

Article **Open Access**

Multi-Domain Adaptation for Autonomous Driving Perception under Diverse Weather

Jinghui Tan ^{1,*}



¹ School of Mathematics and Computational Science, Wuyi University, Guangdong, 529000, China

* Correspondence: Jinghui Tan, School of Mathematics and Computational Science, Wuyi University, Guangdong, 529000, China

Abstract: Autonomous driving perception systems are confronted with substantial robustness challenges under diverse weather conditions, where sensor data distortion caused by rain, fog, snow, or intense illumination often leads to degraded performance in critical tasks such as object detection and semantic segmentation. Existing approaches predominantly depend on single-domain models trained under ideal environmental conditions, which suffer from poor generalization across weather domains due to inherent domain shifts. This study explores the application of multi-domain adaptation techniques to enhance perception stability by integrating heterogeneous sensor data, including RGB (Red, Green, Blue) images, LiDAR (Light Detection and Ranging) point clouds, and thermal imaging, while leveraging cross-domain feature alignment mechanisms. The proposed framework employs domain-specific encoders combined with adversarial learning to mitigate weather-induced domain gaps, alongside a multi-task learning objective that simultaneously optimizes perception accuracy and domain invariance. Experimental validation demonstrates that the framework achieves superior performance compared to conventional single-domain and shallow adaptation models, with interpretability analyses revealing key weather-robust features such as thermal edge consistency and LiDAR (Light Detection and Ranging) point density patterns. Its ability to adapt to unseen weather conditions could enable reliable autonomous driving in complex real-world environments and reduce weather-related accidents. By bridging domain adaptation theory with automotive perception requirements, this work advances the translation of robust AI (Artificial Intelligence)-driven systems into practical autonomous driving applications.

Keywords: multi-domain adaptation; autonomous driving perception; diverse weather conditions; domain shift; robust sensor fusion; adversarial learning; cross-modal feature alignment

Received: 01 June 2025

Revised: 09 June 2025

Accepted: 17 June 2025

Published: 06 August 2025



Copyright: © 2025 by the authors. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Autonomous driving perception systems serve as the foundational component for environmental understanding, enabling vehicles to interpret surrounding conditions and support critical decision-making processes such as obstacle avoidance, lane keeping, and trajectory planning. However, real-world driving environments are inherently dynamic, with diverse weather conditions, with diverse weather conditions including rain, fog, snow, and intense sunlight, posing significant challenges to perception robustness. These weather-related perturbations introduce substantial distortions in sensor data: rain droplets scatter light and create motion blur in RGB images, fog reduces atmospheric visibility and degrades depth estimation accuracy, snowflakes occlude foreground objects and introduce noise in LiDAR point clouds, while strong sunlight causes glare that washes out

critical features in visual data. Such distortions often lead to performance degradation in core perception tasks, including object detection, semantic segmentation, and pedestrian tracking, thereby increasing the risk of accidents in adverse weather [1].

Current autonomous driving perception methods predominantly rely on single-domain models trained under ideal or controlled weather conditions, which exhibit poor generalization across weather domains due to the "domain shift" phenomenon, "which is the distribution mismatch between training (source domain) and deployment (target domain) data. For instance, a model optimized for clear weather may experience a 30 to 50 percent drop in detection accuracy when deployed in heavy rain, as it fails to recognize weather-induced variations in object appearance and sensor characteristics [2]. Existing solutions to address this issue fall into three categories, each with inherent limitations: data augmentation techniques, which simulate weather effects (e.g., adding rain streaks) during training, often fail to capture the complexity of real-world weather patterns and thus limit generalization; single-domain fine-tuning, which retrains models on weather-specific datasets, requires massive annotated data for each weather type and lacks cross-domain adaptability; and shallow domain adaptation methods (e.g., CORAL, DANN), which align low-level features but overlook the hierarchical nature of perception tasks, resulting in suboptimal performance when facing multiple weather domains [3].

The emergence of multi-domain adaptation, a subfield of transfer learning, offers a promising direction to enhance perception robustness by enabling models to learn invariant features across multiple source domains (e.g., different weather conditions) and generalize to unseen target domains. By aligning feature distributions across weather domains while preserving task-specific information, multi-domain adaptation can mitigate domain shift and improve performance in diverse environments. However, its application to autonomous driving perception under diverse weather conditions remains underexplored, with three key challenges: first, heterogeneous sensor data fusion, since RGB images, LiDAR point clouds, and thermal imaging each have distinct data structures and weather sensitivities, they require specialized alignment strategies; second, dynamic domain shifts, as weather intensity (e.g., light vs. heavy rain) varies continuously, demanding adaptive mechanisms rather than static feature alignment; third, real-time inference constraints because automotive systems require perception results within milliseconds to support timely decision-making [4].

This study proposes a multi-domain adaptation framework for autonomous driving perception under diverse weather conditions, which integrates heterogeneous sensor data (RGB images, LiDAR point clouds, and thermal imaging) through domain-specific encoders and cross-domain adversarial alignment mechanisms. The framework employs a multi-task learning objective that simultaneously optimizes perception accuracy (for object detection and semantic segmentation) and domain invariance (to minimize weather-induced distribution gaps), with novel regularization techniques to handle partial sensor failure (e.g., LiDAR malfunction in heavy snow). The goal is to enable robust perception across unseen weather conditions, reduce dependency on large weather-specific datasets, and meet real-time performance requirements for practical deployment.

The remainder of this paper is structured as follows: Section 2 reviews related work on weather-aware perception and multi-domain adaptation; Section 3 details the proposed framework, including data preprocessing, model architecture, and training protocols; Section 4 presents experimental validation, including comparative results and ablation studies, Section 5 discusses the framework's contributions, limitations and practical implications, and Section 6 concludes with future research directions. By bridging multi-domain adaptation theory with automotive perception needs, this work aims to advance the reliability of autonomous driving systems in complex real-world environments.

2. Related Works

The development of weather-robust autonomous driving perception has evolved through four distinct methodological phases, as illustrated in Figure 1, each addressing increasingly complex challenges of domain shift under diverse weather conditions.




Figure 1. the Evolution of Domain Adaptation Methods.

2.1. Phase 1: Single-Modality Rule-Based Methods (Before 2018)

Early efforts relied on handcrafted features and weather-specific heuristics, such as dehazing algorithms for foggy images or threshold-based Light Detection and Ranging (LiDAR) denoising. These methods achieved limited robustness: for example, a histogram equalization-based approach for foggy object detection reported 62 to 68 percent accuracy under light fog but dropped to below 50 percent in dense fog [5]. Such techniques failed to generalize across weather types due to their reliance on fixed rules rather than learned patterns.

2.2. Phase 2: Single-Domain Deep Learning (2018 to 2021)

The adoption of deep neural networks (DNNs) improved performance under ideal conditions, with models like Faster Region-Based Convolutional Neural Network (Faster R-CNN) and PointNet++ achieving over 90 percent mean Average Precision (mAP) in clear weather. However, their single-domain training paradigm led to severe degradation in adverse weather: a study on the nuScenes dataset showed that a state-of-the-art Convolutional Neural Network (CNN) for object detection experienced a 43 percent accuracy drop in heavy rain and 58 percent in snow. Light Detection and Ranging (LiDAR)-based models faced similar issues, with point cloud sparsity in fog reducing semantic segmentation mean Intersection over Union (mIoU) by 35 percent [6].

2.3. Phase 3: Single-Pair Domain Adaptation (2021 to 2023)

Initial domain adaptation methods focused on aligning two domains (e.g., clear to rain) using adversarial learning or metric-based alignment. For instance, a Domain-Adversarial Neural Network (DANN) reduced the performance gap between clear and rainy conditions by 18 percent but failed to generalize to fog or snow [7]. Similarly, Correlation Alignment (CORAL) improved cross-weather segmentation by 12 percent but struggled with non-linear domain shifts [8]. These approaches, as summarized in Table 1, lacked scalability to multiple weather domains.

Table 1. Performance Comparison of Single-Pair Domain Adaptation Methods Under Weather Shifts.

Method	Source Domain	Target Domain	mAP (Target)	Domain Gap Reduction (%)
DANN	Clear	Rain	68.2%	18.3
CORAL	Clear	Fog	61.5%	12.1
Contrastive Adaptation	Clear	Snow	65.7%	15.6

2.4. Phase 4: Multi-Domain Adaptation (2023 to the Present)

Recent work has extended to multi-domain settings, leveraging contrastive learning or meta-learning to align features across three or more weather types. A multi-source domain adaptation framework using shared-private feature disentanglement achieved 72 percent mAP across rain, fog, and snow, outperforming single-pair methods by 15 percent.

However, these studies predominantly focus on Red, Green, Blue (RGB) images, neglecting LiDAR or thermal data, and rarely address dynamic weather intensity variations (e.g., light vs. heavy rain).

Figure 2. shows the distribution of recent studies (2020 to 2024) across these phases, revealing that only 19 percent of publications focus on multi-domain adaptation for weather-robust perception, and only 8 percent integrate heterogeneous sensors. This gap underscores the need for frameworks that unify multi-domain adaptation with cross-modal sensor fusion to address the full complexity of real-world weather challenges.




Figure 2. the Distribution of Recent Studies.

3. Methodology

The proposed multi-domain adaptation framework for weather-robust autonomous driving perception integrates heterogeneous sensor data through domain-specific feature extraction, cross-domain alignment, and multi-task optimization. Figure 3 illustrates the end-to-end pipeline, which addresses three core challenges: heterogeneous sensor fusion, dynamic weather-induced domain shifts, and real-time inference constraints.




Figure 3. Framework Architecture.

3.1. Data Acquisition and Preprocessing

Datasets: The framework is validated on a combined dataset comprising three sources: (1) nuScenes Weather, which includes 1.4M frames across clear, rain, fog, and snow conditions with synchronized Red, Green, Blue (RGB) (1280×720), Light Detection and Ranging (LiDAR) (128 channels), and radar data; (2) WeatherNet, providing 800K annotated thermal images (640×512) under diverse weather; and (3) a custom-collected corpus of 200K frames captured in extreme conditions (e.g., heavy snow, sandstorms) using a roof-mounted sensor suite (ZED 2i camera, Velodyne Alpha Prime LiDAR, FLIR thermal camera).

Temporal alignment is performed via timestamp synchronization (with an error of $\pm 1\text{ms}$) and spatial calibration using extrinsic matrices to map LiDAR points to RGB/thermal image coordinates [8]. LiDAR point clouds are filtered using a statistical outlier removal method:

$$\text{Point}_{\text{filtered}} = \{p \in \text{Cloud} \mid \frac{1}{k} \sum_{i=1}^k \|p - p_i\| < \theta\} \quad (1)$$

Where $k=20$ nearest neighbors and $\theta=0.5$.

Synthetic weather perturbations (e.g., rain streak injection, fog density adjustment) are generated using StyleGAN3 to expand training diversity, with 30% of augmented samples mimicking extreme conditions. Figure 4 shows the distribution of preprocessing steps by computational overhead, with LiDAR denoising and synthetic augmentation accounting for 62% of total preprocessing time.




Figure 4. Preprocessing Computational Overhead.

3.2. Model Architecture

The architecture comprises three key components: modality-specific encoders, cross-domain alignment modules, and multi-task prediction heads (Figure 3).

RGB image is a modified ResNet-50 with weather-aware attention blocks (WABs) that suppress weather-induced noise (e.g., rain streaks) by emphasizing stable features (e.g., road edges). WABs use squeeze-excitation mechanisms to recalibrate channel weights based on local variance [9].

LiDAR point clouds are processed by a spatio-temporal Graph Neural Network (GNN) where nodes represent point clusters and edges encode spatial proximity. This captures density patterns invariant to weather (e.g., vehicle contours preserved in sparse fog) [10].

The thermal images are processed by a U-Net variant with skip connections to preserve temperature gradients, which remain stable under low light or fog [11]. Table 2 summarizes the encoder architectures and their key parameters.

Table 2. Modality-Specific Encoder Parameters.

Modality	Backbone	Input Size	Params (M)	Weather-Aware Modules
RGB Images	ResNet-50 + WABs	1280×720	25.6	4 squeeze-excitation blocks
LiDAR Point Clouds	GNN	100K points	8.2	Spatio-temporal attention
Thermal Images	U-Net	640×512	12.1	Skip connections

Adversarial domain discriminators consist of three discriminators (one per sensor) that distinguish features from source (seen weather) versus target (unseen weather) domains. The adversarial loss minimizes domain discrepancy:

$$\mathcal{L}_{adv} = \mathbb{E}_{x \sim S}[\log D(x)] + \mathbb{E}_{x \sim T}[\log(1 - D(x))] \quad (2)$$

Where S and T denote source and target domains.

Cross-modal contrastive loss aligns features across sensors (e.g., RGB edges with LiDAR contours) using temperature-scaled cosine similarity:

$$\mathcal{L}_{contra} = -\log \frac{\exp(\text{sim}(f_{RGB}, f_{LiDAR})/\tau)}{\sum_{k \neq i} \exp(\text{sim}(f_i, f_k)/\tau)} \quad (3)$$

with $\tau=0.07$ to enhance inter-modal consistency.

3.3. Training and Optimization

The framework is trained end-to-end using a multi-task loss:

$$\mathcal{L}_{total} = \alpha \mathcal{L}_{detect} + \beta \mathcal{L}_{segment} + \gamma (\mathcal{L}_{adv} + \mathcal{L}_{contra}) \quad (4)$$

Where \mathcal{L}_{detect} (Focal Loss) and $\mathcal{L}_{segment}$ (Dice Loss) optimize perception tasks, and $\alpha = 0.4$, $\beta = 0.3$, $\gamma = 0.3$ (weighted via grid search).

Training uses AdamW optimizer with gradient clipping (L_2 norm ≤ 1.0) to prevent instability. A curriculum learning strategy progresses from simple (clear to light rain) to complex (heavy snow to fog) domain shifts.

4. Experiments

To validate the efficacy of the proposed multi-domain adaptation framework, comprehensive experiments were conducted to assess its performance in weather-robust perception tasks, including object detection and semantic segmentation. This section details the experimental setup, baseline comparisons, results analysis, and ablation studies, with key findings visualized in supporting figures and tables.

4.1. Experimental Setup

To ensure the generalizability of the proposed framework across diverse weather scenarios and sensor modalities, the evaluation was conducted on three benchmark datasets, each selected to cover distinct environmental conditions and data types. First, nuScenes Weather served as a foundational multi-modal dataset, comprising 1.2 million annotated frames spanning six weather conditions — clear, light rain, heavy rain, fog, light snow, and heavy snow. It includes synchronized Red, Green, Blue (RGB) images, Light Detection and Ranging (LiDAR) point clouds, and radar data, providing a comprehensive basis for assessing cross-sensor robustness. Second, WeatherNet focused specifically on thermal imaging under extreme weather, offering 800,000 annotated thermal images (640×512) with detailed labels for pedestrian and vehicle detection, thereby addressing the unique challenges of low-visibility and temperature-dependent feature extraction [12]. Additionally, a custom Extreme Weather Corpus was incorporated to fill gaps in rare conditions, featuring 200,000 frames captured during sandstorms and blizzards using a roof-mounted sensor suite (FLIR thermal camera, Velodyne LiDAR, ZED 2i RGB camera) mounted on a test vehicle. This custom dataset includes annotations for 13 semantic segmentation classes (e.g., road, vehicle, pedestrian), ensuring the framework's performance is validated even in infrequent but high-risk weather events. Together, these three datasets collectively cover a broad spectrum of weather intensities, sensor modalities, and annotation tasks, enabling a rigorous and comprehensive evaluation of the framework's robustness.

To comprehensively evaluate the performance of the proposed framework, five state-of-the-art methods were selected for comparison, covering diverse technical paradigms in weather-robust perception to ensure the validity of the benchmarking. Specifically, the Single-domain Residual Network-50 (ResNet-50) served as a representative of conventional single-domain models, trained exclusively on clear weather data and thus reflecting

the limitations of models that lack cross-weather adaptation capabilities. The Domain-Adversarial Neural Network (DANN), a two-domain adversarial adaptation model (clear to rain), was included as a typical example of single-pair domain adaptation, highlighting the constraints of methods designed for only two specific weather scenarios. Correlation Alignment (CORAL), a metric-based alignment method, was chosen to represent techniques that focus on multi-domain feature correlation, offering insights into the effectiveness of statistical feature alignment. The Multi-source Maximum Mean Discrepancy (MMD) [6], a multi-domain adaptation model that uses maximum mean discrepancy for alignment, was selected to benchmark against other multi-source adaptation approaches, enabling comparisons in handling multiple weather domains simultaneously [13]. Finally, CrossWeigh, a recent weather-robust perception model that relies solely on data augmentation without explicit domain adaptation mechanisms, provided a reference for evaluating the added value of domain adaptation beyond data-driven augmentation strategies. Together, these baselines span single-domain, single-pair adaptation, multi-domain alignment, and augmentation-based methods, ensuring a comprehensive assessment of the proposed framework's advancements.

To comprehensively assess the performance of the proposed framework and baselines across perception tasks and practical deployment requirements, three key evaluation metrics were adopted, each tailored to capture distinct aspects of model capability. Specifically, for object detection tasks, mean Average Precision (mAP) at Intersection over Union (IoU) values ranging from 0.5 to 0.95 was used. This standard metric in computer vision quantifies detection accuracy across varying levels of overlap between predicted and ground-truth bounding boxes, ensuring robustness to different object localization precision requirements. For semantic segmentation, mean Intersection over Union (mIoU) was selected, a widely accepted metric that measures the average overlap between predicted and true semantic regions across all classes, effectively reflecting the model's ability to distinguish fine-grained scene elements (e.g., road, vehicle, pedestrian) under diverse weather conditions. Additionally, efficiency was evaluated via inference latency (measured in milliseconds) on an NVIDIA A100 Graphics Processing Unit (GPU), a critical metric for automotive systems where real-time responsiveness (typically ≤ 50 ms) is essential to support timely decision-making and ensure driving safety. Together, these metrics collectively quantify the framework's accuracy in core perception tasks and its practical feasibility for deployment in autonomous driving systems.

A stratified 5-fold cross-validation was adopted, with 80% of data used for training (including all weather conditions) and 20% reserved for testing (unseen weather intensities, e.g., heavy fog not in training). All models were trained for 100 epochs with early stopping (patience=10) to avoid overfitting.

4.2. Main Results

Under heavy fog, it achieved 82.3% mAP for object detection, an improvement of 19.2% compared to the best baseline, Multi-source MMD, and 78.5% mIoU for segmentation, an increase of 16.7% [14]. In terms of efficiency, its inference latency of 42ms met real-time requirements for automotive systems (Table 3).

Table 3. Performance Comparison Across Weather Conditions.

Model	Clear (mAP)	Rain (mAP)	Fog (mAP)	Snow (mAP)	mIoU (All)	Latency (ms)
Proposed Framework	91.2%	85.6%	82.3%	79.1%	78.5%	42
Multi-source MMD [6]	89.5%	76.9%	63.1%	68.5%	61.8%	58
CrossWeigh [7]	88.3%	72.1%	59.8%	62.4%	57.3%	38

Single-domain ResNet	87.9%	65.3%	51.7%	54.2%	52.6%	35
----------------------	-------	-------	-------	-------	-------	----

Figure 5 illustrates performance degradation curves across increasing weather intensity, showing that the proposed framework exhibits the slowest drop in accuracy, with examples such as mAP decreasing by only 11.2% from light to heavy rain, compared to 28.7% for Single-domain ResNet-50. This confirms its ability to handle dynamic domain shifts.



Figure 5. Performance Degradation Under Increasing Weather Intensity.

4.3. Ablation Studies

To quantify the contribution of each key component within the proposed framework, a series of ablation experiments were conducted by iteratively removing individual modules and evaluating the resulting performance changes. Specifically, when the cross-modal contrastive loss was excluded from the framework, the mean Average Precision (mAP) under fog conditions decreased by 8.7%, a result that underscores the critical role of sensor fusion in preserving feature consistency across heterogeneous data sources. Similarly, disabling the adversarial domain alignment mechanism led to a 12.3% drop in mAP under snowy conditions, confirming that cross-domain feature alignment is essential for mitigating weather-induced distribution shifts. Additionally, omitting the weather-aware attention blocks in the RGB encoders resulted in a 9.5% accuracy loss in rainy conditions, which highlights the value of weather-specific feature recalibration in suppressing noise (e.g., rain streaks) and emphasizing stable visual cues (e.g., road edges). Collectively, these findings demonstrate that each component contributes uniquely to the framework's robustness, with their integration enabling the superior performance observed across diverse weather conditions.

4.4. Error Analysis

A pie chart (Figure 6) summarizes error sources in the worst-performing condition (heavy snow). Misclassifications due to severe occlusion accounted for 42%, while sensor noise (LiDAR point sparsity) contributed 35%, indicating potential improvements in occlusion handling and noise reduction.




Figure 6. Error Sources in Heavy Snow Conditions.

5. Discussion

The experimental results demonstrate that the proposed multi-domain adaptation framework significantly enhances autonomous driving perception robustness under diverse weather conditions, addressing critical limitations of existing methods. Three key contributions emerge from this work. First, the integration of heterogeneous sensors (Red, Green, Blue (RGB), Light Detection and Ranging (LiDAR), thermal) with domain-specific encoders enables the capture of complementary weather-robust features. These include cases where thermal imaging preserves edge information in fog, while LiDAR maintains depth cues in rain, thereby overcoming modality-specific vulnerabilities highlighted in prior studies. Second, the cross-modal contrastive loss and adversarial alignment mechanisms effectively mitigate dynamic domain shifts, as evidenced by the framework's 11.2% slower accuracy degradation under increasing weather intensity compared to single-domain models [15]. Third, the multi-task optimization strategy balances perception accuracy and real-time inference efficiency with a latency of 42 ms, meeting the strict timing constraints of automotive systems.

These findings align with broader efforts to advance weather-robust perception in autonomous driving. For example, Müller emphasized the need for multi-sensor fusion to counteract weather-induced sensor degradation, a challenge that is directly addressed by our framework's cross-modal alignment. Similarly, the framework's ability to generalize to unseen weather conditions responds to the call for scalable domain adaptation in real-world driving scenarios.

Despite these strengths, two primary limitations warrant consideration. First, performance in extreme sandstorms remains suboptimal (71.3% mean Average Precision (mAP)), which is attributed to severe LiDAR signal attenuation and thermal image saturation, a gap that is also noted in recent desert driving studies. Second, computational overhead (12.8 Giga Floating-Point Operations per Second (GFLOPs)) exceeds the 10 GFLOPs target for edge deployment, requiring model compression techniques such as knowledge distillation [16,17].

6. Conclusion

This study presents a multi-domain adaptation framework designed to enhance the robustness of autonomous driving perception under diverse weather conditions, addressing a critical bottleneck in real-world deployment of self-driving systems. By integrating heterogeneous sensor data (RGB images, LiDAR point clouds, and thermal imaging) through domain-specific encoders and cross-modal alignment mechanisms, the framework achieves superior performance across rain, fog, snow, and extreme weather conditions, outperforming state-of-the-art baselines in both object detection (up to 82.3% mAP) and semantic segmentation (up to 78.5% mIoU).

The core innovations of this work lie in three interconnected advancements: first, the use of weather-aware modality-specific encoders that leverage the unique strengths of each sensor, which include thermal imaging for preserving edge information in low visibility, LiDAR for stable depth cues in precipitation, and RGB for rich texture under moderate conditions, thereby mitigating the modality-specific vulnerabilities identified in prior research. Second, the integration of adversarial cross-domain alignment and contrastive multi-modal fusion, which enables the model to learn invariant features across weather domains while preserving task-relevant information, a capability critical for generalizing to unseen weather intensities. Third, the multi-task optimization strategy that balances perception accuracy and real-time inference efficiency (42ms latency), meeting the strict timing constraints of automotive systems.

These findings contribute to both theoretical and practical advancements in autonomous driving. Theoretically, they expand the application of multi-domain adaptation to multi-sensor fusion scenarios, providing a blueprint for addressing domain shift in complex, dynamic environments. Practically, the framework's ability to maintain high performance across diverse weather conditions reduces the risk of perception failure in real-world driving, which is linked to 23% of weather-related autonomous driving incidents.

Despite these successes — performance in extreme sandstorms (71.3% mAP) and computational overhead (12.8 GFLOPs) require further optimization. Future work will focus on integrating 77GHz radar data to enhance extreme weather robustness, adopting knowledge distillation for model compression, and developing dynamic adaptation modules that adjust to real-time weather changes. By addressing these challenges, the proposed framework moves closer to enabling reliable autonomous driving in the full spectrum of real-world weather conditions, advancing the goal of safe and widespread deployment of self-driving technology.

References

1. J. Vargas, et al., "An overview of autonomous vehicles sensors and their vulnerability to weather conditions," *Sensors*, vol. 21, no. 16, pp. 5397, 2021, doi: 10.3390/s21165397.
2. C. Xu and R. Sankar, "A comprehensive review of autonomous driving algorithms: Tackling adverse weather conditions, unpredictable traffic violations, blind spot monitoring, and emergency maneuvers," *Algorithms*, vol. 17, no. 11, pp. 526, 2024, doi: 10.3390/a17110526.
3. A. Khosravian, et al., "Multi-domain autonomous driving dataset: Towards enhancing the generalization of the convolutional neural networks in new environments," *IET Image Process.*, vol. 17, no. 4, pp. 1253–1266, 2023, doi: 10.1049/ipr2.12710.
4. L. Han, et al., "A Novel Multi-Object Tracking Framework Based on Multi-Sensor Data Fusion for Autonomous Driving in Adverse Weather Environments," *IEEE Sens. J.*, 2025, doi: 10.1109/JSEN.2025.3550506.
5. K. Muhammad, et al., "Vision-based semantic segmentation in scene understanding for autonomous driving: Recent achievements, challenges, and outlooks," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 12, pp. 22694–22715, 2022, doi: 10.1109/TITS.2022.3207665.
6. A. Piroli, et al., "Label-efficient semantic segmentation of LiDAR point clouds in adverse weather conditions," *IEEE Robot. Autom. Lett.*, vol. 9, no. 6, pp. 5575–5582, 2024, doi: 10.1109/LRA.2024.3396099.
7. B. Sun, J. Feng and K. Saenko, "Return of frustratingly easy domain adaptation," *Proc. AAAI Conf. Artif. Intell.*, vol. 30, no. 1, 2016, doi: 10.1609/aaai.v30i1.10306.
8. W. M. Kouw and M. Loog, "A review of domain adaptation without target labels," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 43, no. 3, pp. 766–785, 2019, doi: 10.1109/TPAMI.2019.2945942.

9. J. Wang, et al., "Superpixel segmentation with squeeze-and-excitation networks," *Signal Image Video Process.*, vol. 16, no. 5, pp. 1161–1168, 2022, doi: 10.1007/s11760-021-02066-2.
10. Y. Wang, et al., "Dynamic graph cnn for learning on point clouds," *ACM Trans. Graph.*, vol. 38, no. 5, pp. 1–12, 2019, doi: 10.1145/3326362.
11. S. Yang, "The Impact of Continuous Integration and Continuous Delivery on Software Development Efficiency", *J. Comput. Signal Syst. Res.*, vol. 2, no. 3, pp. 59–68, Apr. 2025, doi: 10.71222/pzvfqm21.
12. U. Tatli and C. Budak, "Biomedical image segmentation with modified U-Net," *Traitement du Signal*, vol. 40, no. 2, pp. 523–531, 2023, doi: 10.18280/ts.400211.
13. Z. Gu, et al., "Point cloud processing under adverse weather: a survey of datasets, enhancement, and denoising methods," *Signal Image Video Process.*, vol. 19, no. 9, pp. 718, 2025, doi: 10.1007/s11760-025-04352-9.
14. X. Chen, *Domain adaptation for autonomous driving*, Diss., University of Waterloo, 2020.
15. M. Alcon, et al., "Timing of Autonomous Driving Software: Problem Analysis and Prospects for," *Proc. 26th IEEE Real-Time Embedded Technol. Appl. Symp. (RTAS'20)*, 2020.
16. J. Li, et al., "Domain adaptive object detection for autonomous driving under foggy weather," *Proc. IEEE/CVF Winter Conf. Appl. Comput. Vis.*, 2023.
17. P. Chen, et al., "Distilling knowledge via knowledge review," *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021.

Disclaimer/Publisher's Note: The views, opinions, and data expressed in all publications are solely those of the individual author(s) and contributor(s) and do not necessarily reflect the views of PAP and/or the editor(s). PAP and/or the editor(s) disclaim any responsibility for any injury to individuals or damage to property arising from the ideas, methods, instructions, or products mentioned in the content.