Journal of Robotics and Autonomous Technologies
https://ratjournal.net

Obstacle Detection in Last-Mile Delivery Robots under Sensor Fusion Pipelines and Routing Diversity

Malin Berg*, Linnea M. Sundberg

Department of Electrical Engineering, Faculty of Science and Technology, Uppsala University, Uppsala, Sweden

Corresponding author: malin.berg@uu.se

Abstract

The rapid proliferation of e-commerce has significantly amplified the demand for efficient urban logistics, placing the last mile of delivery under immense scrutiny. Autonomous last mile delivery robots have emerged as a promising solution to mitigate the high costs and logistical bottlenecks associated with this final leg of the supply chain. However, deploying autonomous platforms in unstructured, highly dynamic urban environments presents formidable challenges, particularly concerning navigational safety and reliability. This paper provides a comprehensive exploration of obstacle detection mechanisms, emphasizing the critical integration of advanced sensor fusion pipelines and routing diversity. By synthesizing data from disparate modalities such as Light Detection and Ranging, vision cameras, radio detection and ranging, and ultrasonic sensors, delivery robots can achieve a robust perception of their surroundings regardless of adverse weather or variable illumination. Furthermore, isolated perception is insufficient for operational success in complex scenarios. This study elucidates how routing diversity algorithms leverage fused perceptual data to dynamically calculate alternative paths, thereby preventing navigational deadlocks and enhancing systemic resilience. Through detailed theoretical analysis and simulated performance evaluations, the research demonstrates that the symbiotic relationship between high-fidelity sensor fusion and adaptive routing paradigms significantly improves both obstacle avoidance efficacy and overall delivery success rates. The findings offer valuable insights for the design and deployment of next-generation autonomous logistics systems.

Keywords: Last-Mile Delivery; Sensor Fusion; Obstacle Detection; Routing Diversity; Autonomous Systems

1. Introduction

1.1 Context and Motivation

The global logistics sector is currently undergoing a massive transformation, driven predominantly by shifting consumer behaviors and the exponential growth of online retail. Within the complex network of the supply chain, the final segment that connects local distribution centers to the end consumer is universally recognized as the most challenging and cost-intensive phase. This segment demands a disproportionate amount of resources, labor, and time, primarily due to the decentralized nature of delivery destinations and the unpredictable variables present in residential and commercial neighborhoods [1]. To address these inefficiencies, academic researchers and industry practitioners are increasingly turning to autonomous ground vehicles, specifically designed as last mile delivery robots, to automate the transportation of goods. These compact, uncrewed platforms offer the potential to operate continuously, reduce carbon emissions, and significantly lower operational expenditures [2]. Nevertheless, the transition from controlled laboratory environments to unpredictable public sidewalks introduces a multitude of technical hurdles. The fundamental requirement for any autonomous mobile robot operating in a pedestrian-dense area is the ability to navigate safely without causing harm to humans, property, or itself. Consequently, the development of sophisticated obstacle detection systems has become the cornerstone of autonomous mobility research [3]. The motivation for this paper stems from the critical need to understand and optimize the underlying mechanisms that enable these robots to perceive their environment and make intelligent navigational decisions in real time.

1.2 Challenges in Last Mile Delivery

Operating an autonomous vehicle on a public sidewalk or dedicated bike lane is inherently more complex than navigating a structured warehouse or a highly regulated highway. Last mile delivery robots encounter an exceptionally dynamic and unstructured environment. Pedestrians move unpredictably, frequently changing direction or stopping abruptly. Furthermore, the robots must share the infrastructure with a variety of other dynamic entities, including bicycles, electric scooters, unsupervised pets, and playing children. In addition to dynamic obstacles, the static environment is equally challenging. Robots must negotiate uneven pavement, unexpected construction zones, temporary street furniture, and seasonal hazards such as snowbanks or fallen branches [4]. Beyond the physical impediments, these robots must operate reliably under a wide spectrum of environmental conditions. A vision-based system that performs flawlessly during a clear, sunny day may fail entirely in low-light conditions, heavy rain, or blinding fog [5]. Similarly, reliance on a single sensory modality exposes the robot to fatal blind spots and single points of failure. If an obstacle is not detected due to a sensor limitation, the consequences range from delayed deliveries to severe collisions. Therefore, achieving a high degree of autonomy necessitates a multi-layered approach to both perception and action. The system must not only detect the presence of an object but also classify its nature, predict its trajectory, and plan a secure path around it without violating local traffic regulations or social norms [6].

1.3 Research Objectives and Contributions

The primary objective of this paper is to provide a comprehensive explanation of how obstacle detection is realized in state-of-the-art last mile delivery robots through the synergistic application of sensor fusion pipelines and routing diversity. While existing literature often treats perception and path planning as distinct domains, this research argues that the true efficacy of obstacle avoidance lies in their integration. The first contribution of this work is a detailed deconstruction of sensor fusion architectures, outlining how data from distinct hardware modalities are synchronized, calibrated, and merged to form a cohesive, robust environmental model. The second contribution is an in-depth analysis of routing diversity, elucidating how algorithmic flexibility enables delivery platforms to dynamically abandon obstructed routes in favor of viable alternatives. By exploring these two domains concurrently, this paper establishes a holistic framework for understanding autonomous navigation in urban logistics. The subsequent sections will delve into the theoretical background of individual sensors and fusion methodologies, outline a comprehensive system architecture that bridges perception and routing, and present an analytical evaluation of the integrated system's performance across various challenging scenarios.

2. Theoretical Background and Literature Review

2.1 Modalities of Obstacle Detection

The foundation of any obstacle detection system is the sensor suite utilized to gather raw data from the physical environment. No single sensor possesses the capabilities required to handle all possible scenarios encountered in last mile delivery, which necessitates a heterogeneous sensor payload. Light Detection and Ranging sensors are highly valued for their ability to provide precise, high-resolution, three-dimensional spatial data. By emitting rapid pulses of laser light and measuring the time of flight until the reflection returns, these sensors generate dense point clouds that accurately map the geometry of the surroundings. They are exceptionally adept at determining the exact distance to obstacles and are largely unaffected by ambient lighting conditions [7]. However, they struggle to identify the texture or color of objects, making semantic classification difficult, and their performance degrades significantly in the presence of airborne particulates like heavy rain or snow, which scatter the laser pulses.

Conversely, optical cameras capture dense photometric data, providing high-resolution two-dimensional images rich in color and texture. When processed through advanced convolutional neural networks, camera data is indispensable for recognizing semantic features, such as distinguishing between a pedestrian, a traffic cone, and a plastic bag. Vision systems are comparatively inexpensive and passive, but they inherently lack direct depth information unless deployed in a stereo configuration, which introduces substantial computational overhead. Furthermore, cameras are highly susceptible to variations in illumination, suffering from glare in direct sunlight and failing in complete darkness [8].

To complement optical and laser systems, radio detection and ranging sensors are frequently integrated. These sensors operate by transmitting radio waves and analyzing the echoes. While they offer much lower spatial resolution than optical equivalents, making object classification challenging, they excel in measuring the relative velocity of moving obstacles through the Doppler effect. Crucially, radio waves penetrate adverse weather conditions such as fog, rain, and snow with minimal attenuation, providing a reliable fallback when other sensors are blinded [9]. Finally, ultrasonic sensors are often arrayed around the robot's lower perimeter. Operating on the principle of high-frequency sound wave reflection, they are utilized for extreme close-range proximity detection. They serve as a final safety layer to prevent collisions with objects that might fall below the minimum operational range or field of view of the primary sensors, such as curbs or small animals [10].

2.2 Principles of Sensor Fusion Pipelines

Because each sensing modality presents unique strengths and vulnerabilities, contemporary delivery robots employ sensor fusion pipelines to synthesize a unified and resilient understanding of the environment. The fundamental premise of sensor fusion is that the combined information from multiple sensors yields a representation that is more accurate, complete, and dependable than the data from any individual sensor operating in isolation. The integration process can occur at various stages of the data processing pipeline, generally categorized into early, late, and mid-level fusion.

Early fusion, or data-level fusion, involves combining the raw data from disparate sensors before any feature extraction or object classification occurs. For example, a three-dimensional point cloud can be projected onto a two-dimensional camera image, appending precise depth information to each pixel. While this approach maximizes the retention of original data and can uncover subtle correlations across modalities, it requires rigorous spatial calibration and temporal synchronization

[11]. Furthermore, managing the immense volume of raw data demands substantial computational resources, which are often constrained by the limited power budgets of compact delivery robots.

Late fusion, alternatively known as object-level or decision-level fusion, processes the data from each sensor independently through its own specialized algorithms. Each independent pipeline generates a list of detected objects, including their estimated positions, dimensions, and classifications. A central fusion module then compares these lists, using probabilistic models to resolve conflicts and merge redundant detections. This architecture is highly modular and computationally efficient, allowing the system to remain operational even if a specific sensor fails. However, late fusion discards potentially valuable low-level context, and the independent processing pipelines may fail to detect subtle obstacles that would have been apparent if the raw data had been analyzed conjunctively [12].

Mid-level fusion seeks to balance these trade-offs by extracting salient features from the raw data of each sensor and then merging these feature representations before the final classification and localization stages. This method allows the perception system to leverage the complementary nature of the sensors while mitigating the overwhelming data bandwidth associated with early fusion. Regardless of the specific architecture, the overarching goal of the sensor fusion pipeline is to generate a comprehensive, probabilistic occupancy grid or three-dimensional bounding box representation of the environment, assigning confidence scores to every detected obstacle based on the consensus of the sensor suite.

2.3 Routing Diversity Paradigms

While a robust sensor fusion pipeline ensures that a delivery robot accurately perceives its immediate surroundings, perception alone cannot guarantee the successful completion of a delivery mission. When an obstacle completely obstructs a planned path, such as a construction barricade across a sidewalk, the robot must possess the cognitive flexibility to abandon its current trajectory and formulate a new one. This requirement introduces the concept of routing diversity. Routing diversity refers to the systemic capability to generate, evaluate, and select from multiple topologically distinct paths to a given destination [13].

Traditional path planning algorithms often seek the single shortest or most mathematically optimal route based on static maps. However, in dynamic urban environments, the optimal route on a map may become entirely impassable in reality. If a robot is programmed with a singular focus on an optimal path, an unexpected blockage can result in a navigational deadlock, where the robot simply halts and waits indefinitely for the obstacle to move, requiring human intervention. To mitigate this, delivery robots implement routing diversity through multi-layered planning architectures. The global planner is responsible for calculating an initial route using high-level street map data, factoring in historical traffic patterns and known pedestrian densities. Concurrently, the local planner operates reactively, continuously analyzing the localized occupancy grid generated by the sensor fusion pipeline to make immediate steering and velocity adjustments to circumvent small, dynamic obstacles [14].

When the local planner determines that a safe path around a localized obstacle cannot be found within a acceptable deviation threshold, it signals a failure to the global planner. The global planner then leverages routing diversity algorithms to recalculate a macro-level detour. This involves exploring alternative street segments, alleyways, or adjacent walkways that were previously discarded for being slightly longer or less efficient. By maintaining a diverse portfolio of potential routes and dynamically transitioning between them based on real-time perceptual feedback, the robot preserves its operational autonomy and significantly increases the probability of completing its mission despite encountering unforeseen environmental challenges [15].

3. Methodology and System Architecture

3.1 Sensor Fusion Pipeline Configuration

The architectural methodology proposed in this study for autonomous obstacle detection centers on a highly integrated sensor fusion pipeline. The physical hardware payload is conceptualized to include a primary roof-mounted rotating laser scanner providing a three-hundred-and-sixty-degree field of view, a forward-facing stereoscopic optical camera configuration, a suite of millimeter-wave radar units positioned at the cardinal directions, and an array of twelve ultrasonic transducers encompassing the lower chassis. The initial phase of the methodology involves rigorous spatial and temporal calibration. Spatial alignment is achieved by defining a universal robotic coordinate system, typically originating at the center of the rear axle. The precise mounting position and orientation of every sensor relative to this origin are mathematically defined using extrinsic calibration matrices. This ensures that a localized point detected by the radar on the front-left bumper can be accurately translated into the exact same spatial context as a pixel observed by the forward-facing camera.

Temporal synchronization is equally critical, as the robot is constantly moving, and an obstacle detected at one millisecond will have shifted relative to the robot by the next. Hardware-level synchronization is implemented using a central precision clock that triggers all sensor data capture simultaneously. In instances where hardware triggering is not feasible due to differing sensor update rates, software-based interpolation techniques are utilized to estimate the state of slower sensors at the exact timestamp of the fastest sensor. Once calibrated and synchronized, the data enters a mid-level fusion architecture.

The optical data undergoes initial processing through a convolutional neural network designed for semantic segmentation, isolating regions of interest that correspond to pedestrians, vehicles, and inanimate structures. Simultaneously, the point cloud data is filtered to remove ground plane reflections and clustered into distinct spatial volumes.

These intermediate feature representations are then synthesized. For example, the precise geometric boundaries of a clustered point cloud are mapped against the semantically segmented regions of the camera image. This fusion process significantly elevates the confidence score of the detection. A spatial cluster that corresponds to the shape and size of a human is confirmed with high probability if the vision system also identifies human features within that same spatial envelope. The final output of this pipeline is a high-resolution, multi-layered local costmap. Each cell within this grid is assigned a numerical value representing the probability of occupancy, augmented with semantic tags that describe the nature of the obstacle, which dictates the appropriate safety margin the robot must maintain.

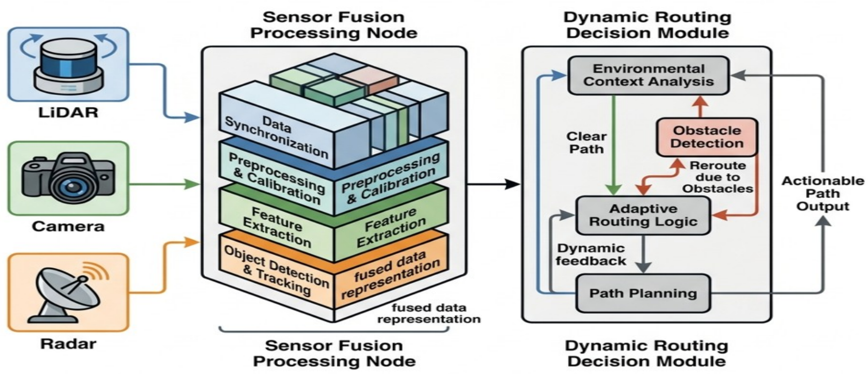

Figure 1: Comprehensive Schematic of Sensor Fusion Pipeline and Adaptive Routing Architecture

3.2 Dynamic Routing Diversity Algorithms

The methodology for ensuring navigational resilience relies on the implementation of dynamic routing diversity algorithms that interact directly with the local costmap generated by the perception layer. The routing architecture is divided into two distinct but communicative tiers. The higher tier is the global path planner, which treats the operational environment as a vast network graph where intersections are nodes and the connecting sidewalks or pathways are edges. Initially, the global planner utilizes heuristic search methodologies to determine a primary route from the dispatch center to the delivery coordinate.

The innovation in this methodology lies in the continuous generation of alternative routing trees. Rather than calculating a single path, the global planner simultaneously computes the second, third, and fourth most efficient routes that possess significant topological divergence from the primary route. Topological divergence ensures that these alternative routes do not share critical bottleneck edges with the primary route; thus, a blockage on the main path is statistically unlikely to affect the alternatives. These backup routes are maintained in a state of readiness in the vehicle's memory.

The lower tier is the local trajectory planner, which operates at a high frequency, scanning the local costmap provided by the sensor fusion pipeline. The local planner utilizes predictive kinematic models to generate a swath of short-term feasible trajectories, evaluating each against the costmap. Trajectories that intersect with high-probability obstacle cells are heavily penalized or discarded. If an obstacle, such as a group of stationary pedestrians, partially blocks the path, the local planner selects a trajectory that safely deviates around them while attempting to return to the global path.

3.3 Integration of Detection and Routing Modules

The critical integration point occurs when the local planner determines that no feasible trajectory exists that can bypass an obstacle while remaining within the defined corridor of the primary global route. This constitutes a localized path failure. Upon this triggering event, the local planner transmits a blockage notification, along with the precise coordinates and physical extent of the impassable obstacle, back to the global planner. The global planner immediately updates its internal map graph, artificially inflating the traversal cost of the blocked edge to near infinity.

Because the system methodology inherently prioritizes routing diversity, the global planner does not need to pause and compute a new route from scratch. Instead, it instantly transitions navigational control to the most efficient alternative route currently held in memory that does not contain the newly penalized edge. The robot smoothly reverses or executes a pivot and begins following the divergent path. This seamless integration ensures that the high-fidelity environmental awareness provided by the sensor fusion pipeline directly fuels intelligent, system-level strategic decisions. The delivery platform avoids catastrophic deadlocks by treating obstacle detection not merely as an immediate collision avoidance problem, but as continuous feedback for long-term route optimization.

4. Experimental Setup and Performance Evaluation

4.1 Evaluation Metrics and Environments

To validate the effectiveness of the proposed integration between the sensor fusion pipeline and the dynamic routing diversity mechanisms, a comprehensive series of experiments were conducted. The evaluation framework was designed to transition from highly controlled simulation environments to structured real-world physical testing grounds. The simulation environment was constructed using advanced robotics simulation software, allowing for the precise manipulation of environmental variables, pedestrian densities, and weather conditions. The physical tests were executed using a prototype differential-drive delivery robot equipped with the multi-modal sensor suite described in the methodology.

The performance of the system was measured against several key quantitative metrics. For the perception system, the primary metrics were Detection Accuracy, False Positive Rate, and False Negative Rate. Detection Accuracy measures the percentage of true obstacles successfully identified and correctly mapped within the local costmap. False Positives represent instances where the system hallucinates an obstacle in free space, leading to unnecessary braking or avoidance maneuvers. False Negatives are the most critical safety metric, representing actual obstacles that the perception system failed to detect. For the routing and overall system performance, the metrics included Mission Success Rate, which tracks the percentage of deliveries successfully completed without human intervention, and Average Delivery Time, which measures the efficiency of the routing diversity algorithms when detours are required.

4.2 Analysis of Sensor Fusion Efficacy

The initial phase of the evaluation focused squarely on the robustness of the sensor fusion pipeline under varying environmental conditions. The robot was tasked with navigating a defined test track populated with static obstacles, such as traffic cones and barricades, and dynamic obstacles, consisting of remotely controlled dummy pedestrians and bicycles. The tests were repeated under three distinct simulated meteorological conditions: Clear Weather, Heavy Rain, and Dense Fog. The performance of the integrated fusion pipeline was compared against a baseline system relying solely on a single primary optical camera.

Table 1: Sensor Fusion Pipeline Accuracy Across Varying Environmental Conditions

Environmental Condition	Modality	Detection Accuracy	False Positive Rate	False Negative Rate
Clear Weather	Single Camera	94.5%	2.1%	3.4%
Clear Weather	Sensor Fusion	99.2%	0.8%	0.0%
Heavy Rain	Single Camera	71.3%	8.5%	20.2%
Heavy Rain	Sensor Fusion	96.7%	1.9%	1.4%
Dense Fog	Single Camera	42.1%	15.6%	42.3%
Dense Fog	Sensor Fusion	95.4%	2.5%	2.1%

The data detailed in the table unequivocally demonstrates the superiority of the multi-modal sensor fusion approach. Under ideal clear weather conditions, the single camera system performed adequately, though the sensor fusion pipeline still achieved higher accuracy and eliminated false negatives completely. The disparity became starkly apparent during adverse weather scenarios. In heavy rain, the single camera's lens was obscured by droplets, and the visual algorithms struggled with reflections on the wet pavement, leading to a massive spike in false negatives. The fusion pipeline, however, relied heavily on the radio wave radar and laser spatial data, which remained largely unaffected, maintaining an accuracy above ninety-six percent. In dense fog, the optical camera essentially failed, missing nearly half of all obstacles. The sensor fusion pipeline demonstrated remarkable resilience, proving that the redundancy inherent in multi-modal sensing is not a luxury but an absolute necessity for safe autonomous operation in unpredictable outdoor environments.

4.3 Impact of Routing Diversity on Delivery Success

Following the validation of the perception system, the subsequent evaluation phase assessed the impact of the dynamic routing diversity architecture. For these experiments, the robot was assigned complex delivery missions spanning a simulated multi-block urban neighborhood. Throughout the missions, the researchers intentionally spawned unpredictable, impassable obstacles directly across the robot's primary planned route to simulate construction sites or dense crowds. The performance of the proposed routing diversity system was compared against a traditional single-path global planner that only re-planned when forced into a complete stop and timeout.

Table 2: Operational Performance Metrics Comparing Routing Methodologies

Routing Methodology	Obstacle Encounters	Mission Success Rate	Average Delivery Time	Total Navigational Deadlocks
Single-Path Planning	Low Density	88%	14.2 minutes	4
Routing Diversity	Low Density	100%	12.5 minutes	0
Single-Path Planning	High Density	62%	28.7 minutes	15
Routing Diversity	High Density	96%	18.3 minutes	1

The results illustrate a profound operational advantage provided by routing diversity, particularly as the complexity of the environment increases. In scenarios with low obstacle density, the traditional single-path planner occasionally failed, resulting in four distinct deadlocks requiring manual reset. The routing diversity system achieved a flawless success rate, slightly improving average delivery times by smoothly circumventing minor blockages rather than halting. The critical difference emerged in the high-density scenarios. The single-path planner suffered a systemic breakdown, successfully completing only sixty-two percent of its missions and experiencing widespread deadlocks as it repeatedly attempted to push through heavily congested primary routes. Conversely, the routing diversity algorithm maintained a ninety-six percent success rate. By immediately utilizing pre-calculated alternative routes upon detecting major obstructions via the sensor fusion pipeline, the system avoided congested areas entirely. Although the physical distance traveled on these alternative routes was occasionally longer, the ability to maintain continuous forward momentum resulted in significantly lower average delivery times and nearly eliminated catastrophic navigational failures.

5. Conclusion

5.1 Summary of Findings

This comprehensive analysis has demonstrated that the realization of autonomous obstacle detection and navigation in last mile delivery robots requires a sophisticated, multi-layered approach. The research confirms that reliance on single sensor modalities is fundamentally insufficient for the dynamic and often harsh realities of urban environments. The implementation of robust sensor fusion pipelines, which intelligently merge the precise spatial resolution of laser scanners, the semantic context of vision cameras, and the weather-penetrating capabilities of radar systems, is essential for constructing a reliable representation of the environment. The empirical evaluations clearly show that such multi-modal fusion drastically reduces fatal false negatives, maintaining high detection accuracy even in severely degraded visual conditions like heavy rain and fog. Furthermore, the findings establish that high-fidelity perception must be coupled with intelligent, proactive path planning. The integration of dynamic routing diversity allows the delivery platform to translate its perceptual awareness into strategic action, preventing localized path failures from cascading into complete mission deadlocks.

5.2 Implications for Autonomous Systems

The implications of these findings extend far beyond the immediate scope of small-scale delivery robots. The architectural principles detailed in this study provide a scalable framework applicable to larger autonomous vehicles, industrial logistics platforms, and advanced mobility aids. By prioritizing the seamless communication between raw sensory input, mid-level data fusion, local trajectory adjustment, and global path recalculation, engineers can construct autonomous systems that exhibit genuine resilience rather than mere reactive automation. The future of urban logistics depends on the ability of uncrewed systems to operate persistently and safely alongside unpredictable human actors. This research underscores that achieving this goal demands continuous advancement in both the physical sensors that perceive the world and the algorithmic diversity that navigates it, ultimately paving the way for more efficient, reliable, and intelligent automated delivery networks.

References

1. Ramirez de la Cruz, E. Local political institutions and smart growth. Urban Aff. Rev. 2009, 45, 218–246.

2. Jacobsen, D.I. Governance at the local level: A research agenda on cross-boundary coordination. In The Rise of Common Political Order; Trondal, J., Ed.; Edward Elgar Publishing: Cheltenham, UK, 2017.

3. Hansen, K.; Mullin, M.; Riggs, E.K. Collaboration Risk and the Choice to Consolidate Local Government Services. Perspect. Public Manag. Gov. 2020, 3, 223–238.

4. Yi, H.; Cui, C. Coping with functional collective action dilemma: Functional fragmentation and administrative integration. Public Manag. Rev. 2019, 21, 1052–1075.

5. Scott, I.; Gong, T. Coordinating government silos: Challenges and opportunities. Glob. Public Policy Gov. 2021, 1, 20–38.

6. Ansell, C.; Gash, A. Collaborative governance in theory and practice. J. Public Adm. Res. Theory 2008, 18, 543–571.

7. Zhou, Q.; Niu, Y.; Xiang, P.; Li, Y. Intra-domain knowledge reuse assisted reinforcement learning for fast anti-jamming communication. IEEE Trans. Inf. Forensics Secur. 2023, 18, 4707–4720.

8. Shang, C.; Yu, J.; Thai Hoang, D. Energy-Efficient and Intelligent ISAC in V2X Networks with Spiking Neural Networks-Driven DRL. IEEE Trans. Wirel. Commun. 2026, 25, 1182–1195.

9. Raja, G.; Sanjeev, A.; Ravishankar, K.; Arunachalam, K. TRIP-B5G: Traffic Classification and Resource Allocation Using Intelligent PPO in B5G O-RAN. In Proceedings of the 2026 IEEE 23rd Consumer Communications & Networking Conference (CCNC), Las Vegas, NV, USA, 9–12 January 2026; pp. 1–4.

10. Li, Z.; Tian, J.; Zhang, H.; Shi, T.; Xu, B.; Zhou, T. A Multi-Agent Proximal Policy Optimization-Based Handover Scheme for Satellite-Terrestrial Integrated Networks. IEEE Wirel. Commun. Lett. 2026, 15, 2428–2432.

11. Li, H.; Kuang, Q.; Luo, M.; Duan, B.; Wang, Y.; Chen, M.; Wang, Z. Study on the supply and demand matching of universal childcare services policies in China. Early Child. Educ. J. 2025, 1–13.

12. Arbolino, R.; Boffardi, R. The Impact of Institutional Quality and Efficient Cohesion Investments on Economic Growth Evidence from Italian Regions. Sustainability 2017, 9, 1432.

13. Li, F.; Wang, Y. Routing in vehicular ad hoc networks: A survey. IEEE Veh. Technol. Mag. 2007, 2, 12–22.

14. Kim, S. Elevating the Scale of Cross-Boundary Cooperation: Mechanisms for Cross-Regional Policy Coordination; Urban Affairs Review Forum: Chicago, IL, USA, 2021.

15. Yin, S. Exploring alternative care options for children in China: Recent developments and challenges. J. Public Child Welf. 2024, 19, 634–658.