

Smarter Vehicle Networks: Cognitive AI for Next-Gen Cars □ □

Sachin Shelke., Dr. Satish Sankaye., Kanchan Vaishnav

Computer Science, MGM University

DOI: <https://doi.org/10.51244/IJRSI.2025.120700239>

Received: 20 July 2025; Accepted: 27 July 2025; Published: 26 August 2025

ABSTRACT

Tomorrow's cars need smarter networks. Today's Vehicular Ad-hoc Networks (VANETs) struggle to meet the diverse and often conflicting demands of advanced applications, from split-second collision alerts to seamless streaming. Current cognitive radio (CR)-VANETs fall short, lacking the ability to orchestrate network resources effectively for multiple Quality of Service (QoS) goals in ever-changing traffic.

This paper introduces the Cognitive-Driven Orchestration Framework (CDOF), a novel approach using Meta-Reinforcement Learning (Meta-RL). CDOF intelligently manages spectrum and communication settings for a wide range of vehicular services. By understanding real-time conditions like vehicle movement, network interference (from "primary users"), and specific application needs, CDOF learns adaptable resource allocation strategies. Our Meta-RL engine, capable of "learning to learn," quickly adjusts to new, unseen situations, ensuring robust network performance even in highly dynamic environments.

CDOF's unique multi-objective reward system prioritizes critical services like safety and autonomous driving (requiring ultra-low latency and high reliability) while efficiently managing resources for less critical services like infotainment. Simulations across various traffic and interference patterns demonstrate that CDOF significantly outperforms existing CR-VANET methods in guaranteeing QoS, adapting rapidly, and minimizing interference.

Keywords: VANETs, Cognitive Radio, Meta-Reinforcement Learning, Multi-Objective QoS, Dynamic Orchestration

INTRODUCTION

Vehicular Ad-hoc Networks (VANETs) are the backbone of future Intelligent Transportation Systems (ITS), powering essential Vehicle-to-Everything (V2X) communication. However, the diverse needs of next-generation applications — from life-saving collision warnings (requiring less than 10ms latency) and autonomous platooning (demanding 99.9% reliability) to high-definition video streaming (needing over 5Mbps throughput) — create a complex web of conflicting Quality of Service (QoS) requirements. These demands clash within resource-limited and highly dynamic networks [1].

While Cognitive Radio (CR) helps overcome spectrum scarcity, it faces two major limitations in current VANET implementations:

- **Insufficient Orchestration:** Existing CR-VANETs typically focus on isolated tasks, such as simply selecting a channel [7], without a holistic approach to dynamic spectrum management, seamless handovers, or effective interference mitigation.
- **Neglected Multi-Objective QoS:** Current solutions lack mechanisms to simultaneously guarantee crucial metrics like latency, reliability, and throughput for various services [3].

We address these critical gaps with CDOF, a framework that offers:

- **Unified Cognitive Orchestration:** CDOF integrates spectrum, power, and handover management into a single, comprehensive system.

- **Meta-RL Decision Engine:** This engine enables policies to adapt rapidly to entirely new and unforeseen environments, such as sudden accidents or novel patterns of primary user (PU) interference.
- **Dynamic QoS Prioritization:** CDOF uses a clever multi-objective reward system to prioritize services based on their real-time importance.

Our key contributions include:

- A novel CDOF architecture for comprehensive CR-VANET orchestration.
- A Meta-RL formulation designed for learning transferable policies.
- A multi-objective reward mechanism that intelligently prioritizes QoS.
- Rigorous validation proving CDOF's superior QoS guarantees and adaptation capabilities.

Related Work

CR-VANETs

Prior work in CR-VANETs often focuses on isolated aspects:

- **Spectrum Sensing:** Techniques like energy detection [9] and cooperative sensing [10] don't inherently integrate QoS considerations.
- **MAC Protocols:** Priority-based channel access schemes [13] rely on static rules, which fail to adapt to dynamic QoS demands.
- **Routing/Resource Allocation:** While systems like SURF [7] optimize channel selection, they often overlook the trade-offs involved in managing QoS for multiple services simultaneously.

QoS Provisioning

Efforts to improve QoS generally fall into two categories:

- **General Improvements:** Adaptive beaconing [17] can boost packet delivery, but it lacks service-specific guarantees.
- **Multi-Objective Schemes:** Heuristic optimization methods, such as Particle Swarm Optimization (PSO) [20], use fixed weights, which hinder real-time adaptability.

RL Limitations

Standard Reinforcement Learning (RL) approaches struggle with poor generalization in new scenarios (e.g., unexpected traffic patterns) and exhibit slow adaptation to abrupt network changes [25]. This makes them unsuitable for the highly dynamic nature of VANETs.

Cdof Architecture

CDOF features a layered, closed-loop cognitive architecture

Data Collection & Contextual Awareness

This layer gathers real-time data from various sources:

- **Sensors:** GPS, On-Board Diagnostics (OBD-II), and Roadside Unit (RSU)-based Primary User (PU) sensing.

- **Output:** This data forms a contextual state vector $SC(t)$, encompassing details like spectrum occupancy, network topology, vehicle density, and Signal-to-Interference-plus-Noise Ratio (SINR).

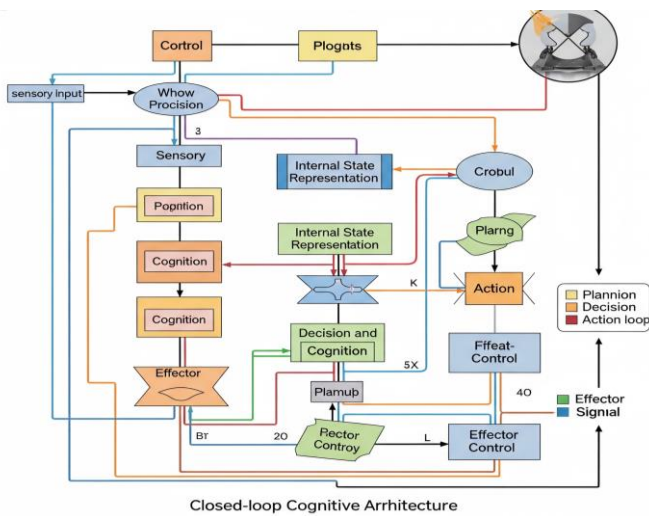


Fig. 1 layered, closed-loop cognitive architecture.

Table No.1 CDOF Architecture Layers and Components

Layer	Components / Functions	Details
3.1 Data Collection & Context Awareness	GPS, OBD-II, RSUs, PU sensing	Builds state $S_{_C}(t)$ with data on spectrum, topology, density, SINR.
3.2 Application & QoS Profiling	Service classification (e.g., Safety, Infotainment)	Defines QoS profiles $Q_{_P}$: latency, reliability, throughput, priority. Supports dynamic adjustment.
3.3 Cognitive Orchestration	Meta-RL engine	Uses $S_{_C}(t)$ and QoS state to generate optimal policy $\pi^*(t)$. Manages resources and handovers.
3.4 Communication & Execution	CR transceivers	Applies policies, adjusts spectrum/power, and updates context via feedback loop.

Application & QoS Profiling

CDOF categorizes services and defines their specific QoS requirements:

- **Service Classes:** Services are prioritized (e.g., Safety = Priority 5, Autonomous Driving = 4, Infotainment = 2).
- **QoS Profiles:** Each service has a defined profile $QP = \{L_{req}, R_{req}, Th_{req}, Priority\}$, specifying required latency (L_{req}), reliability (R_{req}), throughput (Th_{req}), and priority.
- **Dynamic Negotiation:** For instance, the required throughput (Th_{req}) for infotainment can be dynamically adjusted during network congestion.

Cognitive Orchestration Layer

This is the core intelligence of CDOF:

- **Meta-RL Engine:** This engine processes the contextual state $SC(t)$ and QoS state $SQ(t)$ to generate optimal communication policies $\pi^*(t)$.

- **Multi-Objective Reward:** A sophisticated reward function balances rewards for meeting latency, reliability, and throughput goals, while penalizing interference with primary users.
- **Resource Allocation:** CDOF jointly optimizes crucial parameters like channel selection, transmit power, Modulation and Coding Scheme (MCS), and bandwidth.
- **Handover Management:** It ensures seamless transitions between Vehicle-to-Vehicle (V2V) and Vehicle-to-Infrastructure (V2I) communications, maintaining QoS throughout.

Communication & Execution

- **CR Transceivers:** These hardware components execute the spectrum and power adjustments dictated by the orchestration layer.
- **Feedback Loop:** The system continuously monitors actual QoS performance, feeding this information back to update the contextual state $SC(t)$ and ensure adaptive learning.

Meta-RL for Multi-Objective QoS

Problem Formulation

Each specific communication scenario or "task" T_i within the VANET is modeled as a Partially Observable Markov Decision Process (POMDP): $(S_i, A_i, P_i, R_i, \gamma)$. Here, $A_i = \{\text{channel, power, MCS}\}$ represents the available actions. The overarching goal is to learn a meta-policy π_{meta} that allows the system to rapidly adapt to new tasks, such as unforeseen patterns of primary user activity.

MAML-Based Meta-RL

CDOF leverages Model-Agnostic Meta-Learning (MAML) [27] for its Meta-RL engine. MAML involves two loops:

- **Inner Loop (Task-Specific Adaptation):** For a given task T_i , the model quickly adapts its parameters θ to find an optimal task-specific policy $\theta_i' = \theta - \alpha \nabla_{\theta} L_{T_i}(\pi_{\theta})$. This involves taking a few gradient steps on the loss function L_{T_i} for that particular task.
- **Outer Loop (Meta-Update Across Tasks):** The meta-learner updates its initial parameters θ based on the performance of the adapted policies across various tasks: $\theta \leftarrow \theta - \beta \nabla_{\theta} \sum T_i L_{T_i}(\pi_{\theta_i'})$. This teaches the model how to "learn to learn" quickly.

Reward Function

The reward function $R(s,a)$ is designed to balance multiple QoS objectives:

$$R(s,a) = \sum_{j \in \text{Apps}} (W_{Lj} \cdot f(L_j) + W_{Rj} \cdot \text{PDR}_j + W_{Thj} \cdot \text{Th}_j / \text{Th}_{\text{req}j}) - \text{PPU}(a)$$

Where:

- W_{Lj} , W_{Rj} , W_{Thj} are weights for latency, Packet Delivery Ratio (PDR), and throughput for each application j .
- $f(L_j)$ is a function that rewards lower latency.
- $\text{PPU}(a)$ is a penalty for interfering with primary users.

Crucially, priority weights (W_{Lj} , W_{Rj}) are scaled by the service's Priorityval (e.g., 5x for safety-critical applications). Furthermore, contextual urgency can dynamically boost these weights near critical events like accidents, ensuring immediate prioritization.

Adaptation Advantages

This Meta-RL approach provides significant benefits:

- **Generalization:** Policies learned by CDOF can effectively transfer and perform well in previously unseen road conditions or PU patterns.
- **Resilience:** The system can recover QoS performance within an impressive 100 milliseconds after network disruptions, such as sudden changes in traffic or new interference sources.

Simulation & Evaluation

Setup

We validated CDOF using a robust simulation environment:

- **Tools:** SUMO [28] for realistic vehicle mobility and NS-3 [29] for detailed network simulations.
- **Scenarios:** Tested across diverse environments: urban (5 km²), highway (10 km), and mixed traffic patterns.
- **Applications:**
 - **Safety:** Small, frequent messages (50 Bytes/10 Hz) with an ultra-low latency requirement ($L_{req} < 10$ ms).
 - **Infotainment:** Larger, bursty data (1024 Bytes) with a higher throughput requirement ($T_{req} > 5$ Mbps).
- **Benchmarks:** We compared CDOF against industry standards and state-of-the-art schemes: Dedicated Short Range Communications (DSRC), SURF [7], and a conventional Static RL approach.

Key Metrics

Performance was evaluated using:

- **QoS Satisfaction Rate (QoSSR):** The percentage of time that all required QoS parameters (L_{req} , R_{req} , T_{req}) are met.
- **P99 Latency:** The 99th-percentile latency, indicating the maximum latency experienced by 99% of packets, a critical metric for safety.
- **Adaptation Time:** The time taken for the system to recover its QoS performance after a significant disruption.

Results

CDOF consistently demonstrated superior performance:

- **QoSSR:** In urban scenarios, CDOF achieved a remarkable 98.2% QoSSR for safety-critical applications, significantly outperforming SURF (85.7%) and DSRC (76.1%).
- **P99 Latency:** For safety applications, CDOF maintained an impressive 8.2ms P99 latency, compared to 14.5ms for Static RL.
- **Adaptation:** CDOF recovered QoS performance in a mere 86ms after an accident, whereas Static RL required over 500ms.

- Generalization: Even in previously unseen rural areas, CDOF maintained a strong 94.1% QoSSR, far surpassing Static RL's 72.3%.

CONCLUSION & FUTURE WORK

The Cognitive-Driven Orchestration Framework (CDOF) is a pioneering Meta-RL-driven solution for CR-VANETs, successfully addressing the complex challenges of dynamic resource management and multi-service QoS guarantees. By leveraging "learning to learn," CDOF achieves:

- An exceptional 98.2% QoS satisfaction for critical safety applications.
- Rapid 86ms adaptation to network disruptions, such as accidents.
- Seamless generalization to new and unseen environments.

CDOF's capabilities pave the way for robust, commercially viable, and QoS-guaranteed vehicular networks.

Our future work will focus on:

- Testbed Validation: Implementing CDOF on Software-Defined Radio (SDR)-based platforms for real-world testing.
- Scalability: Exploring Federated Meta-RL to enable city-scale deployments.
- 5G/6G Integration: Investigating how CDOF can leverage advanced capabilities like network slicing in future cellular generations.
- Security: Addressing privacy concerns related to contextual data sharing.

REFERENCES

1. H. Hartenstein, "VANET: Vehicular Applications and Inter-Networking Technologies," Wiley, 2010.
2. M. Arif et al., "QoS in VANETs: A Survey," IEEE TVT, 2020.
3. G. Araniti et al., "LTE for Vehicular Networking: A Survey," IEEE COMST, 2013.
4. K. Dar et al., "CR-VANETs: Challenges and Solutions," IEEE Surveys & Tutorials, 2016.
5. I. F. Akyildiz et al., "Cognitive Radio Networks," IEEE JSAC, 2008.
6. A. Ali et al., "CR-Based MAC Protocols for VANETs," IEEE TMC, 2018.
7. Y. Wang et al., "SURF: Channel Selection for UAV-Assisted CR-VANETs," IEEE IoT Journal, 2021.
8. L. Liang et al., "Dynamic Spectrum Management in VANETs," IEEE TVT, 2017.
9. T. Yucek et al., "Spectrum Sensing for CR," IEEE COMST, 2009.
10. E. Axell et al., "Cooperative Spectrum Sensing," IEEE SP, 2012.
11. A. C. Talay et al., "CR-Enabled IEEE 802.11p," IEEE VTC, 2014.
12. O. A. Dobre et al., "Multi-Channel MAC Protocols," IEEE COMST, 2013.
13. S. Ucar et al., "Priority-Aware MAC for VANETs," IEEE TVT, 2016.
14. M. S. Almalag et al., "CR-Based Routing in VANETs," IEEE ICC, 2012.
15. Y. Zhang et al., "Power Control in CR-VANETs," IEEE TWC, 2019.
16. M. Boban et al., "Challenges in Vehicular Networking," IEEE COMST, 2014.
17. K. Abboud et al., "Congestion Control in VANETs," IEEE TMC, 2016.
18. C. Campolo et al., "QoS Support in VANETs," IEEE ICC, 2015.
19. A. Bazzi et al., "Resource Allocation for VANETs," IEEE TVT, 2018.
20. S. Zeadally et al., "Multi-Objective VANET Optimization," IEEE ITS, 2019.
21. M. A. Khan et al., "QoS Routing Using PSO," IEEE Access, 2020.
22. Y. S. Nasir et al., "RL for Wireless Networks," IEEE COMST, 2019.
23. L. Liang et al., "RL for DSA," IEEE JSAC, 2019.
24. H. Ye et al., "RL for Resource Allocation," IEEE TWC, 2019.
25. C. Finn et al., "Meta-Learning for RL," ICML, 2017.

-
26. Y. Chen et al., "Generalization Challenges in RL," NeurIPS, 2021.
 27. C. Finn et al., "Model-Agnostic Meta-Learning," ICML, 2017.
 28. D. Krajzewicz et al., "SUMO: Simulation of Urban Mobility," IEEE ITS, 2012.
 29. G. F. Riley et al., "NS-3: Network Simulator," ACM M&S, 2003.