

Quantum-Inspired Road Analysis and AoI-Aware Transmission for Intelligent V2I

Juhui Heo and Soohyun Park

Division of Computer Science, Sookmyung Women's University, Seoul, Republic of Korea

E-mails: {heojuhui, soohyun.park}@sookmyung.ac.kr

Abstract—While the advancement of autonomous driving technologies has improved traffic efficiency, maintaining reliable information sharing in cooperative adaptive cruise control (CACC) systems remains a significant challenge because stand-alone vehicle-to-vehicle (V2V) communications in a CACC platoon cannot guarantee the reliable dissemination of road information across all vehicles. To tackle the limitation, this paper proposes a novel framework for optimizing road information transmission in proactive non-platoon vehicles (e.g., connected vehicles (CVs)). The proposed framework using CVs consists of two stages. In Stage 1, a quantum graph-based object detection (QGOD) model is adopted for fewer parameter utilization in each CV. The QGOD-based model can dynamically control the detection accuracy and inference latency under trade-off modeling using Lyapunov optimization. In Stage 2, the road situation transmission is formulated using a Markov decision process, and the corresponding optimal information transmission from CVs to roadside units (RSUs) can be realized via proximal policy optimization, considering both accuracy and age-of-information (AoI). Simulation results demonstrate that the proposed framework achieves comparable detection accuracy with significantly fewer parameters and no inference-time penalty compared to conventional CNN-based methods. Furthermore, by scheduling road information based on AoI and accuracy, the proposed framework provides the most stable freshness performance while improving efficiency and maintaining training stability.

Index Terms—Age-of-Information (AoI), Vehicle-to-Infrastructure (V2I), Quantum AI, Reinforcement Learning

I. INTRODUCTION

The advancement of autonomous driving technologies has opened a new paradigm for improving traffic efficiency. In particular, cooperative adaptive cruise control (CACC) has emerged as a key technology that ensures the stability of platooning by enabling vehicles to maintain short headways through vehicle-to-vehicle (V2V) communications [1]. Maintaining such short headways, however, makes the platoon vulnerable to imperfections inherent in wireless vehicle-to-everything (V2X) communication. As CACC reviews note, transmission delay, limited range, and packet loss can all degrade string stability, allowing disturbances to propagate and amplify along the platoon and thereby raising the risk of rear-end collisions [2]. These vulnerabilities are most consequential in on-ramp merging, where a lost V2V link or a failed vehicle-identification step prevents vehicles from cooperating to merge safely [3]. Stand-alone V2V is therefore insufficient to guarantee reliable dissemination of road information, motivating infrastructure-assisted approaches based on roadside units (RSUs) and vehicle-to-infrastructure (V2I) communication. Placing RSUs at such safety-critical points




Fig. 1: Overview of the proposed two-stage framework

extends sensing and communication coverage beyond V2V, enabling each vehicle to obtain both the positions of surrounding traffic and the identifiers required to initiate cooperation via RSU-relayed safety messages.

To this end, we consider a V2I environment, comprising connected vehicles (CVs) and RSUs, as illustrated in Fig. 1. In the system, the role of CVs is to analyze road situations, and RSUs store and manage the collected data from CVs before broadcasting it to the associated CACC platoon. The road situation contents each CV analyzes and forwards to its RSU serve as safety messages for the platoon, and producing and scheduling them in a lightweight, fresh, and accurate manner is the central problem of this work. Furthermore, the freshness and accuracy of road information are crucial for maintaining cooperative analysis and low-latency coordination among vehicles [4]. Therefore, (i) lightweight and real-time road analysis on CVs and (ii) optimal data transmission considering data quality (e.g., freshness and accuracy) are indispensable.

As a promising solution, we propose a two-stage novel framework for road analysis and data transmission. For road analysis in Stage 1, a quantum graph-based object detection (QGOD) model is employed, which integrates the lightweight feature extraction capability of quantum convolutional neural networks (QCNNs) with the relational learning ability of graph convolutional networks (GCNs) [5], [6]. This hybrid architecture makes it suitable for real-time vehicular environments. Aim-

ing for lightweight operation, QGOD is designed to achieve efficient and real-time detection under resource-constrained conditions. Accordingly, the objective of Stage 1 is not to maximize detection accuracy at any computational cost, but to achieve a balanced trade-off between accuracy, inference latency, and model complexity suitable for vehicular environments. Furthermore, Lyapunov-based control dynamically adjusts the depth of the quantum neural network (QNN), modeled by parameterized quantum circuits (PQC), to balance accuracy and delay, ensuring stable performance under heterogeneous and fluctuating vehicular hardware.

Consequently, road information, including freshness and accuracy detected through the proposed QGOD model, is transmitted to RSUs. In Stage 2, the transmission problem is optimized to ensure information credibility and temporal relevance. To achieve this, the age-of-information (AoI), which represents the freshness of the data, and accuracy determined by the analysis model in Stage 1, are considered as key components for the transmission problem. The problem can be formulated as a Markov decision process (MDP). Accordingly, among various deep reinforcement learning (RL) algorithms, a proximal policy optimization (PPO) is utilized to optimize it. The PPO prevents training instability by explicitly constraining policy updates through a clipped surrogate objective and achieves stable, low-variance updates via generalized advantage estimation (GAE) [7]. In particular, the proposed PPO utilizes a continuous latent space to represent discrete actions, thereby improving gradient smoothness. Furthermore, it incorporates threshold discretization and mask-based correction to reflect realistic communication constraints effectively.

II. PRELIMINARIES

A. System Models

This study considers a vehicular network environment composed of multiple CVs and RSUs. The total highway length is set to 1,000 m, where five RSUs are deployed, each with a 500 m coverage radius [8]. Each RSU coverage area is further divided into 5 subregions, resulting in a total of $H=25$ subregions. For each subregion h , the maximum AoI threshold $A_h^{\max}$ is assigned according to its traffic sensitivity, which reflects the density and mobility characteristics of vehicles and motorcycles. Each CV is equipped with a QCNN-based analysis module for lightweight road analysis and drives at a constant speed of 50 km/h [9], generating one piece of road situation content at every time slot. A content item is defined as $C_h(t) = (h, \text{type}, s_h, A_h(t))$, where h , type, s_h , and $A_h(t)$ denote the subregion index, situation level, detection accuracy, and current AoI, respectively. Each CV stores at most one content per subregion, replacing outdated content upon regeneration and removing it once transmitted to the corresponding RSU.

B. AoI Models

When each CV detects the road situation in its corresponding subregion, the freshness of the generated content is quantified using the AoI model. Let $A_h(t)$ denote the AoI of the content associated with subregion h at time t . If no update is received,




Fig. 2: Architectures of the proposed quantum analysis modules

the AoI evolves linearly as $A_h(t+1) = A_h(t) + 1$, whereas upon the generation or successful reception of new content, the AoI is reset according to $A_h(t+1) = 0$. This AoI formulation is applied in both Stage 1 and Stage 2, providing a unified metric for balancing freshness and system efficiency. The AoI increases by one unit at each time slot and resets to zero when the corresponding subregion content is updated.

C. Quantum-Inspired Detection Models

The analysis model embedded in CVs must support both real-time processing and lightweight deployment. CNNs have been widely adopted for vision-based autonomous driving owing to their high accuracy and fast inference [10], yet their large parameter counts and heavy computation hinder real-time deployment on resource-constrained vehicular hardware, a difficulty further compounded by low-resolution inputs and dynamic road conditions [11]. QCNNs replace classical convolutions with quantum gates and measurements, enabling compact feature extraction with fewer parameters [12], but they inherently exhibit an accuracy–latency trade-off in which deeper PQCs improve accuracy at the cost of increased inference delay [5]. To mitigate this, we employ the QGOD model, which combines QCNN-based lightweight feature extraction with GCN-based relational learning: by encoding local quantum features via QCNNs and modeling inter-patch correlations via GCNs, QGOD attains compact representations while maintaining competitive analysis performance. The proposed architecture thus balances real-time performance, computational efficiency, and analysis effectiveness in resource-constrained vehicular environments.

III. PROPOSED METHOD

A. Stage1. Queue-Aware QGOD Analysis Model Control

In Stage 1, we consider CV-side analysis, where incoming image frames are queued for processing.

1) *Quantum Analysis Modules*: The proposed quantum analysis framework consists of two variants, as illustrated in Fig. 2, i.e., the quantum convolution-based object detection (QCOD) model and its extended version, QGOD.

Quantum Convolutional Layer. Both QCOD and QGOD first employ a lightweight classical CNN backbone to extract initial feature maps, which are subsequently processed by the quantum convolutional layer. This layer divides the input image into patches, encodes them into quantum states, and performs quantum operations using rotation (R_x, R_y, R_z), controlled-NOT (CNOT), and CU3 gates. Rotation gates map real-valued inputs to quantum states, CNOT gates capture local correlations through entanglement, and CU3 gates perform data-dependent transformations. This design achieves high representational efficiency with significantly fewer parameters while effectively capturing nonlinear relationships within each patch.

Graph Convolutional Layer (QGOD only). In QGOD, the output of the quantum convolutional layer is further processed by a graph convolutional layer, which transforms the feature map into a grid-based graph representation. Each grid cell is defined as a node connected to its four neighboring nodes and itself. The GCN aggregates neighboring node features according to $h_v^{(t+1)} = \sigma \left(\sum_{u \in \mathcal{N}(v) \cup \{v\}} \hat{A}_{vu} W^{(l)} h_u^{(t)} \right)$, thereby expanding local quantum features into global contextual representations. This additional layer enhances spatial consistency and multi-scale contextual reasoning, improving robustness under complex road environments.

Classical Head. Finally, the enhanced multi-scale features—either the QCOD output or its GCN-enhanced counterpart in QGOD—are fed into the classical head. At each grid location, multiple anchors with different scales and aspect ratios are generated, and the detection head jointly predicts objectness, bounding box offsets, and class probabilities based on IoU-based matching.

2) *Queue Modeling:* The inference accuracy improves with deeper PQC layers at the cost of increased latency. To model this trade-off, we introduce a discrete-time queue defined as $Q[t+1] = \max\{Q[t] - b(\alpha[t]), 0\} + a[t]$, where $Q[t]$ denotes the queue length at time t , $a[t]$ represents the arrival process (new images), and $b(\alpha[t])$ is the service rate determined by the PQC depth $\alpha[t]$. Maintaining queue stability ensures real-time responsiveness, while larger $\alpha[t]$ improves accuracy.

3) *Lyapunov Optimization:* Based on the queue model, we define the Lyapunov function as the quadratic queue backlog, $L(Q[t]) = \frac{1}{2}Q[t]^2$, and the one-step conditional drift as $\Delta(Q[t]) = \mathbb{E}\{L(Q[t+1]) - L(Q[t]) \mid Q[t]\}$. To jointly consider queue stability and inference accuracy, we adopt the drift-plus-penalty (DPP) framework defined as $\Delta(Q[t]) - V \cdot \mathbb{E}\{A(\alpha[t]) \mid Q[t]\}$, where $A(\alpha[t])$ denotes the accuracy under PQC depth $\alpha[t]$, and V is a tunable parameter controlling the trade-off. Minimizing an upper bound of the DPP leads to the following decision rule,

$$\alpha^*[t] = \arg \max_{\alpha[t] \in \mathcal{A}} \left\{ V \cdot A(\alpha[t]) - Q[t] \cdot b(\alpha[t]) \right\}. \quad (1)$$

A larger V prioritizes higher inference accuracy with deeper PQC layers, while a smaller V emphasizes queue stability and low latency. This ensures that the time-averaged queue backlog remains bounded (i.e., $\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T Q[t] < \infty$), while the long-term average analysis accuracy asymptotically approaches

near-optimality as $\max \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T A(\alpha[t])$.

Intuitively, when the queue is heavily loaded, the controller selects a shallower PQC depth to accelerate processing at the cost of reduced accuracy. Conversely, when the queue is lightly loaded, deeper PQC layers are applied to maximize accuracy. This adaptive mechanism guarantees bounded latency while enhancing the analysis quality, serving as the foundation for Stage 2, where AoI-aware content scheduling is performed.

B. Stage2. AoI-Aware Content Scheduling with PPO

1) *MDP Formulation:* In Stage 2, the transmission scheduling is formulated as MDP.

- **State:** The state at time t is defined as $s(t) = \{A_h(t), s_h(t)\}_{h=1}^H$, where $A_h(t)$ and $s_h(t)$ represent the AoI and accuracy score of the buffered content for subregion h , respectively. The normalized detection accuracy s_h from Stage 1 serves as the accuracy indicator in Stage 2.
- **Action:** The policy outputs a latent variable $z(t)$ that is discretized into one action $a(t) \in \{0, 1, \dots, H\}$, where $a(t) = 0$ indicates no transmission, and $a(t) = h$ denotes transmission of the content for subregion h . Only one content can be transmitted per time slot ($\mathbf{1}^\top a(t) \leq 1$).
- **Reward:** The instantaneous reward is defined as

$$R(t) = R^{\text{pos}}(t) + \beta_{\text{eff}}(t) R^{\text{neg}}(t), \quad (2)$$

where $R^{\text{pos}}(t)$ measures the positive contribution of a successful transmission, and $R^{\text{neg}}(t)$ penalizes the accumulation of outdated or unserved contents in other subregions.

To explicitly represent the transmission decision, we introduce a binary indicator $\delta_{a(t),h}$, which equals 1 if the selected action $a(t)$ corresponds to subregion h , and 0 otherwise. Using $\delta_{a(t),h}$, the reward components can be reformulated accordingly, where η_s , η_f , and η_g denote the weighting factors for accuracy, freshness, and AoI gain, respectively. In all experiments, we set $\eta_f > \eta_s$ to prioritize freshness over accuracy. The adaptive penalty weight $\beta_{\text{eff}}(t)$ dynamically scales the negative reward using the clipped EMA ratio of recent negative to positive rewards, reducing it when negatives dominate and increasing it otherwise to maintain reward balance and training stability. The clipping bounds $[0.5, 2.0]$ are introduced to prevent abrupt fluctuations in $\beta_{\text{eff}}(t)$. Finally, the AoI of subregion h is reset if selected at time t and otherwise increases by one. Hence, the long-term average objective is formulated as follows, $\max \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T R(t)$.

2) *Learning Algorithm with PPO:* The proposed PPO-based learning algorithm is designed with three key components. First, the **continuous latent** ($z(t)$) represents the action as a one-dimensional continuous variable, which improves lightweight efficiency and enhances gradient smoothness, thereby stabilizing the optimization process. Second, the **threshold discretization** maps the continuous variable z into a discrete action space, enabling non-continuous transmission decisions while maintaining policy stability. Third, the **mask-based correction** adjusts invalid actions to the nearest feasible ones, reflecting realistic communication constraints and enhancing exploration stability.

$$R^{\text{pos}}(t) = \sum_{h=1}^H \delta_{a(t),h} \left[\eta_s s_h(t) + \eta_f \left(1 - \frac{A_h(t)}{A_h^{\max}} \right) + \eta_g \Delta A_h(t) \right], \quad (3)$$

$$R^{\text{neg}}(t) = -\frac{1}{H-1} \sum_{h=1}^H \left[(1 - \delta_{a(t),h}) s_h(t) \times \tanh\left(\frac{A_h(t)}{A_h^{\max}}\right) \right], \quad (4)$$

TABLE I: PPO variants for content scheduling

Algorithm	Removed Component(s)
Proposed	None (full structure)
Ordinal Softmax	(1) Continuous latent
Continuous-z	(2) Threshold discretization
No-Mask Softmax	(3) Mask-based correction

By combining these three components, the proposed PPO framework is tailored for AoI-, content-, and accuracy-aware scheduling problems. The policy network takes the system state $s(t)$ as input and outputs a continuous variable $z(t) \in [-1, 1]$ following a tanh-squashed Gaussian distribution. The interval $[-1, 1]$ is uniformly divided into $H+1$ bins, each corresponding to a discrete action $\{a(t) = 0, 1, \dots, H\}$. Invalid actions are corrected to the nearest valid ones through the masking mechanism described above. The PPO objective combines a clipped policy loss, value function regression (i.e., mean square error), and entropy regularization, while employing GAE and gradient clipping for training stability. Through this design, the agent learns to balance information freshness (i.e., AoI) and content accuracy, supporting effective road analysis and real-time responsiveness while maintaining lightweight communication efficiency.

IV. PERFORMANCE EVALUATION

A. Simulation Setting

All experiments were conducted in PyTorch on an NVIDIA RTX 4080 SUPER GPU. In **Stage 1**, classical CNN, QCOD, and QGOD were implemented to evaluate the impact of varying PQC depths and post-activation channel configurations. The proposed QGOD employed Lyapunov-based optimization to dynamically adjust PQC depth for balancing accuracy and queue stability. The KITTI dataset (5,223 training, 746 testing) was converted to grayscale before augmentation, and its six object classes were merged into three (Vehicle, Pedestrian, Cyclist) for evaluation, with square-root inverse-frequency positive weights (normalized to vehicle) applied to mitigate the vehicle-dominant class imbalance. The class-agnostic mAP was used as the evaluation metric, emphasizing localization performance. In **Stage 2**, the transmission process was formulated as an MDP and optimized using various RL algorithms. An ablation study on PPO-based variants (Table I) was conducted to analyze the contribution of each architectural component.

B. Experimental Results

1) *Stage1*: As presented in Table II, the proposed QGOD model achieved an mAP of 89.10% with only 87,631 parameters and an inference time of 1.46×10^{-2} seconds. Compared to

TABLE II: Performance and complexity comparison (QCOD and QGOD use one PQC layer, while CNN has none)

Model	mAP@0.5 (%)	Number of Parameters	Inference Time (sec)
Classical CNN	89.01	143,215	1.51×10^{-2}
QCOD (32 channels)	88.24	86,754	1.82×10^{-2}
QCOD (64 channels)	89.01	92,252	2.08×10^{-2}
QGOD (32 channels)	89.10	87,631	1.46×10^{-2}


TABLE III: QGOD accuracy and inference time vs. PQC depth

PQC Layers	Accuracy (%)	Inference Time (sec)
1	89.10	1.46×10^{-2}
3	89.90	1.67×10^{-2}
5	90.81	1.91×10^{-2}

the classical CNN, QGOD reduces the parameter count by about 39% (87,631 vs. 143,215) with no inference-time penalty (1.46 vs. 1.51×10^{-2} seconds) at comparable accuracy. It also outperforms the 64-channel QCOD in both accuracy and speed, while the 32-channel QCOD degrades in both.

To enable real-time inference, the proposed framework dynamically adjusts the number of PQC layers in QGOD within a Lyapunov-based optimization framework. As shown in Table III, increasing PQC depth improves accuracy from 89.10 to 90.81% but also increases inference time from 1.46×10^{-2} to 1.91×10^{-2} seconds, revealing a clear accuracy-latency trade-off. To address this trade-off, the proposed system employs a Lyapunov controller that adaptively adjusts the PQC layer configuration according to the real-time queue state. As illustrated in Fig. 3(a), a larger control parameter V prioritizes accuracy by selecting deeper models but increases delay, while a smaller V improves latency efficiency at the cost of accuracy. Therefore, appropriately tuning V is essential to maintain a stable balance between accuracy and latency.

2) *Stage2*: Figure 3(b) compares the training reward trends of the RL algorithms, and the dotted line presents their normalized variance values for stability analysis. The Proposed Algorithm converges fastest, stabilizing within about 300 episodes and maintaining an average reward of 800-850 thereafter. In contrast, Ordinal Softmax remains below 50, Continuous-z gradually declines to around 400, and No-Mask Softmax increases steadily to about 800-850. SAC and DQN stagnate near 700-720 and 30-50, respectively, showing poor learning performance. These results stem from the complex vehicular environment, which features a 26-dimensional discrete action space (25 transmission candidates and one non-transmission action) whose valid action set dynamically changes over time. Such a non-stationary structure poses challenges for conventional algorithms like DQN and SAC, particularly under dynamically varying



(a) Queue backlog dynamics of the Lyapunov-optimized QGOD (b) RL reward and convergence comparison

Fig. 3: Lyapunov and RL Performance Analysis

TABLE IV: Statistical Summary of Freshness and Accuracy

Method	Freshness			Accuracy		
	Mean	Std	μ/σ	Mean	Std	μ/σ
Proposed	0.9429	0.0032	294.66	0.4298	0.0203	21.17
Ordinal Softmax	0.6324	0.0474	13.34	0.4313	0.0635	6.79
Continuous-z	0.9222	0.0071	129.88	0.4288	0.0315	13.61
No-Mask Softmax	0.9436	0.0038	248.32	0.5724	0.0249	22.99
SAC	0.9304	0.0059	157.69	0.4763	0.0217	21.95
DQN	0.0180	0.0223	0.81	0.4269	0.0974	4.38

discrete constraints. This may limit their ability to learn stable policies in dynamically constrained action spaces. In contrast, the proposed algorithm maintains a low variance (≈ 0.25 - 0.35) throughout training, while all baselines exhibit higher fluctuations (≈ 0.6 - 0.8). Overall, it achieves higher final rewards, faster convergence, and more stable training behavior compared to the baselines. Without continuous latent representation and threshold discretization, other algorithms fail to handle the high-dimensional action space effectively. Although the No-Mask Softmax achieved moderate performance, its slower convergence highlights the critical role of mask-based correction in accelerating convergence and stabilizing training.

As shown in Table IV, Freshness is defined as the normalized AoI, while Accuracy is measured by mAP from the QGOD detector. The μ/σ ratio (mean over standard deviation) is reported, where a higher value indicates lower variability. The proposed method achieves the highest Freshness μ/σ (294.66) and competitive Accuracy μ/σ (21.17), close to No-Mask Softmax (22.99). These results indicate that the proposed algorithm achieves the most stable freshness performance while maintaining comparable accuracy, demonstrating a favorable trade-off between performance and variability.

V. CONCLUDING REMARKS

This paper proposes the QGOD model, which integrates QCNN-based lightweight quantum feature extraction with GCN-based relational reasoning. Experimental results show that QGOD achieves detection accuracy comparable to the classical CNN and higher than the QCNN-based QCOD model, while requiring significantly fewer parameters and a competitive inference time. In addition, the proposed Lyapunov-based adaptive control mechanism effectively mitigates the accuracy-latency trade-off by dynamically adjusting the number of PQC layers according to the queue state, ensuring stable and real-time performance under varying computational loads. In

Stage 2, the proposed RL-based content scheduling algorithm achieves the highest reward as well as the fastest convergence, supported by the complementary effects of continuous latent representation, threshold-based discretization, and mask-based correction. In particular, the mask-based correction enables immediate adjustment from invalid to valid actions, reducing redundant exploration and accelerating convergence. Overall, the proposed framework provides an effective solution for real-time perception and scheduling in intelligent vehicular systems, achieving competitive accuracy while improving parameter efficiency and maintaining stability.

ACKNOWLEDGMENT

This research was supported by MSIT (Ministry of Science and ICT), Korea, under ITRC (Information Technology Research Center) (IITP-2026-RS-2024-00436887) supervised by IITP (Institute for Information & Communications Technology Planning & Evaluation). Soohyun Park is a corresponding author (e-mail: soohyun.park@sookmyung.ac.kr).

REFERENCES

- [1] G. Ma, P. R. Pagilla, and S. Darbha, "Selection of Time Headway in Connected and Autonomous Vehicle Platoons Under Noisy V2V Communication," *IEEE Transactions on Intelligent Transportation Systems*, vol. 26, no. 1, pp. 1029–1038, Nov. 2025.
- [2] M. A. Tahiri, A. Rachid, B. Boudmane, I. Mortabit, and S. Laaroussi, "Toward Cooperative Adaptive Cruise Control: A Mini-review," in *Proc. IEEE International Conference on Circuits, Systems and Communication (ICCS)*, Fes, Morocco, Aug. 2024.
- [3] H. Liu, F. Jahedini, Z. Mu, and B. B. Park, "Evaluation of Connected Vehicle Identification-Aware Mixed Traffic Freeway Cooperative Merging," in *Proc. IEEE Intelligent Vehicles Symposium (IV)*, Jeju, Republic of Korea, Jul. 2024, pp. 1222–1229.
- [4] J. Rodríguez-Piñeiro, Z. Wei, J. Wang, C. A. Gutiérrez, and L. M. Correia, "6G-Enabled Vehicle-to-Everything Communications: Current Research Trends and Open Challenges," *IEEE Open Journal of Vehicular Technology*, vol. 6, pp. 2358–2391, Aug. 2025.
- [5] M. Kim, J. Heo, E. J. Roh, and S. Park, "Dynamic Multi-PQC Quantum Convolutional Neural Network for Real-Time Pothole Detection: Invited Paper," in *Proc. IEEE International Workshop on Signal Processing and Artificial Intelligence for Wireless Communications (SPAWC)*, Surrey, United Kingdom, Jul. 2025.
- [6] Q. Gao, T. Xu, X. Li, W. Gao, H. Shi, Y. Zhang, J. Chen, and Z. Yue, "Interpretable Dynamic Directed Graph Convolutional Network for Multi-Relational Prediction of Missense Mutation and Drug Response," *IEEE Journal of Biomedical and Health Informatics*, vol. 29, no. 2, pp. 1514–1524, Oct. 2025.
- [7] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal Policy Optimization Algorithms," *arXiv preprint*, 2017.
- [8] R. Chen, L. Gao, Y. Liu, Y. L. Guan, and Y. Zhang, "Smart Roads: Roadside Perception, Vehicle-Road Cooperation, and Business Model," *IEEE Network*, vol. 39, no. 3, pp. 311–318, Oct. 2025.
- [9] M. M. Balas, V. E. Balas, F.-M. Petcut, C. Barna, and R. Dey, "Automation of the Two-Second Rule via Constant Time-to-Collision Curves," in *Proc. IEEE International Conference on Intelligent Engineering Systems (INES)*, Palermo, Italy, Jun. 2025, pp. 155–158.
- [10] D. Chun, J. Choi, H.-J. Lee, and H. Kim, "CP-CNN: Computational Parallelization of CNN-Based Object Detectors in Heterogeneous Embedded Systems for Autonomous Driving," *IEEE Access*, vol. 11, pp. 52 812–52 823, May 2023.
- [11] M. Cheshfar, M. H. Maghami, P. Amiri, H. G. Garakani, and L. Lavagno, "Comparative Survey of Embedded System Implementations of Convolutional Neural Networks in Autonomous Cars Applications," *IEEE Access*, vol. 12, pp. 182 410–182 437, Dec. 2024.
- [12] E. J. Roh, H. Baek, D. Kim, and J. Kim, "Fast Quantum Convolutional Neural Networks for Low-Complexity Object Detection in Autonomous Driving Applications," *IEEE Transactions on Mobile Computing*, vol. 24, no. 2, pp. 1031–1042, Sep. 2025.