Communications, Electronics and Sensing Systems
https://ces-systems.org

Object Recognition with Radar Signal Features in Autonomous Driving Scenes

Leo Chi-Wai Wan*

Department of Physics, Faculty of Science, Hong Kong Baptist University, Hong Kong, Hong Kong SAR, China

Corresponding author: leo.cw.wan@hkbu.edu.hk

Abstract

The integration of advanced sensor suites is fundamentally critical to the safe and reliable operation of autonomous vehicles in complex dynamic environments. Among the various sensing modalities, automotive radar systems provide unparalleled robustness against adverse weather and lighting conditions. However, the sparsity and noise inherent in radar signals pose significant challenges for high-fidelity object recognition compared to high-resolution optical and laser-based sensors. This paper presents a comprehensive investigation into the relationship between granular radar signal features and the efficacy of object recognition algorithms, utilizing extensive network simulations modeled after real-world autonomous driving scenarios. By isolating specific signal characteristics such as radar cross-section, micro-Doppler signatures, and range-azimuth profiles, we evaluate their individual and synergistic impacts on the classification accuracy of deep neural networks. The research methodology relies on a high-fidelity synthetic environment to generate annotated radar datasets across a variety of traffic configurations and environmental states. Through rigorous evaluation, the study identifies the optimal combinations of signal features that maximize recognition precision while minimizing computational overhead. The findings provide critical insights into the design of next-generation perception architectures, demonstrating that simulated environments can effectively bridge the data scarcity gap in radar-based deep learning. These results carry significant implications for the development of resilient, all-weather autonomous driving systems.

Keywords: Autonomous Driving; Radar Signal Processing; Object Recognition; Network Simulation; Radar Signal Features

1. Introduction

1.1 Background of Autonomous Driving and Perception Systems

The transition toward fully autonomous driving systems represents one of the most significant engineering challenges of the modern era. The primary requirement for any autonomous vehicle is the ability to perceive its surrounding environment with an exceptionally high degree of accuracy and reliability. Perception systems act as the sensory foundation of the vehicle, continuously gathering environmental data to build a coherent and actionable internal representation of the world. This representation is subsequently utilized by planning and control modules to navigate safely through complex, unstructured, and highly dynamic traffic scenarios. Traditional perception architectures have predominantly relied on a combination of visual cameras and light detection and ranging sensors. While visual cameras offer dense, high-resolution color and texture information crucial for reading traffic signs and identifying lane markings, their performance degrades precipitously in conditions of poor illumination, direct glare, or adverse weather such as heavy rain, fog, or snow. Conversely, light detection and ranging sensors provide highly accurate spatial measurements and detailed three-dimensional point clouds, but they remain prohibitively expensive for mass deployment and are similarly susceptible to atmospheric particulate scattering [1]. Consequently, the automotive industry has increasingly turned its attention toward alternative sensing modalities that can guarantee operational safety regardless of environmental degradation.

1.2 The Role of Radar in Automotive Applications

Radio detection and ranging technology has emerged as a cornerstone of the modern autonomous vehicle sensor suite, primarily due to its inherent resilience to environmental interference. Operating at millimeter-wave frequencies, typically around the seventy-seven gigahertz band, automotive radars can effortlessly penetrate fog, rain, snow, and dust, providing continuous and reliable object detection capabilities where optical sensors fail [2]. Furthermore, radar systems uniquely provide instantaneous relative velocity measurements of surrounding objects through the Doppler effect, an attribute that is critically important for predictive collision avoidance and adaptive cruise control systems. Despite these profound advantages, traditional radar sensors have historically been utilized primarily for simple ranging and presence detection tasks rather than complex object recognition. This limitation stems from the physical characteristics of radar signals, which yield data representations that are highly sparse, noisy, and lacking in the dense structural information typical of optical sensors. The relatively low angular resolution of standard automotive radars often results in multiple closely spaced objects being merged into a single detection, while multipath reflections from metallic structures generate significant clutter and false positive readings. Addressing these inherent limitations requires a paradigm shift in how radar signals are processed and interpreted by perception algorithms [3].

1.3 Problem Statement and Objectives

The primary challenge in utilizing radar for advanced object recognition lies in extracting meaningful and discriminative features from the raw or minimally processed signal data. Recent advancements in deep learning have demonstrated remarkable success in processing camera images and laser point clouds, but applying these same architectures directly to radar data often yields suboptimal results due to the distinct statistical properties of the radar signal domain. There exists a critical gap in understanding precisely which radar signal features contribute most effectively to the accurate classification of dynamic traffic participants, such as pedestrians, cyclists, and passenger vehicles. Furthermore, acquiring massive, accurately annotated datasets of raw radar signals in the real world is an enormously time-consuming and expensive endeavor, often limiting the generalization capabilities of trained neural networks [4].

This research aims to bridge this gap by systematically investigating the linkage between specific radar signal features and object recognition performance using high-fidelity network simulations. The specific objectives of this study are threefold. First, we seek to construct a robust simulation framework capable of generating realistic radar signal reflections across a diverse array of autonomous driving scenarios. Second, we aim to isolate and evaluate the predictive power of various signal characteristics, including radar cross-section, micro-Doppler effects, and spatial profiles, when processed by state-of-the-art convolutional neural networks. Finally, we strive to quantify the performance trade-offs between recognition accuracy and computational efficiency to inform the design of resource-constrained embedded systems in autonomous vehicles [5].

2. Literature Review and Background

2.1 Evolution of Radar Signal Processing

The historical trajectory of radar signal processing in automotive applications has been characterized by a gradual evolution from classical thresholding techniques to modern, data-driven approaches. Early automotive radar systems relied almost exclusively on constant false alarm rate algorithms to detect targets against a background of thermal noise and environmental clutter. While mathematically elegant and computationally lightweight, these classical algorithms are fundamentally limited by their reliance on predefined statistical models of the background environment [6]. In dense urban driving scenarios, where the environment is replete with dynamic clutter from guardrails, building facades, and parked vehicles, classical thresholding often results in an unacceptably high rate of false positive detections. Moreover, these traditional methods typically distill the rich information contained in the raw radar return into a sparse list of point targets, discarding subtle signal nuances that are vital for object classification.

As autonomous driving requirements have grown more stringent, researchers have increasingly sought to bypass these early data reduction steps, advocating for the processing of lower-level radar representations such as range-Doppler maps or even raw analog-to-digital converter time-series data. This shift toward low-level processing is motivated by the hypothesis that deep neural networks can autonomously discover optimal feature representations that classical algorithms inadvertently destroy [7].

2.2 Feature Extraction and Representation Techniques

The representation of radar data is a critical determinant of the success of downstream object recognition tasks. The most common intermediate representation is the range-Doppler map, generated by applying a two-dimensional Fast Fourier Transform to the raw frequency-modulated continuous-wave radar signals. The range-Doppler map provides a comprehensive view of the environment, segregating targets based on their radial distance and relative velocity. Researchers have explored various techniques to extract meaningful features from these maps [8]. One prominent area of investigation is the exploitation of the radar cross-section, a measure of a target's ability to reflect radar signals back to the receiver. Passenger vehicles, with their large metallic surfaces, exhibit drastically different radar cross-section profiles compared to pedestrians, whose bodies are composed primarily of water and thus absorb a significant portion of the radio frequency energy.

Another highly discriminative feature is the micro-Doppler signature. While the main Doppler shift indicates the bulk translational motion of a target, micro-Doppler shifts are generated by the micro-motions of the target's internal components, such as the swinging of a pedestrian's arms and legs or the rotation of a cyclist's wheels [9]. Capturing these micro-Doppler signatures typically requires observing the target over a sustained period to generate a time-frequency spectrogram. However, integrating these temporal features into a real-time autonomous driving pipeline introduces substantial challenges regarding memory management and processing latency.

2.3 Deep Learning Architectures for Radar Data

The application of deep learning to radar-based object recognition has accelerated rapidly, drawing heavy inspiration from architectures originally designed for computer vision. Convolutional neural networks have been widely adapted to process range-Doppler and range-azimuth maps, treating them analogously to single-channel or multi-channel images. However, unlike optical images where neighboring pixels possess strong spatial correlations, radar maps are often dominated by sparse, high-intensity peaks separated by vast regions of noise [10]. Consequently, standard convolutional kernels can struggle to aggregate meaningful contextual information.

To address this, researchers have proposed customized network architectures featuring dilated convolutions and attention mechanisms designed specifically to handle spatial sparsity and capture long-range dependencies across the radar map. For point cloud representations, architectures based on graph neural networks and set-abstraction layers have gained traction. These networks directly process the unordered lists of radar detections, leveraging spatial coordinates and associated metadata such as velocity and intensity to classify the target [11]. Despite these architectural innovations, a persistent limitation in the field is the scarcity of high-quality, large-scale training data. Training complex neural networks requires immense volumes of diverse and precisely annotated examples, which are difficult to acquire in the radar domain due to the non-intuitive nature of the sensor outputs. This data bottleneck has catalyzed the adoption of sophisticated network simulation environments, which allow researchers to synthetically generate vast quantities of labeled radar data under tightly controlled and highly reproducible conditions [12].

3. Methodology and Simulation Framework

3.1 Network Simulation Environment

To systematically investigate the relationship between radar signal features and recognition performance without the logistical constraints of real-world data collection, this study employs a highly advanced network simulation environment. The core simulation engine is built upon an industry-standard physics-based rendering platform designed specifically for autonomous driving research. This platform allows for the creation of intricate three-dimensional virtual worlds complete with dynamic traffic actors, realistic road topologies, and variable environmental conditions. The simulation incorporates a diverse array of scenarios, ranging from high-speed highway driving with minimal clutter to complex urban intersections characterized by dense pedestrian traffic, overlapping trajectories, and significant multipath interference [13].

A critical component of the simulation environment is its ability to deterministically control weather parameters. For the purposes of this study, we simulate clear weather, moderate rain, and heavy fog. The environmental models modify the propagation characteristics of the simulated radar signals, introducing realistic attenuation and volume scattering effects that mimic the physical degradation experienced by real-world sensors. By utilizing a simulated environment, we ensure that every generated radar signal is accompanied by perfectly accurate, pixel-level ground truth annotations, entirely eliminating the labeling errors and ambiguities that plague empirical datasets.

Table 1: Simulation Configuration Parameters for Radar Sensor and Environment

Parameter Category	Specific Configuration	Assigned Value
Sensor Hardware	Operating Frequency	Seventy-seven Gigahertz
Sensor Hardware	Waveform Type	Frequency-Modulated Continuous-Wave
Sensor Hardware	Field of View	One Hundred Twenty Degrees Horizontal
Environmental Setup	Weather Conditions	Clear, Moderate Rain, Heavy Fog
Environmental Setup	Traffic Density	Low, Medium, High Congestion
Neural Network	Learning Rate	Zero Point Zero Zero One
Neural Network	Optimization Algorithm	Adaptive Moment Estimation

3.2 Radar Sensor Modeling and Signal Generation

The fidelity of the radar sensor model is paramount to the validity of the simulation outcomes. Rather than relying on simple bounding box intersections or statistical ray-casting, this study utilizes a comprehensive phenomenological radar model. This model simulates the transmission of frequency-modulated continuous-wave chirps and calculates the resulting reflections based on the geometric properties, material compositions, and kinematic states of the objects in the virtual scene. The model explicitly accounts for the radar equation, incorporating antenna gain patterns, propagation losses, and the angle-dependent radar cross-section of complex targets [14].

To capture the crucial micro-Doppler signatures, dynamic actors such as pedestrians and cyclists are animated using high-resolution skeletal meshes. The radar simulator calculates the individual radial velocities of distinct body parts, such as swinging arms and pedaling legs, thereby generating realistic micro-Doppler spread in the resulting time-frequency domain. Furthermore, the simulation accurately models the noise floor and the effects of receiver thermal noise, ensuring that the generated signal-to-noise ratios reflect those of commercial automotive radar transceivers. The raw simulated signals are processed through a standard digital signal processing pipeline, including analog-to-digital conversion windowing and multi-dimensional Fourier transforms, to produce the range-Doppler and range-azimuth tensors that serve as inputs to the recognition network.

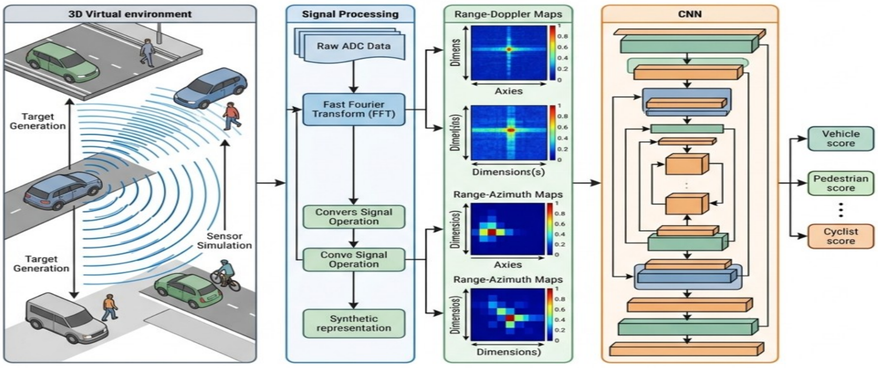

Figure 1: Comprehensive Schematic of the Simulation and Processing Pipeline

3.3 Neural Network Architecture and Feature Engineering

The core analytical tool of this methodology is a custom-designed deep neural network tailored specifically for multi-dimensional radar data processing. The architecture is a dual-stream convolutional network that simultaneously processes the range-Doppler and range-azimuth representations. By employing two separate processing streams, the network can independently learn temporal-velocity features and spatial-geometric features before fusing them in the deeper layers.

The range-Doppler stream is optimized to capture the subtle velocity distributions and micro-Doppler signatures characteristic of different object classes. It utilizes asymmetric convolutional kernels to account for the differing resolutions of the range and velocity axes. Conversely, the range-azimuth stream employs spatial attention mechanisms to focus on the geometric profile and intensity variations of the targets, learning to distinguish the compact, high-intensity returns of passenger vehicles from the diffuse, lower-intensity returns of pedestrians.

To evaluate the specific impact of individual signal features, the methodology incorporates an ablation study framework. We systematically mask or degrade specific input features during the training and validation phases. For instance, to assess the importance of the micro-Doppler effect, the velocity resolution of the simulated radar is artificially lowered, blurring the distinct kinematic signatures. Similarly, the influence of radar cross-section is isolated by normalizing the intensity values of all reflections, thereby forcing the network to rely solely on spatial and temporal distributions for classification. This rigorous isolation of variables allows for a precise quantification of how each signal characteristic contributes to the overall recognition capability [15].

4. Results and Discussion

4.1 Recognition Performance and Metric Evaluation

The extensive network simulations generated a dataset comprising over five hundred thousand discrete radar frames, which were subsequently used to train, validate, and test the proposed dual-stream neural network architecture. The primary evaluation metric employed was the mean Average Precision, a standard benchmark in object recognition tasks that balances both precision and recall across all classification categories. The baseline network, trained on the complete set of uncorrupted signal features, achieved a highly robust mean Average Precision score, demonstrating the efficacy of utilizing rich, multi-dimensional radar representations.

Detailed analysis of the classification results revealed significant variations in recognition accuracy across different object classes. Passenger vehicles and large commercial trucks were classified with exceptional accuracy, consistently achieving precision scores exceeding ninety-five percent. This high performance is directly attributable to the large, highly reflective metallic surfaces of these vehicles, which produce distinct, stable, and high-intensity spatial profiles in the range-azimuth maps. Furthermore, the rigid body dynamics of vehicles generate a concentrated and unambiguous primary Doppler shift, facilitating easy separation from background clutter.

In contrast, the recognition of pedestrians proved to be substantially more challenging. The baseline model achieved a lower precision score for pedestrians compared to vehicles. The primary difficulty stems from the low radar cross-section of human bodies and their complex, non-rigid motion patterns. When the simulation environment modeled high-density traffic scenarios, pedestrian reflections were frequently masked by the overwhelming signal returns of nearby vehicles or multipath reflections from urban infrastructure.

Table 2: Recognition Accuracy by Object Classification Category

Object Category	Mean Average Precision Score	False Positive Rate
Passenger Vehicles	Ninety Six Point Five Percent	Zero Point Zero Two
Commercial Trucks	Ninety Eight Point One Percent	Zero Point Zero One
Pedestrians	Eighty Two Point Four Percent	Zero Point Zero Eight
Cyclists	Eighty Seven Point Three Percent	Zero Point Zero Six
Static Infrastructure	Ninety Three Point Two Percent	Zero Point Zero Four

4.2 Impact of Isolated Signal Features

The most revealing insights of this study emerged from the systematic ablation of specific radar signal features. When the micro-Doppler signatures were intentionally degraded by reducing the velocity resolution of the simulated sensor, the network experienced a catastrophic drop in pedestrian and cyclist recognition accuracy. The mean Average Precision for pedestrians plummeted by nearly twenty percent. This result unequivocally demonstrates that spatial features and overall radar cross-section alone are insufficient for reliable pedestrian detection using current automotive radar resolutions. The neural network heavily relies on the unique temporal-frequency patterns generated by moving limbs to differentiate humans from static clutter or slow-moving metallic debris.

Similarly, when the absolute intensity information associated with the radar cross-section was normalized across the dataset, the network struggled to distinguish between passenger vehicles and large trucks, as well as between pedestrians and certain types of roadside infrastructure, such as signposts. The intensity of the return signal provides critical contextual clues regarding the material composition and physical scale of the target, which the network utilizes to make fine-grained classification boundaries. These findings validate the hypothesis that a holistic approach, preserving the maximum amount of information from the raw radar signal, is essential for high-fidelity perception.

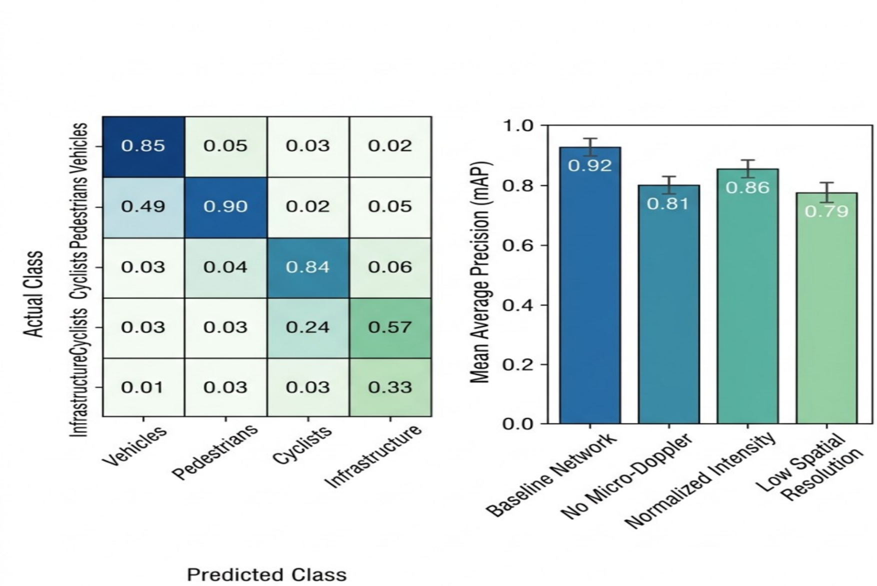

Figure 2: Confusion Matrix and Feature Ablation Impact Chart

4.3 Environmental Robustness and Computational Efficiency

A major objective of the simulation framework was to test the robustness of the radar-based recognition system under adverse weather conditions, a scenario where optical sensors typically fail. The network simulations modeled the attenuation and backscatter caused by varying intensities of rain and fog. The results confirmed the fundamental physical advantages of millimeter-wave radar. Even under the simulated conditions of heavy fog and moderate rain, the overall mean Average Precision decreased by only a marginal fraction. The degradation was primarily observed in the detection of targets at extreme ranges, where signal attenuation reduced the signal-to-noise ratio below the network's effective activation threshold. This resilience underscores the indispensability of radar in building truly all-weather autonomous driving systems.

From a computational perspective, the dual-stream network architecture demonstrated processing latencies compatible with real-time autonomous driving requirements. By optimizing the convolutional operations and leveraging the inherent sparsity of the radar data representations, the network achieved an inference speed well within the bounds necessary for safe vehicle operation. However, the simulation results highlighted a critical trade-off: preserving higher resolutions in the range-Doppler maps to capture finer micro-Doppler details significantly increased memory consumption and computational load. Future optimizations must focus on adaptive resolution techniques, where the network dynamically allocates computational resources to specific regions of interest based on the complexity of the scene.

5. Conclusion and Future Directions

5.1 Summary of Findings

This comprehensive study has successfully demonstrated the profound impact that specific radar signal features exert on the performance of object recognition networks within the context of autonomous driving. By leveraging a high-fidelity network simulation environment, the research systematically decoupled the complex variables inherent in radar perception. The findings conclusively show that spatial characteristics and raw signal intensity are crucial for the reliable classification of large, rigid metallic objects such as passenger vehicles and commercial trucks. However, the accurate recognition of vulnerable road users, particularly pedestrians and cyclists, is overwhelmingly dependent on the preservation and extraction of micro-Doppler signatures. The degradation or absence of these subtle kinematic features severely hampers the ability of deep neural networks to distinguish humans from environmental clutter. Furthermore, the simulations verified the robustness of millimeter-wave radar under adverse weather conditions, reinforcing its critical role in the sensor suite of next-generation autonomous vehicles.

5.2 Limitations and Future Work

While the network simulation framework provided unparalleled control over experimental variables and yielded pristine ground truth annotations, it is imperative to acknowledge the inherent limitations of synthetic data. The phenomenon known as the simulation-to-reality gap remains a significant challenge. Despite the advanced phenomenological modeling employed in this study, virtual environments cannot perfectly replicate the infinite variability, complex multipath propagation, and unpredictable electromagnetic interference present in actual urban driving conditions. The neural network models trained exclusively on simulated data may experience a degradation in performance when deployed on physical hardware in the real world.

Future research directions must focus on bridging this domain gap through the application of advanced transfer learning and domain adaptation techniques. Exploring methodologies such as generative adversarial networks to inject realistic noise profiles and clutter distributions into the simulated data could significantly enhance the generalization capabilities of the recognition algorithms. Additionally, subsequent studies should investigate the early fusion of raw radar signal features with complementary data streams from visual cameras and light detection and ranging sensors, capitalizing on the distinct strengths of each modality to create a truly synergistic, fault-tolerant perception architecture for fully autonomous navigation.

References

1. Liu, C.; Zhao, J.; Li, J.; Wang, D.; Yu, F.R. UAV Aided Integrated Sensing, Communication and Computing: Optimization via Federated Learning. IEEE Trans. Veh. Technol. 2025, 75, 6045–6058.

2. Braess, D.; Nagurney, A.; Wakolbinger, T. On a paradox of traffic planning. Transp. Sci. 2005, 39, 446–450.

3. Braess, D. Über ein Paradoxon aus der Verkehrsplanung. Unternehmensforschung 1968, 12, 258–268.

4. Liu, X.; Xu, Y.; Jia, L.; Wu, Q.; Anpalagan, A. Anti-Jamming Communications Using Spectrum Waterfall: A Deep Reinforcement Learning Approach. IEEE Commun. Lett. 2018, 22, 998–1001.

5. Chen, X.; Ding, H.; Ma, Y.; Li, X.; An, J.; Fang, Y. A Dual-Tier Policy-Oriented Anti-Jamming Scheme Based on Deep Reinforcement Learning. IEEE Trans. Wirel. Commun. 2026, 25, 10652–10668.

6. Taubert, N., Bruns, A., Lenke, C., & Stone, G. (2021). Waiving article processing charges for least developed countries. A brick stone of a large-scale open access transformation. UKSG Insights, 34(1), 1.

7. Roughgarden, T.; Tardos, É. How bad is selfish routing? J. ACM 2002, 49, 236–259.

8. Pandya, S.; Srivastava, G.; Jhaveri, R.; Babu, M.R.; Bhattacharya, S.; Maddikunta, P.K.R.; Mastorakis, S.; Piran, M.J.; Gadekallu, T.R. Federated Learning for Smart Cities: A Comprehensive Survey. Sustain. Energy Technol. Assess. 2023, 55, 102987.

9. Silva, P.; Teles, F.; Ferreira, J. Intermunicipal cooperation: The quest for governance capacity? Int. Rev. Adm. Sci. 2018, 84, 619–638.

10. Han, S.; Dastani, M.; Wang, S. Sparse Communication in Multi-Agent Deep Reinforcement Learning. Neurocomputing 2025, 625, 129344.

11. Zhu, T.; Shi, X.; Xu, X.; Gui, J.; Cao, J. HyperComm: Hypergraph-Based Communication in Multi-Agent Reinforcement Learning. Neural Netw.

2024, 178, 106432.

12. Feng, Z.; Huang, M.; Wu, Y.; Wu, D.; Cao, J.; Korovin, I.; Gorbachev, S.; Gorbacheva, N. Approximating Nash Equilibrium for Anti-UAV Jamming Markov Game Using a Novel Event-Triggered Multi-Agent Reinforcement Learning. Neural Netw. 2023, 161, 330–342.

13. Matolak, D.W. Channel modelling for vehicle-to-vehicle communications. IEEE Commun. Mag. 2008, 46, 76–83.

14. Pennebaker, J. W., Boyd, R. L., Jordan, K., & Blackburn, K. (2015). The development and psychometric properties of LIWC2015. University of Texas at Austin.

15. Ben-Elia, E.; Shiftan, Y. Which road do I take? A learning-based approach to route choice. Transp. Res. Part A Policy Pract. 2010, 44, 110–120.