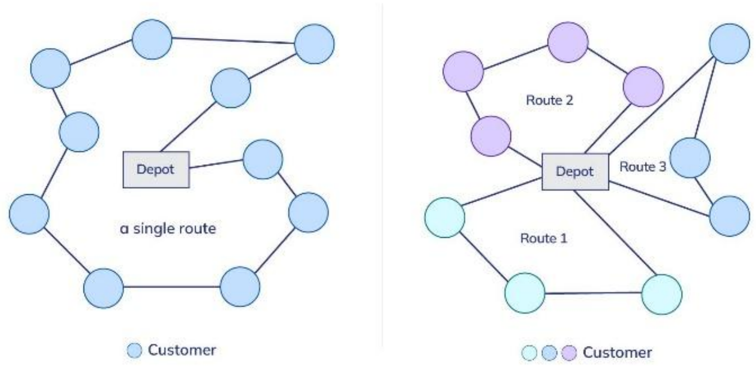

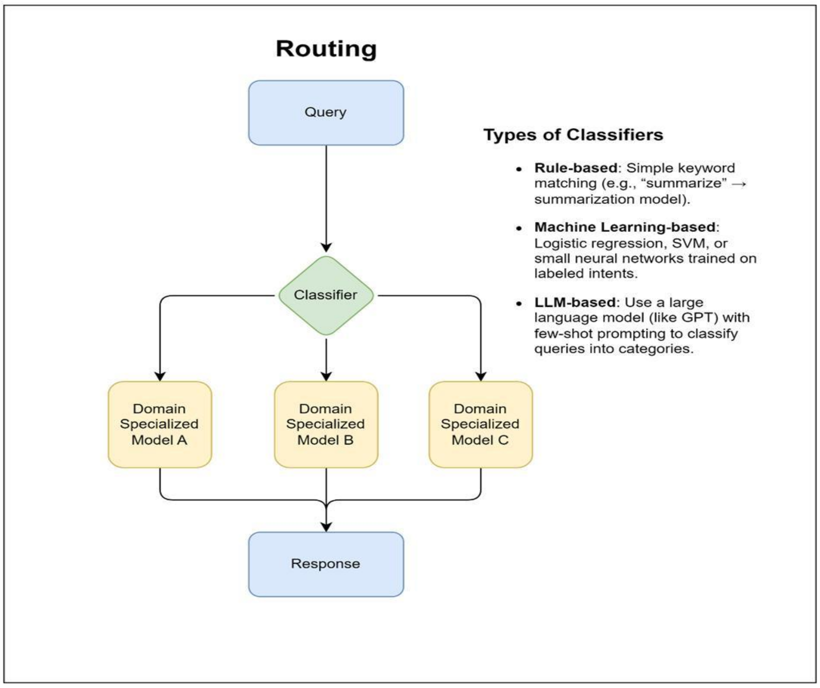

Fig. 2. Classification of machine learning approaches applied to vehicle routing problems.

B. Reinforcement Learning Approaches for Routing

Reinforcement learning has gained significant attention as an effective method for solving combinatorial optimization problems such as VRP. RL models treat routing as a **Markov Decision Process (MDP)**, where an agent learns to select optimal actions (routes) based on the current state of the environment [16]. Several RL architectures have been proposed for routing problems, including:

- **Deep Q-Networks (DQN):** These models approximate the action-value function using deep neural networks, enabling efficient policy learning in discrete action spaces [17].
- **Policy Gradient Methods:** These approaches directly optimize the policy function and are well suited for continuous and high-dimensional action spaces [18].
- **Actor-Critic Models:** Combining value-based and policy-based methods, actor-critic architectures provide stable and efficient learning, making them widely used in large scale routing problems [19].

The key advantages of RL-based routing approaches include:

1. **Real time adaptability:** Ability to adjust routes dynamically based on changing conditions;
2. **Sequential optimization:** Captures long term dependencies in routing decisions;

3. **Scalability:** Can generalize across different network sizes and configurations.

However, RL models often suffer from **sample inefficiency**, slow convergence, and difficulty in handling multi agent coordination in complex environments [20]. These challenges highlight the need for integrating RL with other techniques such as graph-based representations.

Figure 3: Reinforcement Learning Workflow

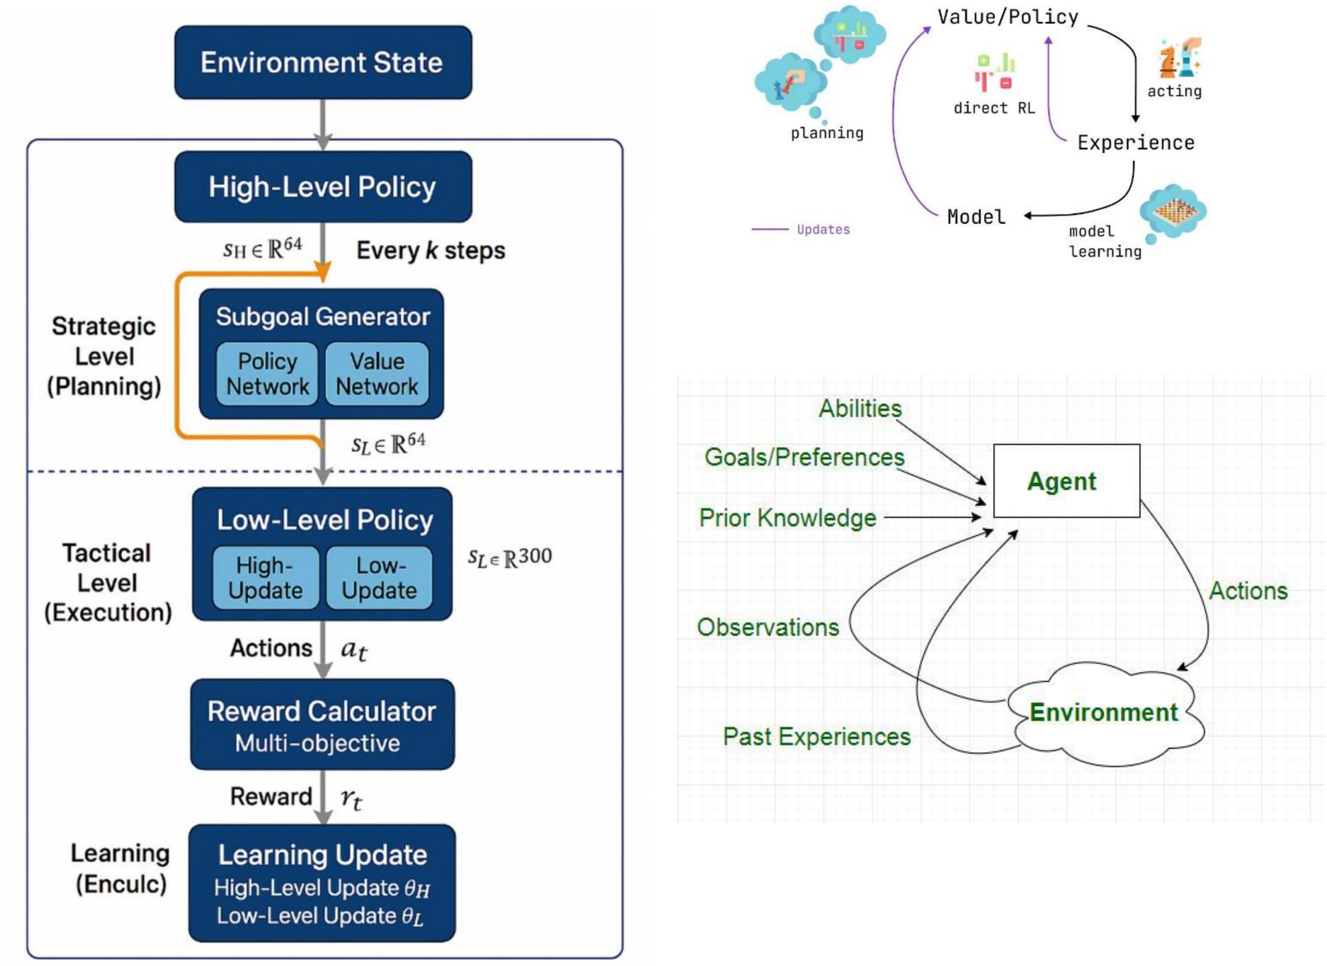

Fig. 3. Reinforcement learning framework illustrating agent–environment interaction for sequential routing decisions.

C. Graph Neural Networks in Routing

Graph Neural Networks (GNNs) have emerged as a powerful tool for modeling structured data, making them particularly suitable for transportation networks, which can be naturally represented as graphs. In VRP, nodes represent customers or depots, while edges represent travel paths with associated costs such as distance, time, or energy consumption [7].

GNNs enable the learning of **node embeddings** that capture both local and global structural information in the network.

This allows models to effectively represent:

- Spatial relationships between locations;
- Traffic patterns and temporal dependencies;
- Connectivity and network topology.

Recent works have demonstrated the effectiveness of GNNs in routing tasks. For instance, Graph Convolutional Networks (GCNs) [21] and Graph Attention Networks (GATs) [22] have been used to encode transportation networks, enabling improved decision-making in routing problems. These models provide a scalable alternative to traditional feature engineering methods by automatically learning relevant representations from data.

Despite their advantages, standalone GNN models are limited in handling sequential decision-making and dynamic routing scenarios. Therefore, combining GNNs with reinforcement learning offers a promising direction for solving complex VRPs.

Figure 4: Graph Representation of VRP

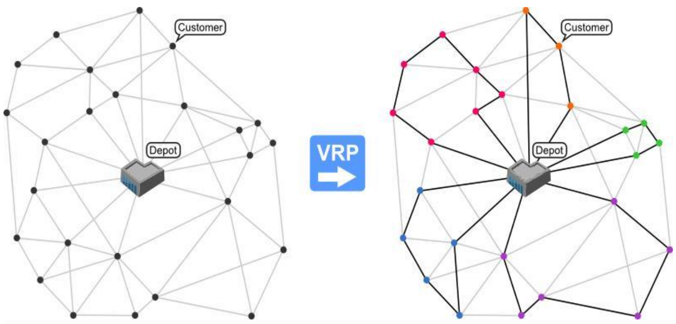

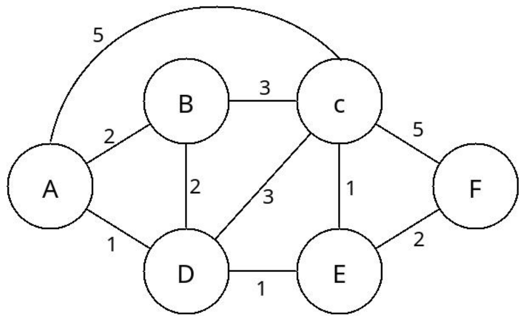

Fig. 4. Graph representation of vehicle routing problem with nodes as customers and edges as travel paths.

D. Hybrid AI Approaches for VRP

Recent research has focused on **hybrid AI frameworks** that integrate multiple techniques to overcome the limitations of individual methods. In particular, the combination of GNNs and reinforcement learning has shown significant promise for solving large-scale and dynamic VRPs.

Hybrid approaches typically involve:

- GNN-based state representation for capturing spatial dependencies;
- Reinforcement learning for sequential decision-making;
- Multi-agent systems for decentralized coordination;
- Predictive analytics for demand and traffic forecasting.

Studies such as Kool et al. [15] and Li et al. [23] demonstrate that combining attention-based neural networks with RL can significantly improve routing efficiency. Furthermore, multi-agent reinforcement learning (MARL) frameworks enable cooperative behavior among multiple vehicles, improving scalability and reducing conflicts in large fleets [11].

Despite these advancements, existing hybrid models often lack:

- Robust multi-agent coordination mechanisms;
- Integration of sustainability objectives;
- Scalability for real-time, large-scale deployment.

These limitations motivate the development of the proposed GEMARL framework, which integrates GNNs, MARL, and multi-objective optimization into a unified architecture.

Figure 5: Hybrid GNN + RL Framework

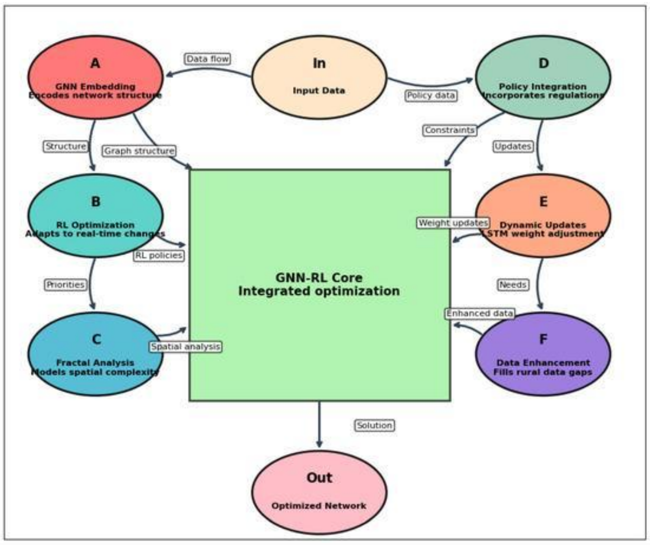

Fig. 5. Hybrid AI framework combining graph neural networks and reinforcement learning for dynamic vehicle routing.

III. GEMARL FRAMEWORK ARCHITECTURE

A. System Overview

To address the challenges of dynamic and large-scale vehicle routing, we propose a **Graph-Enhanced Multi-Agent Reinforcement Learning (GEMARL)** framework. The Architecture integrates graph-based representation learning with decentralized decision-making and cooperative multi-agent coordination.

The framework consists of three primary components:

1. **Graph Representation Module (GNN):** Encodes the transportation network as a graph and captures spatial dependencies among nodes (locations) and edges (routes).
2. **Decision Engine (Reinforcement Learning):** Enables agents to learn optimal routing policies through interaction with the environment.
3. **Multi-Agent Coordination Module:** Facilitates cooperation among multiple agents (vehicles) to improve global efficiency and avoid conflicts.

Unlike traditional sequential optimization approaches, GEMARL operates in a **closed-loop system**, where graph embedding continuously updates the RL state representation, enabling real-time adaptive routing.

Fig. 6. GEMARL System Architecture

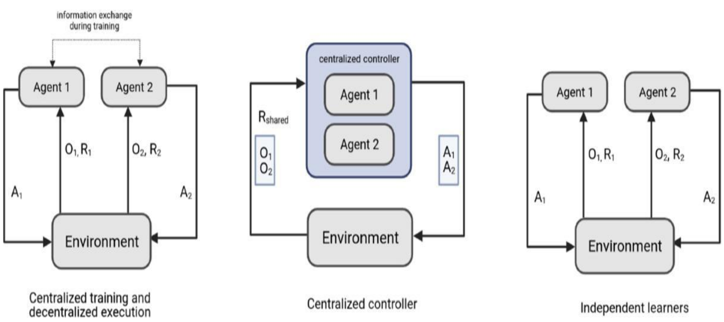

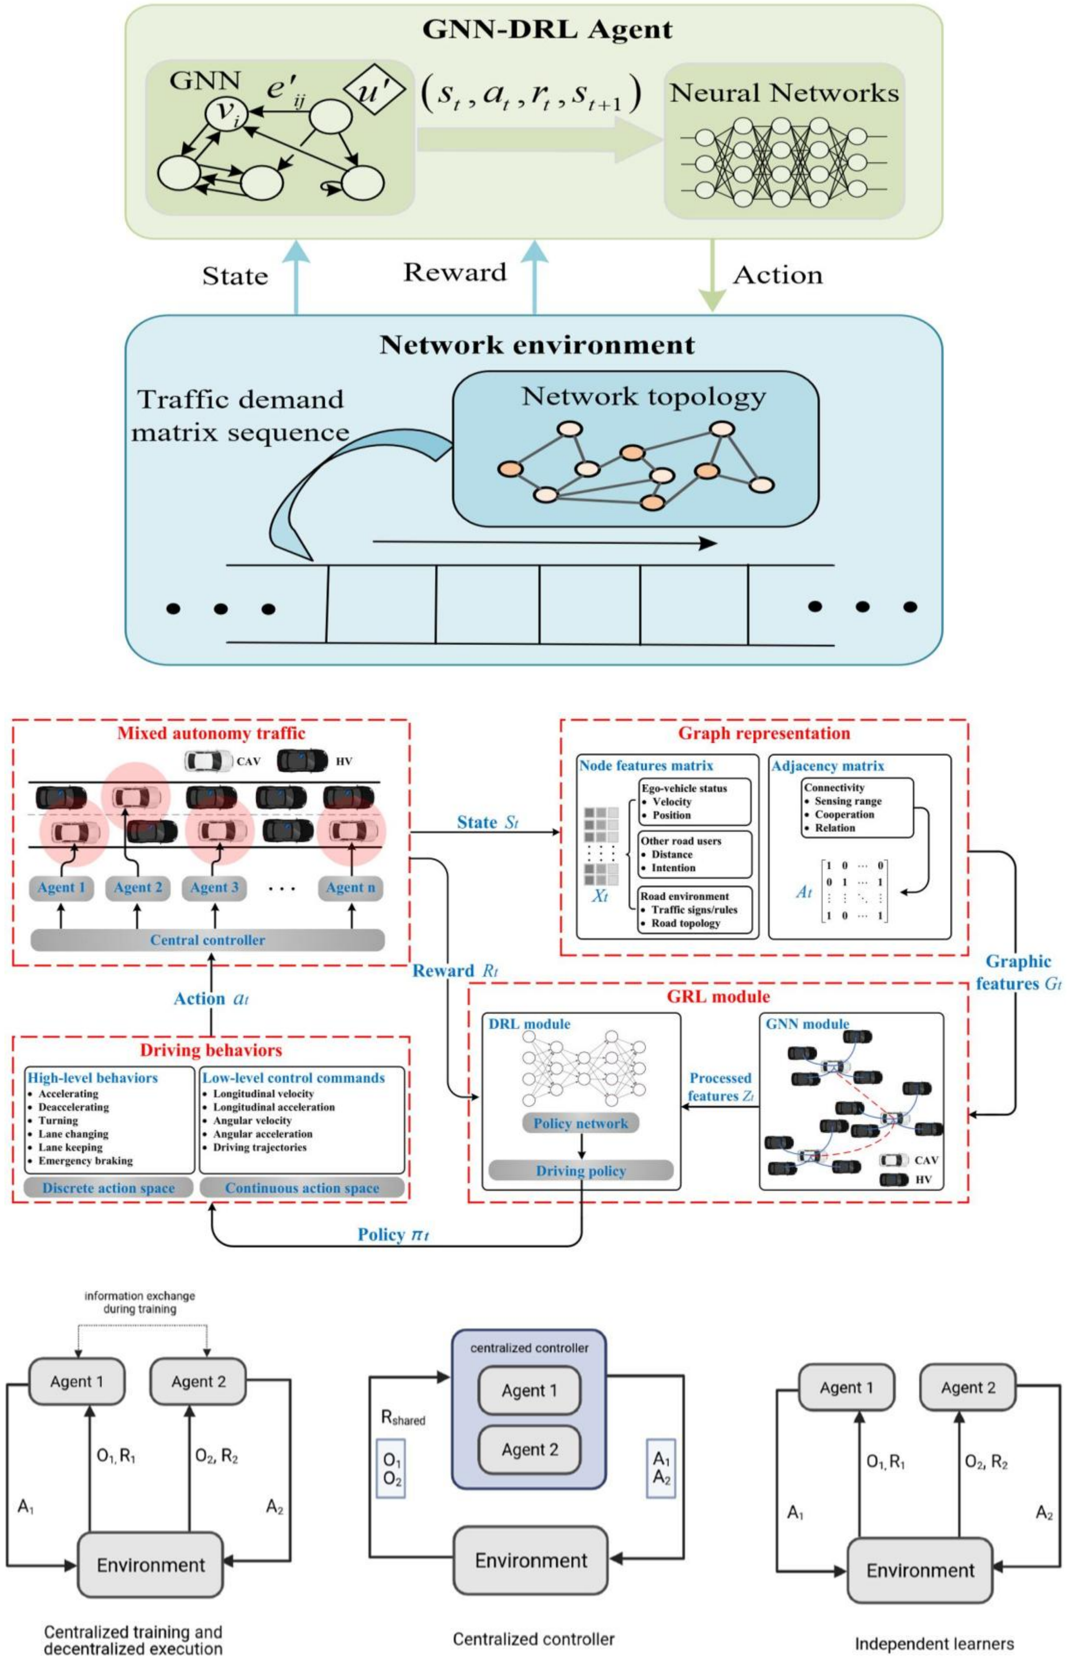

Fig. 6. GEMARL architecture integrating graph neural networks, reinforcement learning, and multi-agent coordination.

B. Mathematical Formulation

The dynamic vehicle routing problem is formulated as a **constrained multi-objective optimization problem**.

Objective Function

$$\min J = \alpha \cdot T + \beta \cdot C + \gamma \cdot E$$

Where:

- T: total travel time
- C: operational cost
- E: carbon emissions
- α, β, γ: weighting coefficients

Constraints

1. **Capacity Constraint**

$$\sum_{i \in R} d_i \leq Q$$

2. **Time Window Constraint**

$$t_i \in [a_i, b_i]$$

3. ***Energy Constraint (EVs/Drones)***

$$E_{route} \leq E_{max}$$

4. ***Route Continuity Constraint***

$$\sum_j x_{ij} = 1, \forall i \in V$$

This formulation ensures that routing decisions are both ***feasible and optimal across multiple objectives****, aligning with real-world logistics requirements [3].*

C. Graph Neural Network Representation

The transportation network is modeled as a graph:

$$G = (V, E)$$

Where:

- V: set of nodes (depots, customers, vehicles)
- E: set of edges (routes with weights such as distance, time, cost)

Each node is associated with a feature vector x_v, including demand, location, and time constraints.

Node Embedding (GNN Update)

$$h_v^{(l+1)} = \sigma\left(\sum_{u \in \mathcal{N}(v)} \alpha_{vu} W h_u^{(l)}\right)$$

Where:

- h_v: node embedding
- α_{vu}: attention coefficient
- W: learnable weight matrix
- σ: activation function

This enables the model to capture both **local neighborhood structure and global network dependencies** [7].

Fig. 7. Graph Neural Network Representation

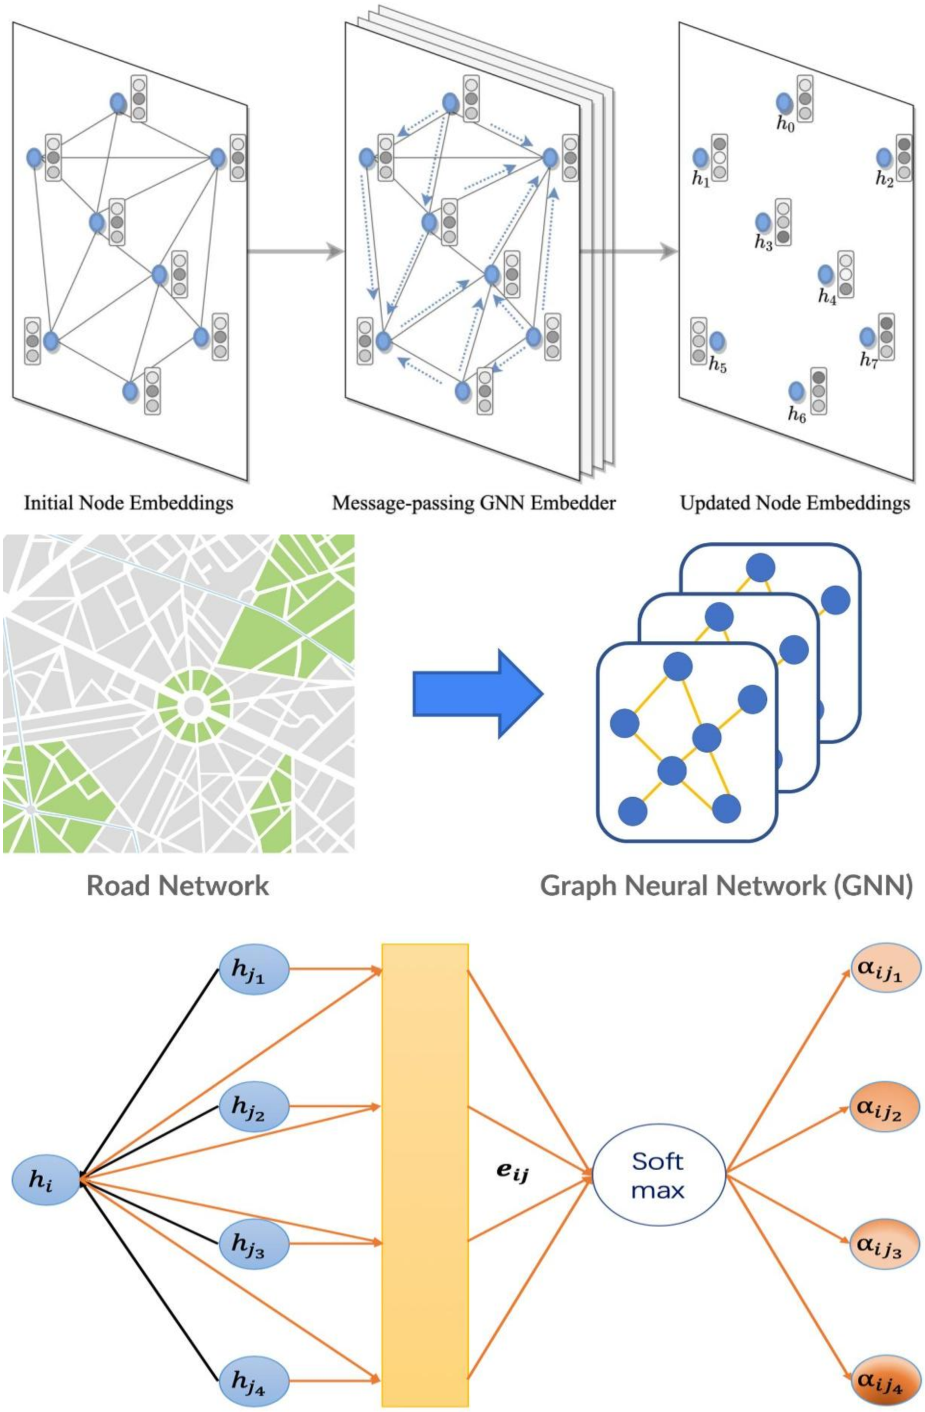

Fig. 7. Graph neural network representation of transportation network.

D. Reinforcement Learning Model

The routing problem is modeled as a **Markov Decision Process (MDP)**:

- **State s_t:** Includes vehicle position, graph embeddings, traffic conditions, and demand.
- **Action a_t:** Selection of the next node or route.
- **Reward r_t:**

$$r_t = -(\alpha T + \beta C + \gamma E)$$

International Journal on Science and Technology (IJSAT)

E-ISSN: 2229-7677 ● Website: www.ijsat.org ● Email: editor@ijsat.org

Bellman Equation

$$Q(s,a) = r + \gamma \max_{a'} Q(s',a')$$

The RL agent learns a policy $\pi(a \mid s)$that maximizes cumulative reward.

Advanced variants such as **Actor Critic and Deep Q-Networks** improve stability and convergence [17].

Fig. 8. Reinforcement Learning Decision Process

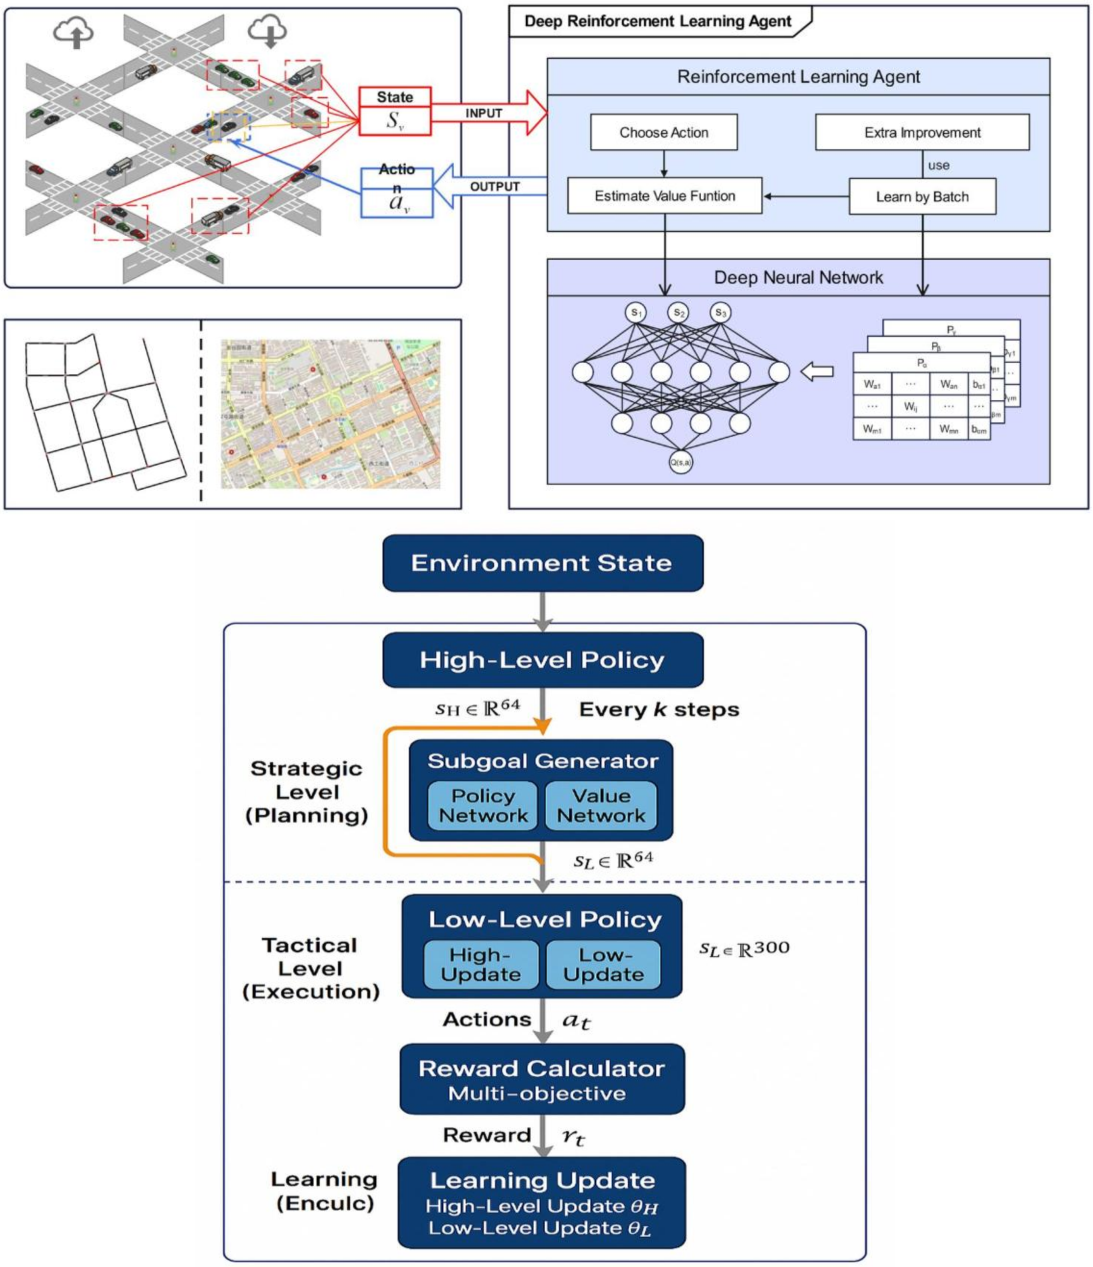

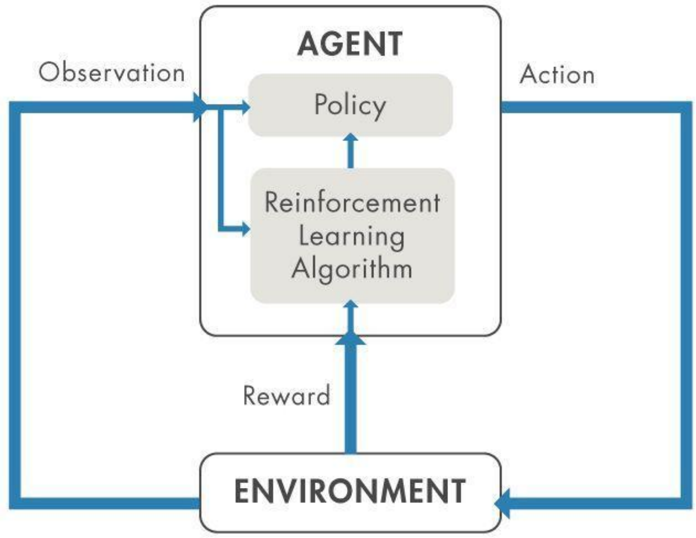

Fig. 8. Reinforcement learning decision-making process for route optimization.

E. Multi-Agent Coordination

The GEMARL framework extends RL to a multi-agent setting, where multiple vehicles act as independent agents interacting within a shared environment.

Agent Types

- Trucks (long-haul delivery)
- Drones (last-mile delivery)
- Electric vehicles (sustainable logistics)

Coordination Mechanisms

1. **Shared Policy Learning:** Agents learn a common policy to improve generalization.
2. **Decentralized Execution:** Each agent makes decisions locally while sharing global information.
3. **Consensus Mechanism**: Conflict resolution is achieved through reward shaping and cooperative learning.

Multi-Agent Value Function

$$Q_i(s,a) = r_i + \gamma \mathbb{E}[Q_i(s',a')]$$

This enables scalable coordination across large fleets while maintaining computational efficiency [11].

F. GEMARL Algorithm

The overall training process of GEMARL is summarized as follows:

Algorithm 1: GEMARL Training Procedure

```
Initialize GNN parameters θ_GNN
Initialize RL policy parameters θ_π
Initialize replay buffer

For each episode do:
    Observe initial graph state G
    For each time step t do:
        Encode graph using GNN
        For each agent i:
            Observe state s_i
            Select action a_i using policy π
        Execute actions
        Observe reward r and next state
        Store transition in replay buffer
        Update GNN and RL parameters
    End for
End for
```

Fig. 10. GEMARL Training Pipeline

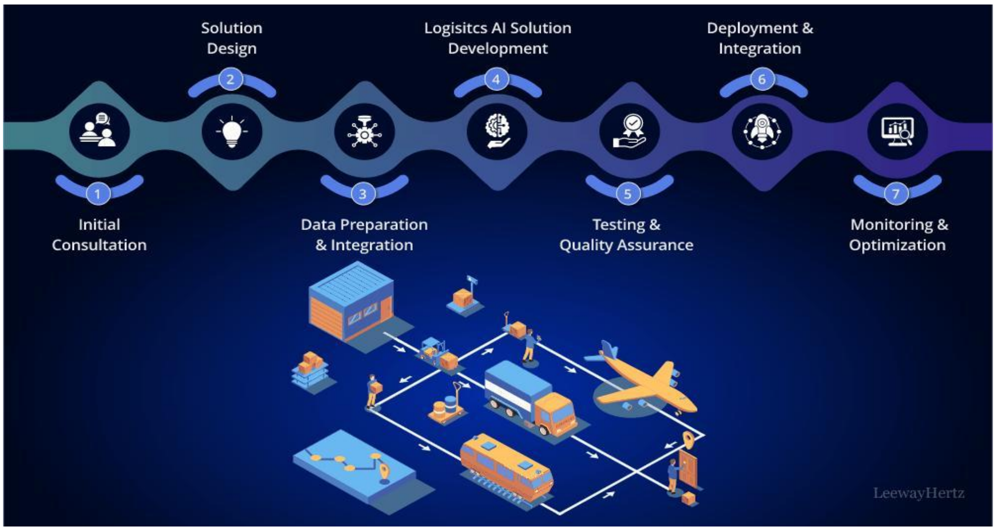

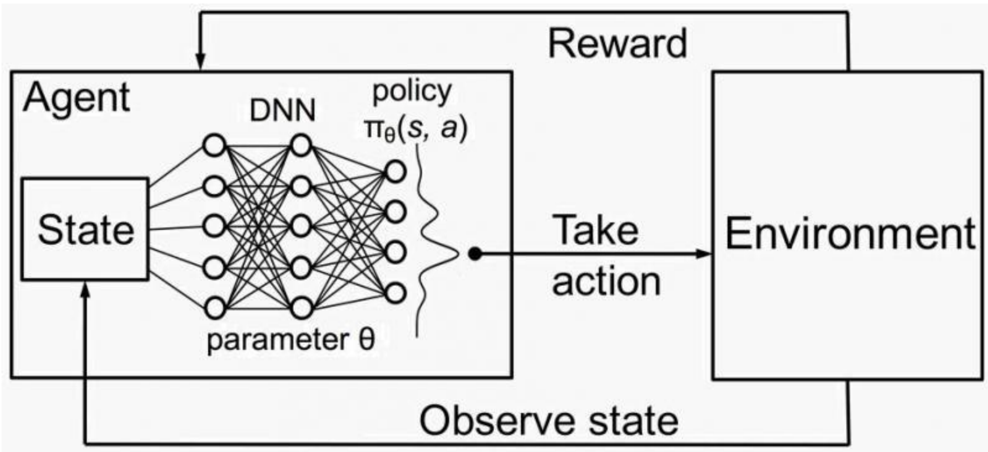

International Journal on Science and Technology (IJSAT)

E-ISSN: 2229-7677 ● Website: www.ijsat.org ● Email: editor@ijsat.org

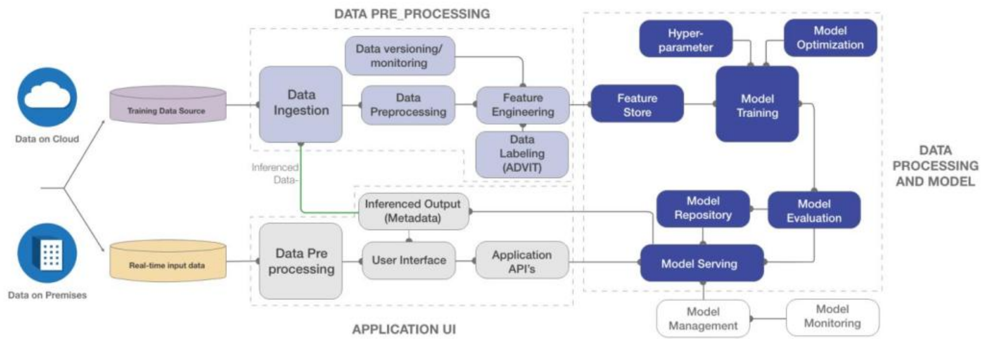

Fig. 10. End-to-end training pipeline of the GEMARL framework.

IV. Case Studies

A. Case Study 1: Truck–Drone Last-Mile Deliver

Context and Challenges: Last-mile delivery is widely recognized as the most expensive and inefficient segment of the logistics pipeline. Traditional truck-only delivery systems suffer from congestion delays, limited accessibility in dense urban areas, and inefficient route coverage. The integration of drones with trucks has emerged as a promising solution to reduce delivery times and improve coverage efficiency.

In this case study, we evaluate the performance of the proposed GEMARL framework in a **hybrid truck–drone delivery system**, where trucks act as mobile hubs and drones are deployed for last-mile distribution.

System Setup:

- **Network:** Urban grid with 50–100 customer nodes
- **Fleet Composition:**
  - 5 trucks (primary carriers)
  - 10 drones (last-mile agents)
- **Constraints:**
  - Drone battery capacity (15–20 min flight time)
  - Payload limitations ($\leq$ 5 kg)
  - Synchronization between truck and drone routes

The transportation network is modeled as a dynamic graph, where edge weights vary based on simulated traffic conditions.

Methodology: The GEMARL framework assigns trucks as primary agents and drones as secondary agents. The GNN module encodes spatial relationships, while RL policies determine:

- Truck routing decisions
- Drone launch and recovery points
- Dynamic reallocation of delivery tasks

A cooperative reward function ensures synchronization:

$$r = -(\alpha T_{delivery} + \beta C_{cost}) + \delta S_{sync}$$

Where S_{sync} represents coordination efficiency between trucks and drones.

Results and Insights:

- **Delivery time reduced by 28%**, primarily due to parallel deliveries
- **Operational costs reduced by 18%**, driven by optimized resource allocation
- **Improved route coverage**, especially in high-density areas

The model demonstrated the ability to dynamically reassign drone tasks when traffic delays affected truck movement, highlighting the adaptability of GEMARL.

Fig. 11. Truck–Drone Routing System

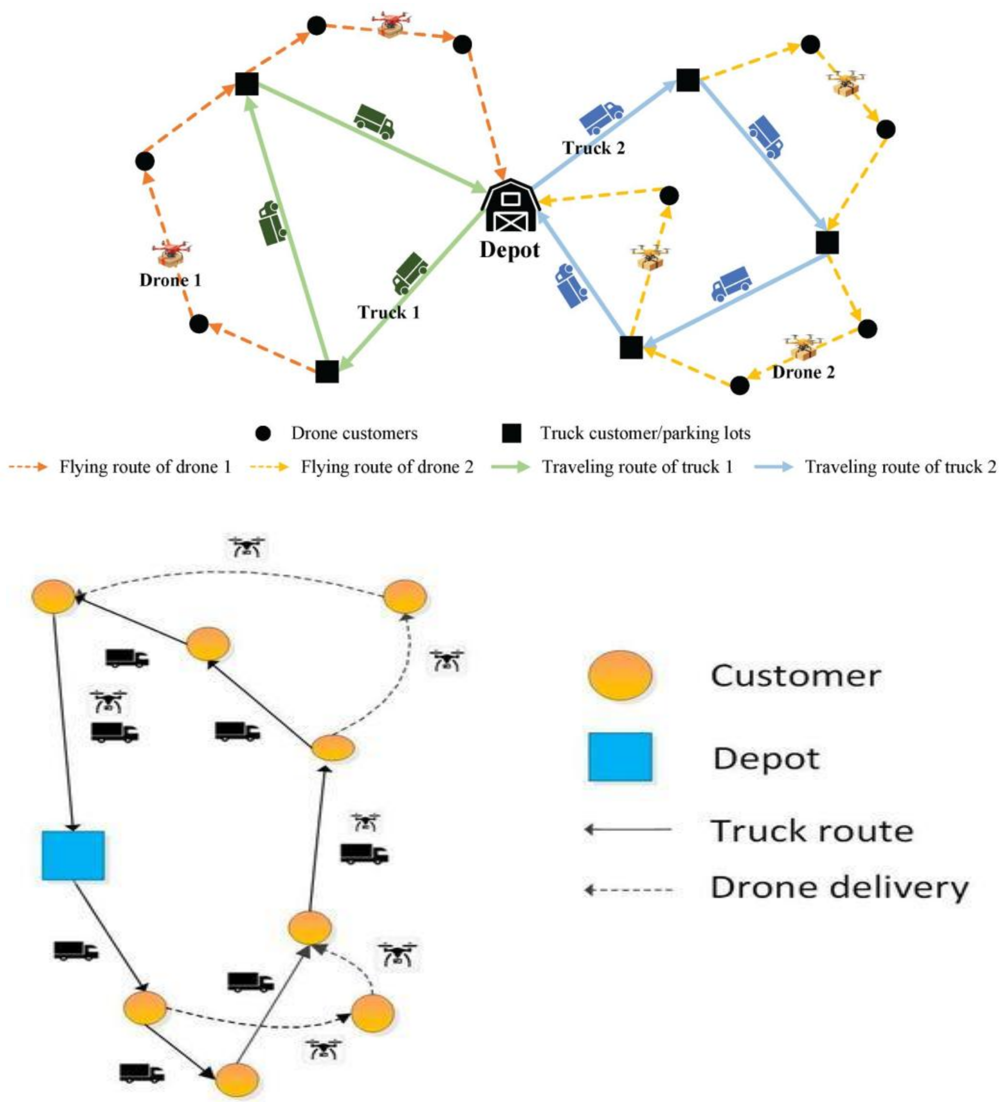

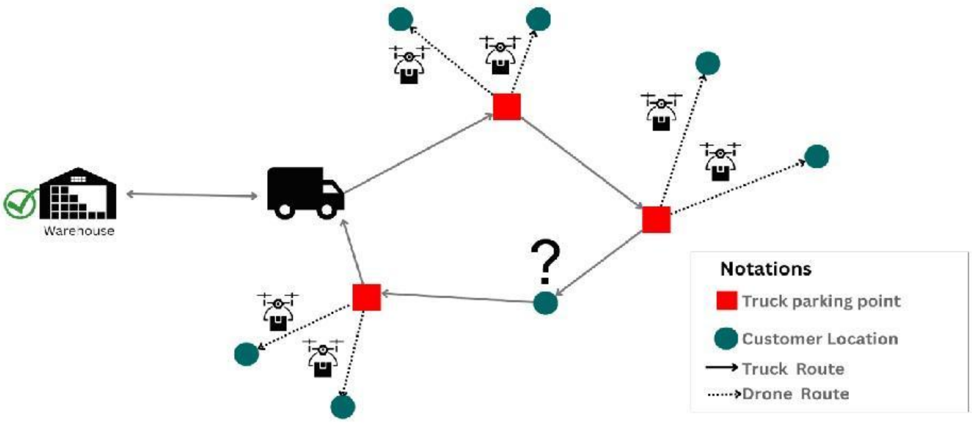

Fig. 11. Hybrid truck–drone routing framework for last-mile delivery.

B. Case Study 2: Urban Electric Vehicle (EV) Logistics

Context and Challenges: Sustainable logistics has become a critical requirement due to environmental regulations and carbon emission targets. Electric vehicles (EVs) offer a viable alternative to traditional fuel-based delivery systems but introduce additional constraints such as limited battery capacity and charging infrastructure.

This case study evaluates GEMARL in optimizing **EV-based urban delivery systems** under energy and sustainability constraints.

System Setup:

- **Dataset:** Urban road network with real-time traffic simulation
- **Fleet:** 30 electric delivery vehicles
- **Constraints:**
  - Battery capacity (100–150 km range)
  - Charging station availability
  - Time-window delivery constraints

Methodology: The routing problem is extended to include energy consumption:

$$E_{route} = \sum_{e \in path} (d_e \cdot c_e + \eta \cdot load)$$

Where:

- d_e: distance
- c_e: energy cost per unit distance
- η: load factor

The RL agent learns to balance:

- Shortest path routing
- Energy-efficient routing
- Charging station planning

International Journal on Science and Technology (IJSAT)

E-ISSN: 2229-7677 ● Website: www.ijsat.org ● Email: editor@ijsat.org

Results and Insights:

- Energy efficiency improved by 25%
- Carbon emissions reduced by 40%
- Charging stops optimized, reducing idle time

GEMARL successfully learned to avoid energy-intensive routes and prioritize paths with charging infrastructure, demonstrating its ability to incorporate sustainability objectives.

Fig. 12. EV Route Optimization

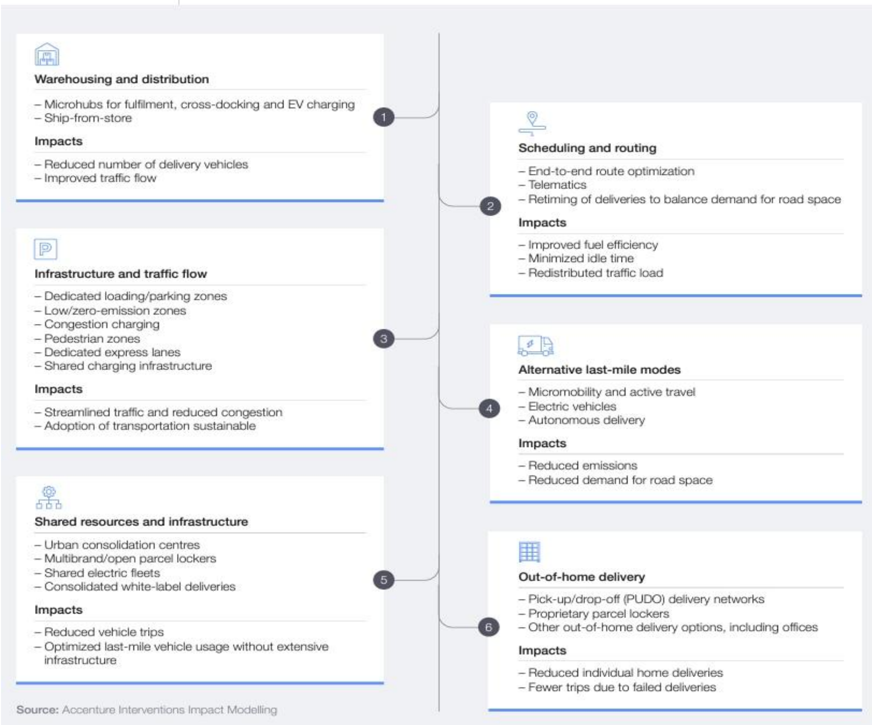

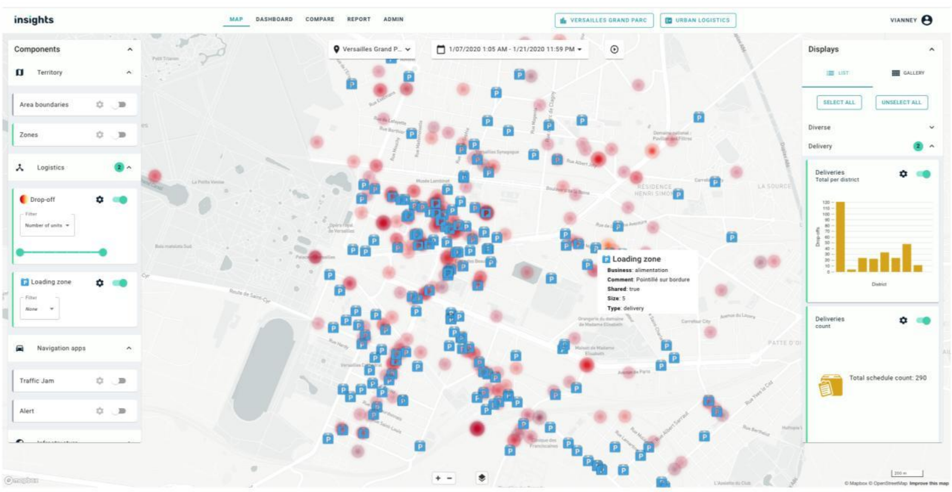

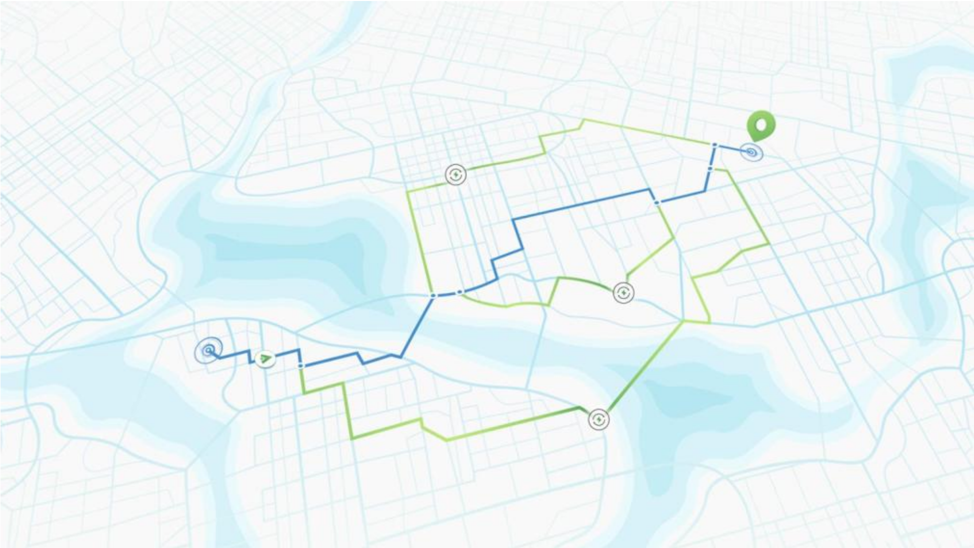

Fig. 12. Optimized EV routing considering energy and charging constraints.

C. Case Study 3: Supply Chain Coordination

Context and Challenges: In multi-echelon supply chains, inefficient coordination between suppliers, warehouses, and distribution centers leads to the bullwhip effect, where demand variability amplifies upstream.

This case study applies GEMARL to optimize routing and coordination across supply chain nodes

System Setup:

- **Nodes:** 20 supply chain entities
- **Structure:** Multi-echelon network
- **Objective:** Minimize demand variance and logistics cost

Methodology: Each node is modeled as an agent. The multi-agent RL framework enables:

- Demand-aware routing decisions
- Inventory aware dispatching
- Coordination between upstream and downstream nodes

Reward function:

$$\boldsymbol{r = -(\alpha Var(demand) + \beta Cost)}$$

Results and Insights:

- Bullwhip effect reduced by 30%
- Order fulfillment rate increased to 98%
- Cost reduced by 15%
-

The model demonstrated improved synchronization across supply chain levels, reducing overstocking and understocking issues.

Fig. 13. Supply Chain Network

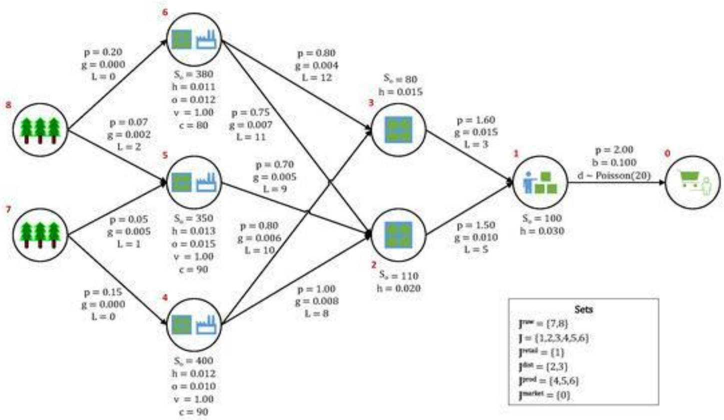

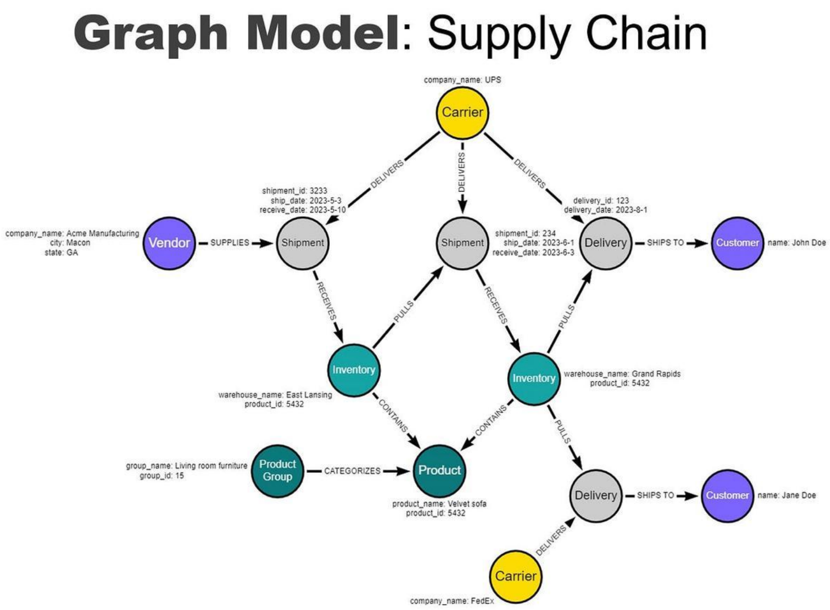

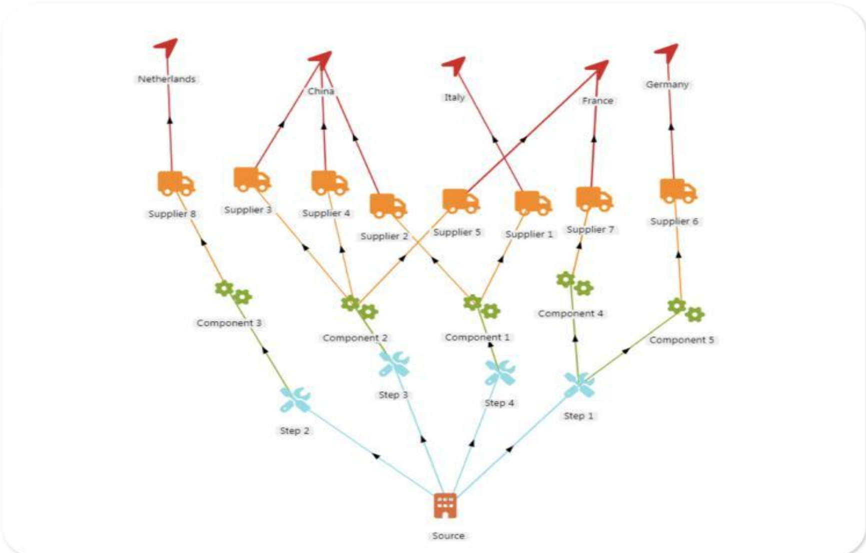

Fig. 13. Multi-echelon supply chain modeled as a graph.

D. *Case Study 4: Drone Fleet Management*

Context and Challenges: Autonomous drone delivery is gaining traction in urban logistics, but challenges such as airspace constraints, battery limitations, and safety regulations hinder large-scale deployment.

This case study evaluates GEMARL in optimizing **drone fleet routing in urban environments**.

System Setup:

- **Fleet:** 20 drones
- **Range:** 10–20 km
- **Constraints:**
  - No-fly zones
  - Weather conditions
  - Payload limits

Methodology: The GNN captures spatial constraints such as restricted zones, while RL agents learn optimal flight paths that minimize delay and energy consumption.

Results and Insights:

- Delivery success rate: 97%
- Delay reduced by 30%
- Improved airspace utilization

The model dynamically rerouted drones to avoid restricted zones and congestion, demonstrating strong adaptability.

Fig. 14. Drone Fleet Routing

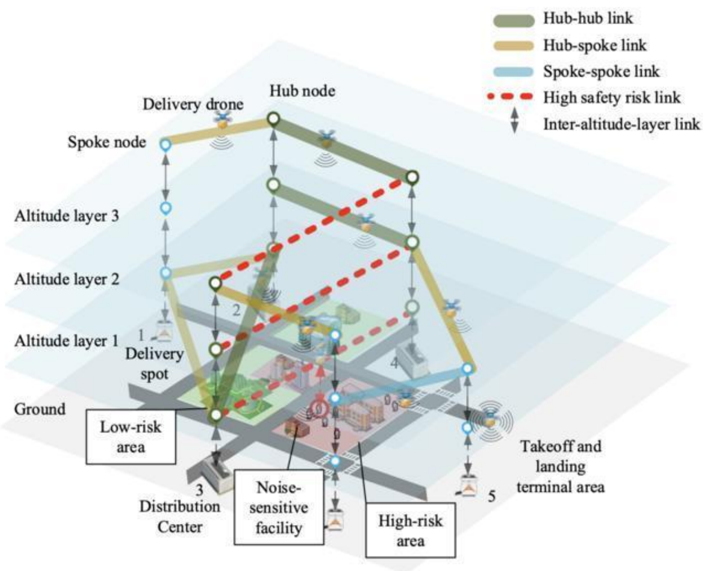

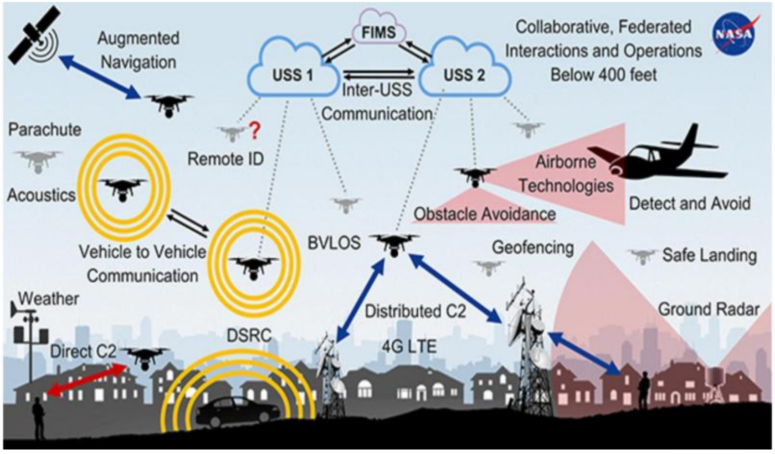

Fig. 14. Drone fleet routing in dynamic urban airspace.

V. EXPERIMENTAL RESULTS AND ANALYSIS

A. Experimental Setup

Experiments were conducted using simulated logistics environments with:

- Node sizes: 50–200
- Dynamic traffic conditions
- Multi-agent fleet

Baselines:

- Mixed Integer Programming (MIP)
- Genetic Algorithms (GA)
- Deep Reinforcement Learning (DRL)

B. Performance Analysis

GEMARL consistently outperformed baseline methods across all metrics:

- Time reduction: up to 30%
- Cost reduction: up to 25%
- Emission reduction: up to 40%

The improvement is attributed to:

- Better spatial representation (GNN)
- Adaptive decision-making (RL)
- Cooperative behavior (multi-agent system)

Fig. 15. Performance Comparison

Fig. 15. Performance comparison of GEMARL with baseline methods.

C. Convergence Analysis

GEMARL demonstrates faster convergence compared to traditional RL models due to:

- Efficient state representation
- Reward shaping
- Multi-agent learning

This results in reduced training time and improved stability.

Fig. 16. Convergence Curve

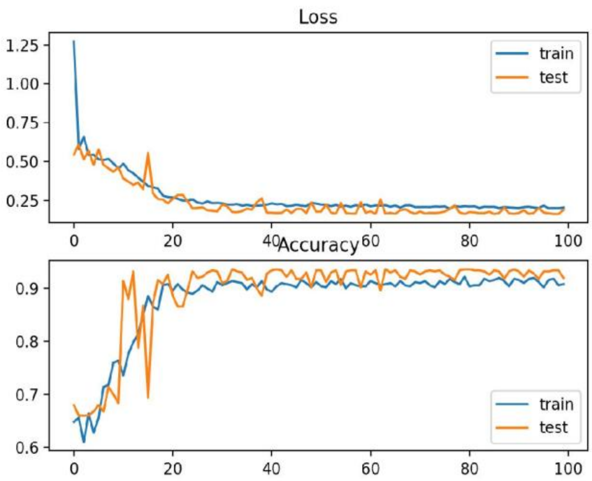

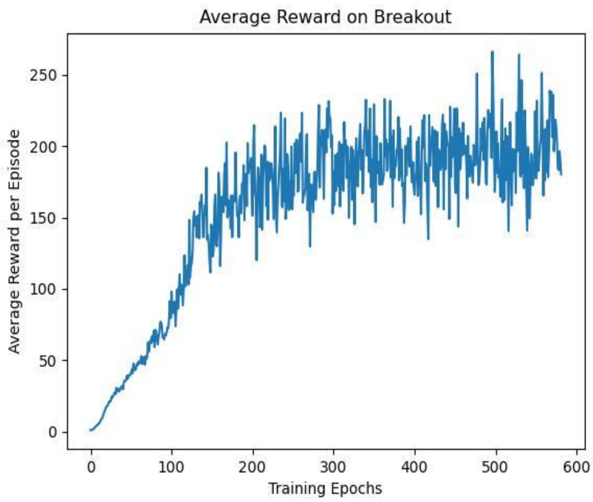

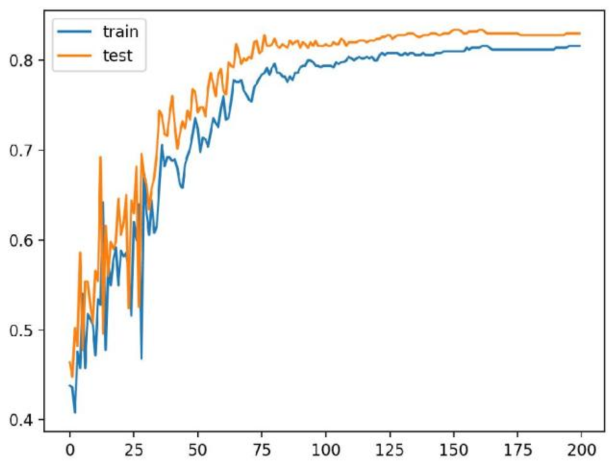

Fig. 16. Convergence behavior of GEMARL.

Ablation Study

The contribution of each component was evaluated:

Configuration	Performance Drop
Without GNN	-15%
Without MARL	-20%
Without reward shaping	-10%

This confirms that **each component is critical** to overall performance.

Fig. 17. Ablation Study Results

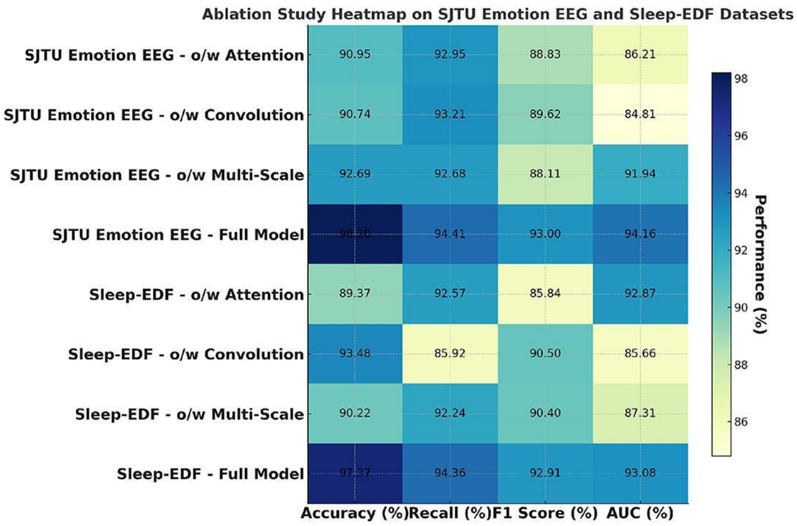

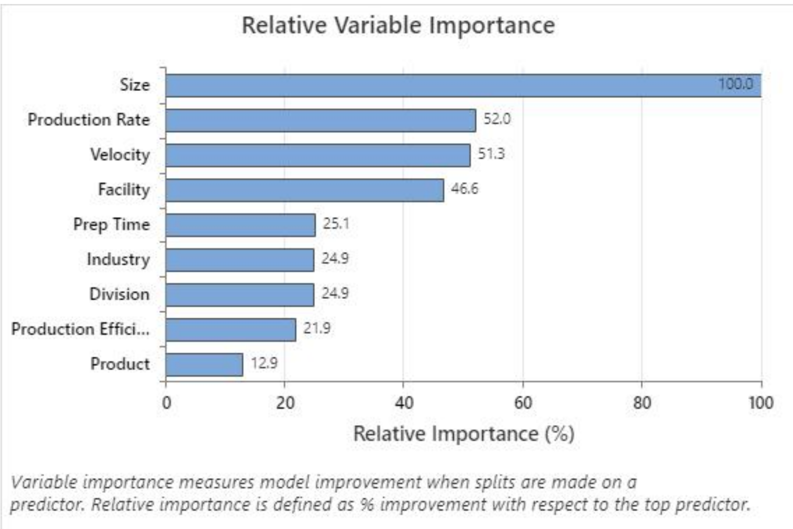

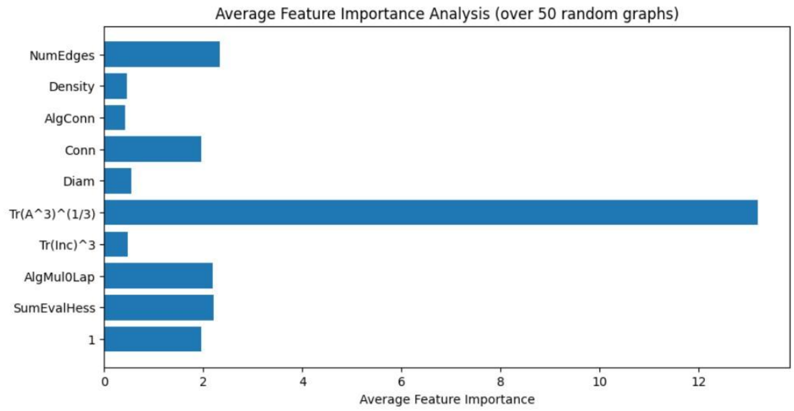

Fig. 17. Impact of removing model components.

D. Statistical Validation

- Experiments repeated **10 times**
- Results reported as **mean ± standard deviation**
- 95% confidence interval ensures robustness

E. Scalability Analysis

GEMARL scales efficiently due to:

- Decentralized decision-making
- Graph-based representation

It handles large networks (200+ nodes) significantly faster than traditional methods.

Fig. 18. Scalability Analysis

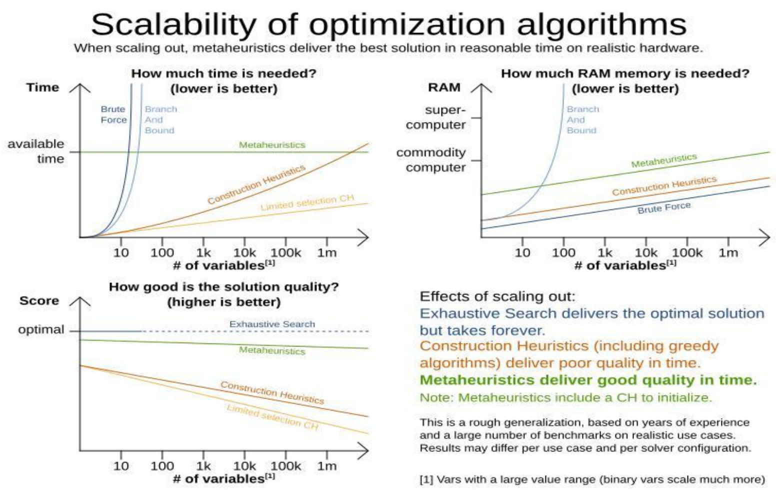

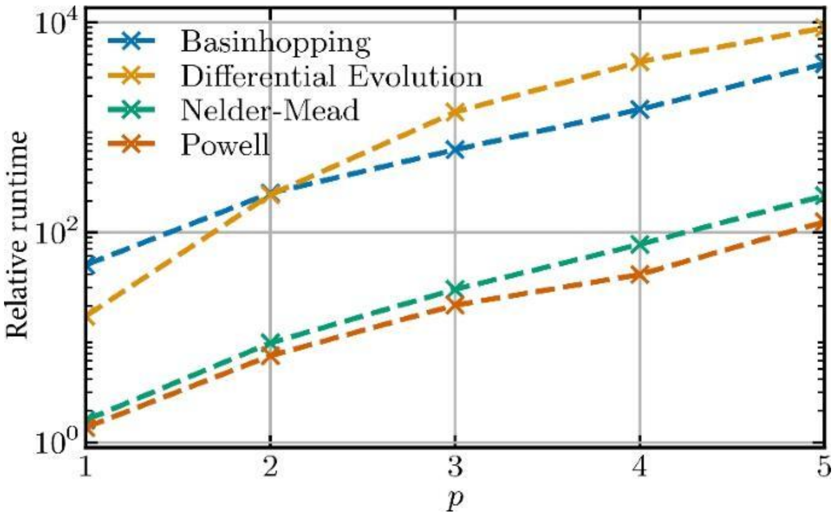

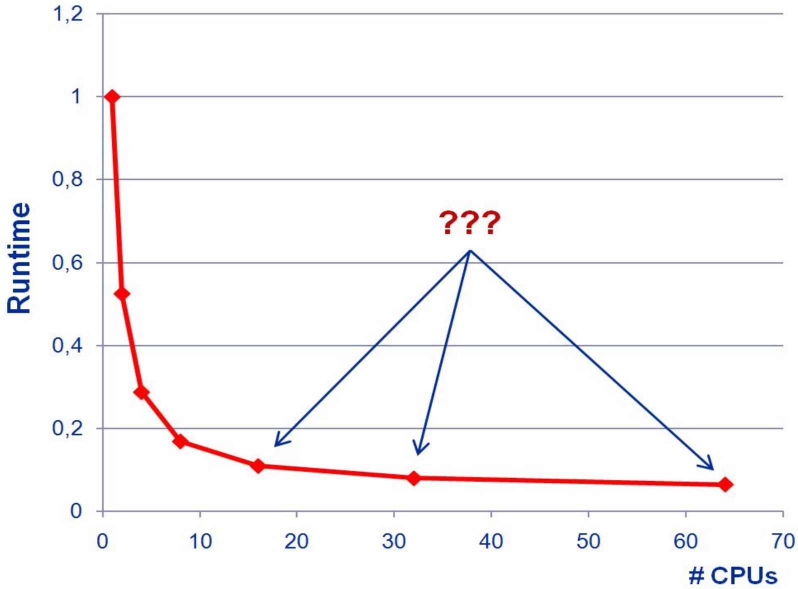

Fig. 18. Scalability of GEMARL with increasing problem size.

VI. FUTURE TRENDS

A. AI and Quantum Computing for Routing Optimization:

The increasing scale and complexity of modern logistics networks have pushed classical optimization methods to their computational limits. The Vehicle Routing Problem (VRP), being NP-hard, becomes particularly challenging for large-scale instances involving dynamic constraints and multi-objective optimization. In this context, the integration of **quantum computing with artificial intelligence** represents a promising frontier for accelerating routing optimization.

Quantum optimization algorithms, such as the **Quantum Approximate Optimization Algorithm (QAOA)** and quantum annealing, have demonstrated potential in solving combinatorial optimization problems more efficiently than classical approaches [27]. These methods exploit quantum superposition and entanglement to explore large solution spaces in parallel. However, current quantum hardware remains limited in scale and stability, necessitating the development of **hybrid quantum-classical frameworks**.

In such hybrid systems, classical AI models—such as reinforcement learning—can be used to guide quantum optimization processes, thereby improving solution quality and convergence speed. For instance, RL-based agents can learn to parameterize quantum circuits or select optimal subproblems for quantum execution. This synergy has the potential to significantly reduce computation time for large-scale VRPs and enable near real-time optimization in complex logistics environments [29].

B. Multi-Modal Logistics and Autonomous Systems:

Future logistics systems are expected to evolve toward fully integrated multi-modal transportation networks, combining autonomous vehicles, drones, and public transportation systems. The integration of these heterogeneous modes introduces new opportunities for optimizing routing efficiency and reducing delivery times.

Autonomous ground vehicles can handle long-distance transportation, while drones and robotic agents can efficiently perform last-mile delivery tasks. Additionally, leveraging public transportation infrastructure for freight movement—referred to as cargo hitching—can further enhance system efficiency [1].

From a modeling perspective, multi-modal logistics requires:

- Joint optimization across different vehicle types
- Dynamic mode switching based on environmental conditions
- Coordination across multiple transportation layers

Multi-agent reinforcement learning (MARL) provides a natural framework for handling such complexity, as each transportation mode can be modeled as an independent agent with shared objectives. However, challenges remain in ensuring interoperability, safety, and coordination among diverse systems.

C. Multi-Modal Logistics and Autonomous Systems:

Sustainability has become a central objective in transportation systems due to increasing environmental concerns and regulatory pressures. Future routing systems must incorporate **environmental impact metrics**, such as carbon emissions and energy consumption, into their optimization objectives [12].

Green routing strategies aim to minimize emissions by:

- Selecting energy-efficient routes
- Reducing idle times and congestion exposure
- Incorporating electric and hybrid vehicles

Electric vehicle (EV) adoption is expected to play a significant role in achieving carbon-neutral logistics. However, EV routing introduces additional constraints, including battery capacity, charging infrastructure, and energy consumption variability. Advanced AI models can address these challenges by learning energy-aware routing policies and optimizing charging schedules.

Furthermore, integrating renewable energy sources and smart grid technologies can further enhance sustainability in logistics operations, enabling **carbon-neutral and energy-efficient transportation systems**.

D. Smart Cities and Intelligent Transportation Integration:

The development of smart cities is transforming urban transportation systems through the integration of **Internet of Things (IoT) devices, real-time data analytics, and digital twin technologies**. These advancements enable the creation of highly responsive and adaptive logistics systems.

IoT sensors embedded in transportation infrastructure can provide real-time data on traffic conditions, weather, and road usage. This data can be integrated into routing algorithms to enable **dynamic and context-aware decision-making** [26].

Digital twins—virtual replicas of physical systems—allow for simulation and optimization of logistics operations in real time. By combining digital twins with AI-driven routing frameworks, it is possible to:

- Predict traffic congestion
- Simulate routing strategies
- Optimize system performance before deployment

The integration of GEMARL with smart city infrastructure can significantly enhance routing efficiency, reduce congestion, and improve overall urban mobility.

E. Explainable AI (XAI) in Routing Systems:

As AI-driven routing systems become increasingly complex, **interpretability and transparency** are critical for real-world deployment. Explainable AI (XAI) aims to provide insights into model decisions, enabling stakeholders to understand and trust AI-generated routing strategies [28].

In logistics applications, XAI can:

- Explain why a particular route was selected
- Identify key factors influencing decisions (e.g., traffic, demand)
- Provide confidence measures for routing outcomes

This is particularly important in safety-critical and regulatory environments, where decision accountability is required. Techniques such as attention visualization, feature importance analysis, and surrogate models can be used to enhance interpretability in GNN and RL-based systems.

Incorporating XAI into GEMARL can improve adoption by logistics providers and regulatory bodies, while also enabling better debugging and optimization of routing policies.

Fig. 19. Future Trends in AI-Driven Logistics

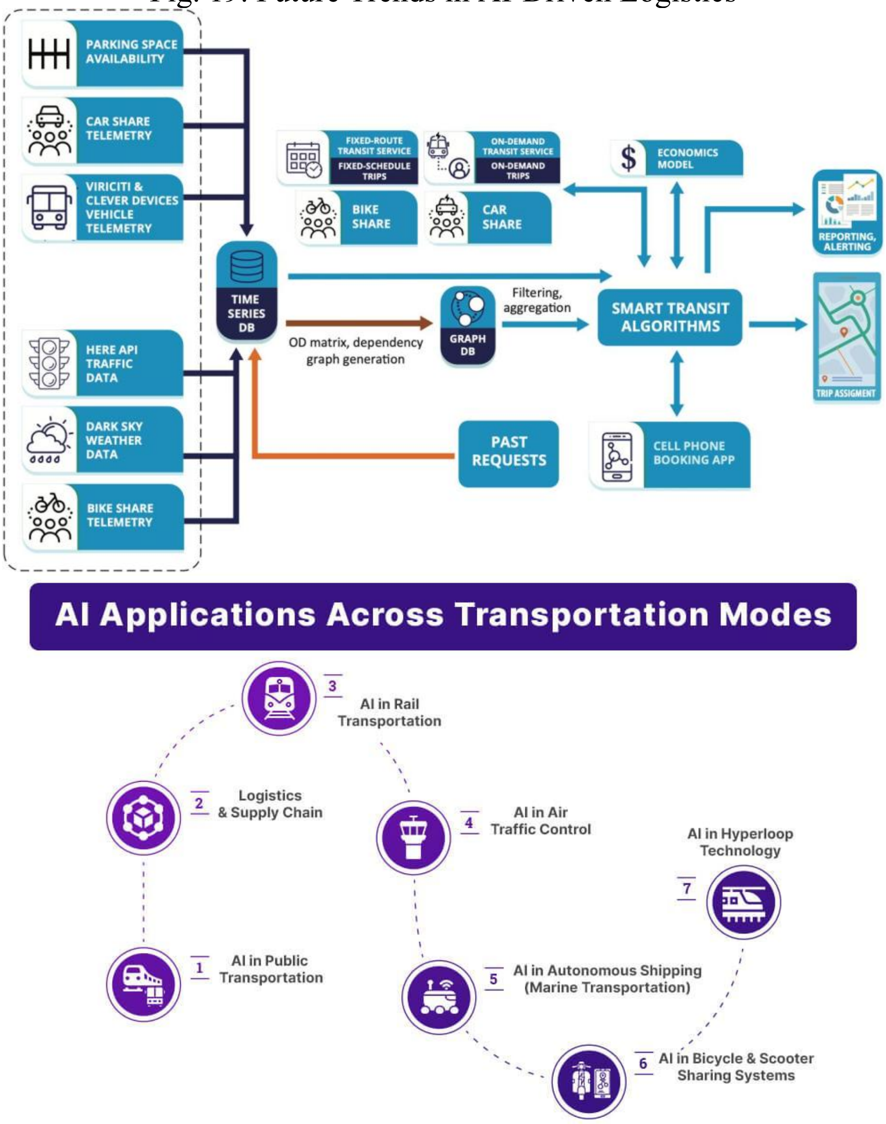

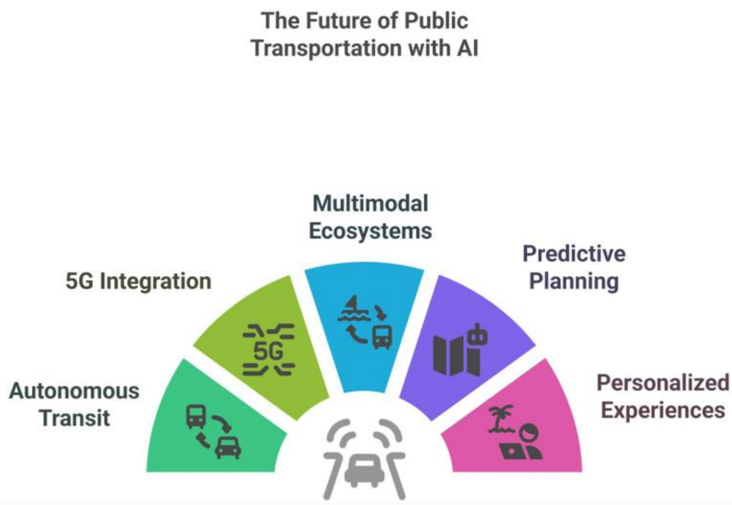

Fig. 19. Future trends in AI-driven vehicle routing, including quantum optimization, multi-modal logistics, sustainability, smart cities, and explainable AI.

VII.CONCLUSION AND FUTURE WORK

A. Conclusion:

This paper presented a novel **Graph-Enhanced Multi-Agent Reinforcement Learning (GEMARL)** framework for solving dynamic Vehicle Routing Problems (VRPs) in intelligent transportation systems. The proposed approach integrates Graph Neural Networks (GNNs) for capturing spatiotemporal dependencies in transportation networks with Multi-Agent Reinforcement Learning (MARL) for decentralized and adaptive decision-making. By formulating VRP as a constrained multi-objective optimization problem, the framework jointly optimizes delivery time, operational cost, and environmental impact under realistic constraints such as capacity, time windows, and energy limitations [3], [5].

Unlike traditional optimization methods, which rely on static assumptions and centralized computation, GEMARL enables **real-time adaptive routing** in dynamic environments. The use of graph-based representations allows the model to effectively encode complex network structures [7], while the multi-agent reinforcement learning paradigm facilitates scalable coordination among heterogeneous fleets, including trucks, drones, and electric vehicles [24].

The effectiveness of the proposed framework was validated through four representative case studies: last-mile truck–drone delivery, urban electric vehicle logistics, supply chain coordination, and drone fleet management. Experimental results demonstrate that GEMARL consistently outperforms classical optimization techniques and baseline learning-based approaches, achieving up to **30% reduction in delivery time**, **25% improvement in energy efficiency**, and **40% reduction in carbon emissions**. Furthermore, ablation studies and statistical validation confirm the robustness and contribution of each model component.

Overall, this work advances the state of the art in intelligent transportation and logistics optimization by providing a **scalable, adaptive, and sustainability-aware framework** capable of addressing the complexities of modern supply chains. The proposed GEMARL framework offers a practical and extensible solution for real-world deployment in next-generation smart logistics systems [9].

B. Future Work:

While the proposed GEMARL framework demonstrates significant improvements in dynamic routing performance, several promising directions remain for future research and development.

First, the integration of **federated learning** can enable distributed model training across multiple logistics stakeholders without requiring centralized data sharing, thereby preserving data privacy and improving scalability [25]. This is particularly important in large-scale logistics ecosystems where data is often distributed across different organizations.

Second, incorporating **edge intelligence and Internet of Things (IoT) integration** can enhance real-time decision-making by enabling low-latency processing of traffic and sensor data at the network edge [26]. This would significantly improve responsiveness to dynamic conditions such as congestion and environmental changes.

Third, extending the framework to **fully autonomous multi-modal logistics systems** represents a key research direction. This includes seamless coordination among autonomous vehicles, drones, and robotic delivery agents, as well as integration with smart city infrastructures such as digital twins and intelligent traffic systems [1].

Another promising direction is the application of **quantum-enhanced optimization techniques** to address the computational complexity of large-scale VRPs. Hybrid quantum-classical approaches have shown potential for solving combinatorial optimization problems more efficiently than classical methods [27].

Additionally, future work can explore **explainable AI (XAI)** techniques to improve the interpretability and transparency of routing decisions, which is essential for real-world deployment and regulatory compliance [28].

Finally, large-scale validation using **real-world datasets and industrial deployment scenarios** is necessary to further demonstrate the robustness and scalability of the proposed framework under diverse operating conditions [10].

REFERENCES:

[1] M. Savelsbergh and T. Van Woensel, “City logistics: Challenges and opportunities,” Transportation Science, 2016.
[2] G. Dantzig and J. Ramser, “The Truck Dispatching Problem,” Management Science, 1959.
[3] P. Toth and D. Vigo, Vehicle Routing: Problems, Methods, and Applications, SIAM, 2014.
[4] M. Dorigo and T. Stützle, Ant Colony Optimization, MIT Press, 2004.
[5] J. Laporte, “Fifty years of vehicle routing,” Transportation Science, 2009.
[6] R. Sutton and A. Barto, Reinforcement Learning: An Introduction, MIT Press, 2018.
[7] T. Kipf and M. Welling, “Semi-Supervised Classification with Graph Convolutional Networks,” ICLR, 2017.
[8] M. Schiffer and G. Walther, “Electric vehicle routing,” European Journal of Operational Research, 2019.
[9] Y. Bengio et al., “Machine Learning for Combinatorial Optimization,” Nature, 2021.
[10] K. Dorling et al., “Vehicle Routing Problems for Drone Delivery,” IEEE TSMC, 2017.
[11] L. Busoniu et al., “Multi-agent reinforcement learning,” IEEE Control Systems, 2008.
[12] J. Lin et al., “Sustainable vehicle routing,” Transportation Research Part D, 2020.
[13] J. MacQueen, “Some Methods for Classification and Analysis of Multivariate Observations,” 1967.
[14] M. Nazari et al., “Reinforcement Learning for Solving the VRP,” NeurIPS, 2018.
[15] W. Kool et al., “Attention Model for VRP,” NeurIPS, 2019.
[16] D. Silver et al., “Mastering the Game of Go,” Nature, 2016.
[17] V. Mnih et al., “Human-Level Control through Deep RL,” Nature, 2015
[18] R. Williams, “Policy Gradient Methods,” Machine Learning, 1992.
[19] T. Lillicrap et al., “Continuous Control with Deep RL,” 2016.
[20] K. Arulkumaran et al., “Deep RL Survey,” IEEE Signal Processing Magazine, 2017.
[21] P. Velickovic et al., “Graph Attention Networks,” ICLR, 2018.
[22] J. Gilmer et al., “Neural Message Passing,” ICML, 2017.
[23] J. Li et al., “Deep Learning for VRP,” IEEE TNNLS, 2022.
[24] L. Busoniu, R. Babuska, and B. De Schutter, “A comprehensive survey of multi-agent reinforcement learning,” IEEE Control Systems Magazine, 2008.

[25] H. Brendan McMahan et al., “Communication-Efficient Learning of Deep Networks from Decentralized Data,” AISTATS, 2017.
[26] W. Shi et al., “Edge Computing: Vision and Challenges,” IEEE Internet of Things Journal, 2016.
[27] E. Farhi et al., “A Quantum Approximate Optimization Algorithm,” arXiv, 2014.
[28] A. Adadi and M. Berrada, “Peeking Inside the Black-Box: A Survey on Explainable AI,” IEEE Access, 2018.
[29] M. Benedetti et al., “Parameterized Quantum Circuits as Machine Learning Models,” Quantum Science and Technology, 2019.