








Contents lists available at ScienceDirect

International Journal of Applied Earth Observation and Geoinformation

journal homepage: www.elsevier.com/locate/jag

PhyDAWS: Physically-inspired data augmentation with weather simulation for domain-generalized point cloud segmentation[☆]

Jing Du^a, John Zelek^a , Chaolei Wang^c , Ting Han^c , Yiping Chen^c, Dedong Zhang^{a,d} , Jonathan Li^{a,b,d} ,^{*,1}

^a Department of Systems Design Engineering, University of Waterloo, Waterloo, ON N2L 3G1, Canada

^b Department of Geography and Environmental Management, University of Waterloo, Waterloo, ON N2L 3G1, Canada

^c School of Geospatial Engineering and Science, Sun Yat-sen University, Zhuhai, Guangdong, 519082, China

^d Tianfu Jiangxi Laboratory, Chengdu, Sichuan, 641419, China

ARTICLE INFO

Keywords:

Point cloud
Semantic segmentation
Deep learning
Domain generalization
Diverse weather scenarios

ABSTRACT

Point cloud semantic segmentation performance in autonomous driving systems degrades notably under adverse weather conditions, posing a challenge for reliable deployment. Current data augmentation methods for domain generalization often inadequately model the physical interactions between LiDAR sensors and atmospheric particles specific to diverse weather phenomena. This paper introduces PhyDAWS, a physically-inspired data augmentation framework designed for domain generalization in point cloud segmentation. It aims to enhance model robustness across various weather conditions using only source domain data for training. PhyDAWS incorporates two complementary weather simulation techniques. The first is a phenomenological method modeling angle-dependent occlusion (rain), height-dependent density variation (snow), and distance-dependent attenuation (fog). The second is a simulation based on Mie scattering theory, accounting for particle size distributions, refractive indices, and wavelength-dependent scattering. These physically-inspired augmentations are integrated with a dual-view contrastive learning strategy to promote the extraction of weather-invariant features. Evaluations on domain generalization benchmarks, including SemanticKITTI to SemanticSTF and SynLiDAR to SemanticSTF transfers, demonstrate the framework's effectiveness. Specifically for the SemanticKITTI to SemanticSTF task, PhyDAWS achieves improvements in overall mean Intersection over Union (mIoU) ranging from 4.1% to 10.3% compared to recent state-of-the-art approaches. By bridging the domain gap between clear and adverse weather conditions, this method advances the capability of autonomous systems for reliable environmental perception under diverse, unseen adverse conditions.

1. Introduction

Autonomous vehicles operate in an open world environment where unpredictable weather conditions pose significant challenges to perception systems (Tsakmakopoulou and Moustakas, 2024; H. Lin et al., 2025). While these systems demonstrate impressive performance in favorable conditions, their reliability deteriorates dramatically during adverse weather events, precisely when robust perception is most critical for maintaining safety (Xiao et al., 2022a; Yao et al., 2022). This performance gap represents one of the most serious barriers to widespread deployment of autonomous driving technology. In particular, weather-related perception failures have been identified as a

major contributor to disengagement events during autonomous driving tests (Zhang et al., 2023; Xu and Sankar, 2024; Grandhi, 2025).

LiDAR sensors, which provide critical 3D information through point clouds, are particularly susceptible to weather-induced degradation (Kettelgerdes et al., 2024; Su et al., 2024; Kuang et al., 2025). Rain introduces anisotropic occlusion patterns and unpredictable specular reflections (Li et al., 2024). Snow creates both geometric alterations through accumulation and visibility challenges through airborne particles (Awasthi and Varade, 2021; Khoshbakhhtnejad et al., 2024). Fog significantly attenuates signals and reduces the effective sensing range in moderate conditions (Zhan et al., 2024). These physical phenomena

[☆] This work was supported in part by the China Scholarship Council under Ph.D. Scholarship 202208350003.

* Corresponding author.

E-mail address: junli@uwaterloo.ca (J. Li).

¹ Given his role as Editor in chief, Jonathan Li had no involvement in the peer review of this article and had no access to information regarding its peer review. Full responsibility for the editorial process for this article was delegated to another journal editor.

<https://doi.org/10.1016/j.jag.2026.105280>

Received 24 September 2025; Received in revised form 5 March 2026; Accepted 3 April 2026

Available online 7 April 2026

1569-8432/© 2026 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

fundamentally transform point cloud characteristics, creating a substantial domain gap between clear and adverse weather data distributions. Conventional deep learning models struggle to bridge this gap (Zhu et al., 2021; Du et al., 2021; Wang et al., 2024a).

Domain generalization, which involves training models to perform effectively on previously unseen domains without adaptation, represents a critical capability for autonomous driving perception systems (Hu et al., 2025). Unlike domain adaptation approaches that require target domain data during training, domain generalization techniques must develop intrinsic robustness solely from source domain data (Liu et al., 2025). This capability is particularly valuable for autonomous vehicles, which must maintain reliable performance across diverse and unpredictable environmental conditions throughout their operational lifetime (B. Li et al., 2025).

Among various approaches to domain generalization, data augmentation has emerged as one of the most promising strategies (Xiao et al., 2023; Xu et al., 2024). By exposing models to synthetically diverse data during training, data augmentation enables the learning of more robust representations that better generalize to unseen conditions (Chen and Ying, 2025; Yu et al., 2025). For point cloud segmentation in autonomous driving, effective data augmentation can leverage the abundance of clear-weather data to improve model robustness under adverse weather conditions, thus alleviating the inherent imbalance in available training data (Xiao et al., 2023).

Existing data augmentation techniques for point clouds, however, suffer from fundamental limitations when applied to weather-related domain generalization:

1. *Geometric Transformations* (e.g., rotation, scaling) (Zhu et al., 2024) effectively adapt to viewpoint variations but overlook the physical interactions between LiDAR beams and atmospheric conditions critical to weather effects.
2. *Point Cloud Perturbations* (e.g., random point dropout, noise addition) (Kim et al., 2021) introduce uniform distortions, failing to replicate the spatially heterogeneous and distance-dependent degradation caused by weather phenomena.
3. *Feature Adjustments* (e.g., global intensity scaling) (Iwana and Uchida, 2021) apply simplistic modifications to point attributes, insufficient for capturing the non-linear signal attenuation patterns of diverse weather conditions.
4. *Generative Model Augmentation* (e.g., GANs, style transfer) (Sapkota et al., 2025; Y. Li et al., 2025) depends on access to adverse weather data for training, contradicting the domain generalization objective of performing without target domain samples.
5. *Sample Mixing Augmentation* (e.g., PointMixup) (Chen et al., 2020) enhances decision boundaries through interpolation, yet generates combinations lacking the physical realism of weather-induced transformations.

These limitations highlight a critical gap in the field: the need for physically accurate, unsupervised data augmentation techniques that can transform clear-weather point clouds into realistic adverse-weather equivalents without requiring real adverse-weather data for calibration or validation.

The degradation of LiDAR measurements under adverse weather follows predictable physical laws governing light-particle interactions. Rain, snow, and fog each induce characteristic patterns of signal attenuation, scattering, and occlusion determined by particle size distributions, refractive indices, and atmospheric density gradients. Simulating these underlying physical processes can inherently capture the spatially heterogeneous, distance-dependent, and material-specific characteristics that distinguish authentic weather effects from simplistic perturbations. This physical grounding offers a principled pathway to generating realistic weather variations using only clear-weather source data.

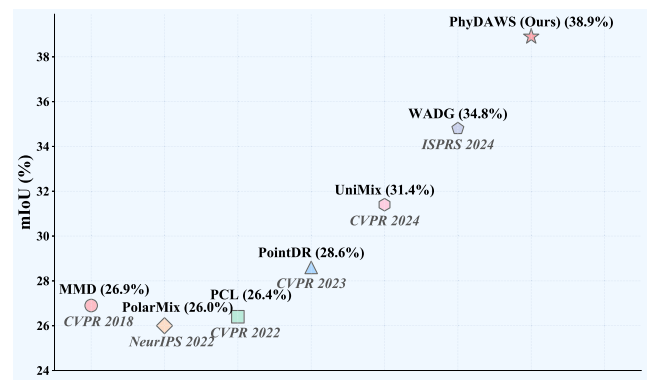


Fig. 1. Comparison of cross-weather domain generalization methods using SemanticKITTI as the source domain and SemanticSTF as the target domain. Our proposed PhyDAWS method achieves the highest mIoU performance (38.9%) across adverse weather conditions, significantly outperforming previous state-of-the-art approaches.

Building on this foundation, we propose PhyDAWS, a novel physically-inspired data augmentation framework that addresses this gap through weather simulation based on LiDAR-atmospheric interactions. As shown in Fig. 1, the proposed physically-inspired approach significantly outperforms existing methods for cross-weather domain generalization in point cloud segmentation. By incorporating these physical models into our data augmentation pipeline, we create realistic synthetic weather variations that help networks learn weather-invariant features. This enables robust generalization to unseen weather conditions.

The core innovation of PhyDAWS lies in its two complementary physically-inspired weather simulation methods. First, our phenomenological weather simulation approach captures weather-specific effects. This method models key phenomena such as angle-dependent occlusion in rain, where falling raindrops create anisotropic occlusion patterns based on their orientation; height-dependent density variation in snow, which accounts for the natural concentration gradient of snowflakes with altitude; and distance-dependent attenuation in fog, which captures the exponential signal decay characteristic of light propagation through particulate media. Second, we develop a more sophisticated Mie scattering-based simulation that rigorously models the interactions between LiDAR beams and atmospheric particles based on electromagnetic theory. This advanced approach accounts for weather-specific particle size distributions, complex refractive indices, and wavelength-dependent scattering properties.

Our main contributions are summarized as follows:

1. We introduce a physically-inspired data augmentation framework that enables point cloud segmentation models to generalize to diverse weather conditions without requiring target domain data during training.
2. We develop two complementary physically-inspired weather simulation methods that model the complex interactions between LiDAR beams and atmospheric particles, generating realistic weather transformations that effectively bridge the domain gap.
3. We demonstrate through extensive experiments that our method significantly outperforms existing approaches on challenging benchmarks, maintaining robust performance across diverse weather conditions without requiring target domain adaptation.

2. Related work

This section reviews recent advancements in point cloud perception systems for autonomous driving, focusing on four key research

areas: semantic segmentation architectures, perception robustness in adverse weather conditions, domain generalization techniques, and data augmentation strategies. These complementary fields collectively address the challenges of building reliable perception systems capable of operating across diverse real-world environments.

2.1. Point cloud semantic segmentation

Point cloud semantic segmentation has evolved from the foundational point-based processing of PointNet (Qi et al., 2017a) and PointNet++ (Qi et al., 2017b) toward architectures optimized for autonomous driving applications. Contemporary research prioritizes computational efficiency and accuracy for large-scale outdoor environments, with sparse convolutional networks (Choy et al., 2019; Peng et al., 2024; Du et al., 2025; F. Lin et al., 2025; Wang et al., 2024c; Deng et al., 2025) and transformer-based approaches (Zhao et al., 2021; Wu et al., 2022, 2024) emerging as dominant paradigms. This evolution reflects the growing demand for high-performance, practical solutions capable of robust 3D scene understanding in dynamic, real-world settings.

LiDAR-based semantic segmentation has emerged as a critical component in autonomous driving systems, enabling precise environmental perception and scene understanding (Yang et al., 2024; Zhuang et al., 2024; Zhao et al., 2025; Kong et al., 2025). Recent research demonstrates a significant trend toward integrating multiple data representations and processing strategies to overcome the inherent challenges of outdoor environments. ViDAR (Yang et al., 2024) introduces an innovative pre-training framework that forecasts future point clouds from historical images, effectively transforming visual inputs into a 3D latent space through a specialized architecture. This approach enhances perception and planning capabilities while addressing limitations in multi-view geometry and temporal modeling. Moving beyond 2D projection methods, Zhuang et al. (2024) integrate temporal residual features with cylindrical partitioning to manage noise and density issues in dynamic environments, improving moving object segmentation and SLAM performance. For road boundary segmentation, Zhao et al. (2025) develop CurbNet, employing a Multi-Scale and Channel Attention module to extract curb features directly from 3D point clouds. LaserMix++ (Kong et al., 2025) addresses the challenge of labor-intensive annotations through semi-supervised learning and multi-modal fusion. By integrating LiDAR spatial priors with camera features through beam mixing and feature distillation, this approach significantly reduces annotation requirements while maintaining competitive accuracy. This represents an important advance toward practical deployment in autonomous driving systems.

2.2. Adverse weather perception in autonomous driving

Environmental robustness remains a critical challenge for autonomous driving perception systems, with adverse weather conditions significantly degrading sensor performance and reliability (Xiao et al., 2023; C. Sun et al., 2024; Du et al., 2024; Heo et al., 2025; Xie et al., 2025; Bi et al., 2025; Nawaz et al., 2025). Rain, fog, and snow introduce complex interference patterns that compromise the accuracy of LiDAR and camera-based perception systems. Recent research has addressed this challenge through several complementary approaches.

The SemanticSTF dataset provides dense point-wise annotations of LiDAR data captured in fog, snow, and rain (Xiao et al., 2023). In the same work, the PointDR framework employs geometry style randomization and embedding aggregation to enhance robustness to weather-induced distortions. Similarly, C. Sun et al. (2024) develop the MSP dataset specifically for snowfall impacts on LiDAR, accompanied by analytical methods that quantify noise distributions and their effects on detection capabilities.

Algorithmic innovations for cross-weather generalization have emerged as another promising direction. The framework in Du et al.

(2024) integrates Adaptive Feature Normalization with Dual-Attention Fusion, effectively calibrating features and incorporating contextual information to improve segmentation accuracy across diverse weather conditions. For bird's-eye-view perception, Xie et al. (2025) introduce RoboBEV, a benchmark for evaluating robustness against natural corruptions and sensor failures, demonstrating that pre-training and temporal fusion significantly enhance performance in challenging conditions.

Sensor-specific approaches have also yielded valuable insights. Research in Heo et al. (2025) utilizes controlled weather chamber experiments to develop theoretical models quantifying signal attenuation, providing a foundation for sensor technology enhancement. Task-specific solutions include (Bi et al., 2025), which combines traditional image processing with deep learning to improve lane detection under poor visibility. Another solution is presented in Nawaz et al. (2025), which integrates LiDAR, RGB, and thermal sensor data through a novel fusion framework based on a modified YOLOv8 architecture to improve object detection in adverse conditions. Collectively, these research directions advance the field toward perception systems capable of maintaining performance regardless of weather conditions.

2.3. Domain generalization for point cloud perception

Point cloud perception models require robust domain generalization (DG) capabilities to maintain performance across diverse deployment environments (Rafi et al., 2024; T. Sun et al., 2024; Zhao et al., 2024; Kim et al., 2024; Liu et al., 2025; Sanchez et al., 2025). Current research addresses several key challenges in this field. Rafi et al. (2024) survey DG methods for semantic segmentation, highlighting techniques like data augmentation and domain randomization that enhance generalization without target domain data. To reduce annotation burden, T. Sun et al. (2024) propose a cross-modal framework leveraging unlabeled images for 3D semantic segmentation, using a dual-branch network with less than 1% labeled point clouds. For environmental robustness, UniMix (Zhao et al., 2024) constructs a Bridge Domain and implements a universal mixing operator to mitigate discrepancies across varying scenes and adverse environmental conditions. Addressing sensor heterogeneity, Kim et al. (2024) introduce the Density Discriminative Feature Embedding module that extracts density-specific features to improve generalization across varying point cloud densities. Liu et al. (2025) tackle rotational variations through orientation mining and contrastive learning, enhancing rotational robustness in cross-domain settings. Finally, COLA (Sanchez et al., 2025) harmonizes heterogeneous LiDAR datasets using coarse labels, providing integrated strategies for domain generalization across varying sensors and environments.

2.4. Data augmentation strategies for point clouds

Point cloud data augmentation techniques address the critical challenge of limited training data in autonomous driving perception systems (Zhu et al., 2024; Wang et al., 2024b; J. Li et al., 2025; Sun et al., 2025; Chen and Ying, 2025; Yu et al., 2025). Comprehensive categorization of augmentation techniques by Zhu et al. (2024) establishes a taxonomy dividing methods into basic transformations and specialized approaches. This systematic classification identifies effective strategies for different perception tasks while highlighting future research directions in generative augmentation techniques. Novel mixing-based augmentation methods have introduced significant improvements in feature preservation. PointPatchMix (Wang et al., 2024b) advances beyond point and block-level approaches by implementing patch-level mixing with a teacher-guided scoring module that preserves critical local structures while generating diverse samples. FPSMix (Chen and Ying, 2025) innovates through farthest point sampling combined with significance-weighted loss functions, specifically enhancing global feature learning while maintaining classification accuracy. For real-world applications, AdaptPoint++ (J. Li et al., 2025) overcomes limitations of

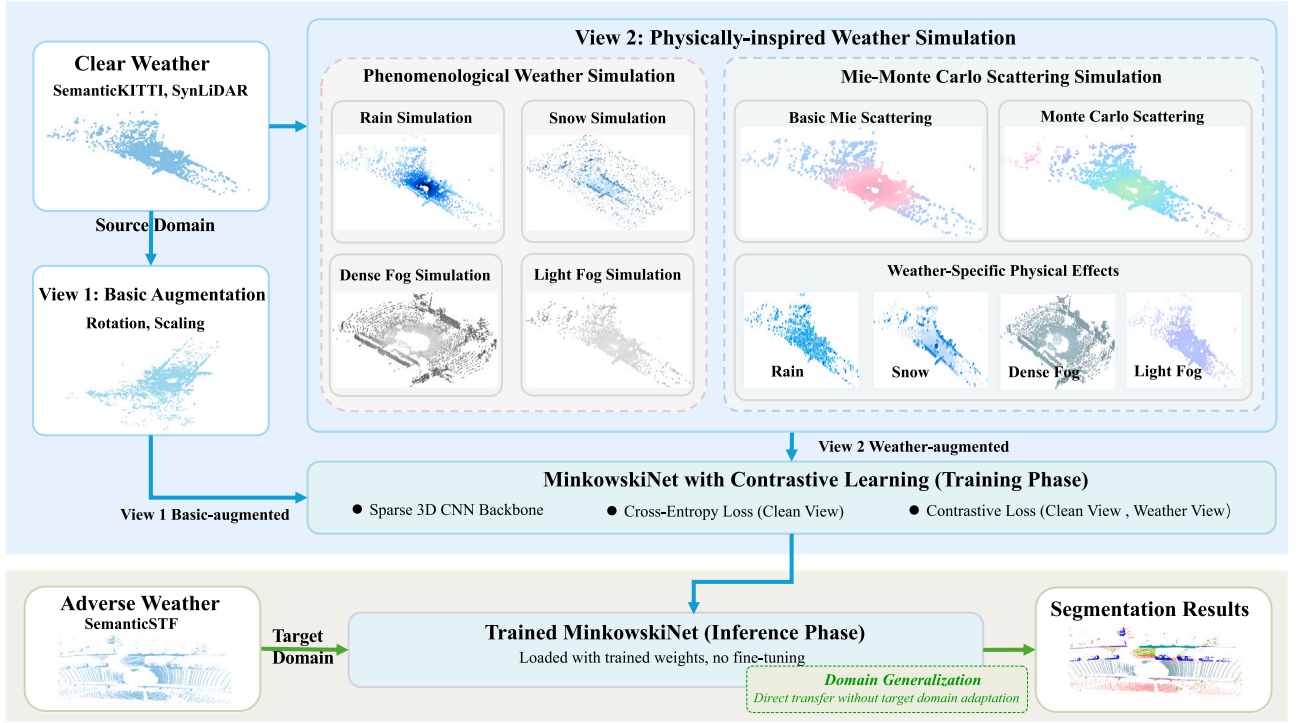


Fig. 2. Overview of the proposed physically-inspired data augmentation with weather simulation framework.

static offline augmentation through a sample-adaptive framework that dynamically generates realistic, challenging examples. Cross-domain adaptation challenges are specifically targeted by LBDA (Sun et al., 2025), which introduces bidirectional ellipsoid-constrained transformations to effectively reduce domain gaps in few-shot scenarios. Similarly, the self-supervised strategies in Yu et al. (2025) directly counter geometric disparities between synthetic and real-world data through translation distance prediction and cascaded relational learning, improving adaptation without relying on conventional adversarial approaches.

3. Method

In this section, we present a physically-inspired approach for domain generalization in point cloud semantic segmentation across diverse weather conditions. As illustrated in Fig. 2, the proposed method (PhyDAWS) mitigates performance degradation of semantic segmentation models in adverse weather by integrating physically-inspired weather simulations into the data augmentation pipeline. This improves robustness in previously unseen conditions without access to target-domain labels.

3.1. Problem formulation

Given a source domain D_S with labeled point cloud data $\{X_S, Y_S\}$ collected in clear weather and a target domain D_T with unlabeled point cloud data X_T from adverse-weather scenarios, we aim to learn a model f_θ using only $\{X_S, Y_S\}$ that generalizes to X_T . The domain generalization objective is to minimize the expected risk on the source domain while encouraging representations that transfer well to the target domain:

$$\min_{\theta} \mathbb{E}_{(X,Y) \sim D_S} [\mathcal{L}(f_\theta(\mathcal{T}(X)), Y)] \quad (1)$$

where θ denotes the model parameters, \mathcal{L} is the complete loss function that promotes domain-invariant feature learning, and \mathcal{T} represents data augmentation transformations. The fundamental challenge lies in the absence of labeled target-domain data during training, as only

unlabeled samples X_T from D_T are available without their corresponding labels. This constraint distinguishes domain generalization from domain adaptation methods, which typically require access to target-domain data during training. This challenge is particularly pronounced in LiDAR point-cloud segmentation because adverse weather systematically alters point density (via occlusions and spurious returns) and reflectance intensity patterns (via absorption and scattering). These alterations create a substantial domain gap that conventional data augmentation techniques often fail to address effectively.

3.2. Network architecture

We employ the MinkowskiNet architecture (Choy et al., 2019), a sparse convolutional neural network designed for efficient 3D point cloud processing. After voxelization with voxel size $v = 0.05$ m, the input point cloud is processed through an encoder-decoder structure with skip connections. Our network outputs both semantic segmentation logits and normalized feature embeddings for contrastive learning:

$$\{\mathbf{p}, \mathbf{z}\} = f_\theta(X), \quad (2)$$

where $\mathbf{p} \in \mathbb{R}^{N \times C}$ represents segmentation logits for C classes and $\mathbf{z} \in \mathbb{R}^{N \times 128}$ represents normalized feature embeddings (produced via a projection head after the final layer).

3.3. Domain generalization strategy

We adopt a dual-view contrastive learning framework that leverages physically-inspired weather simulations to enable robust domain generalization. For each training point cloud, we generate two views with different augmentation strategies:

$$X_1 = \mathcal{T}_{\text{basic}}(X), \quad (3)$$

$$X_2 = \mathcal{T}_{\text{weather}}(X), \quad (4)$$

where X is the input point cloud, $\mathcal{T}_{\text{basic}}$ applies conventional augmentations (random rotation and scaling), and $\mathcal{T}_{\text{weather}}$ additionally

applies our physically-inspired weather simulations. This asymmetric augmentation strategy creates a controlled disparity between the two views that encourages the network to learn weather-invariant features while maintaining geometric consistency. Both views are processed through the network:

$$\{\mathbf{p}_1, \mathbf{z}_1\} = f_\theta(X_1), \quad (5)$$

$$\{\mathbf{p}_2, \mathbf{z}_2\} = f_\theta(X_2), \quad (6)$$

where f_θ represents the neural network with parameters θ , \mathbf{p}_1 and \mathbf{p}_2 are the segmentation logits, and \mathbf{z}_1 and \mathbf{z}_2 are the feature embeddings from the first and second views, respectively.

To facilitate contrastive learning, we utilize a class-wise prototype memory bank \mathcal{M} that stores representative feature vectors for each semantic class, updated via momentum averaging (Xiao et al., 2023):

$$\mathcal{M}_c^{(t)} = m \cdot \mathcal{M}_c^{(t-1)} + (1 - m) \cdot \frac{1}{|\Omega_c|} \sum_{i \in \Omega_c} \mathbf{z}_{1,i}, \quad (7)$$

where $\mathcal{M}_c^{(t)}$ is the prototype vector for class c at iteration t , m is a momentum coefficient, $\Omega_c = \{i | y_i = c\}$ is the set of points with ground truth label c , $|\Omega_c|$ is the cardinality of this set, and $\mathbf{z}_{1,i}$ is the feature embedding of point i from the first view.

Our training objective realizes the formulation in Eq. (1), where the complete loss function is defined as:

$$\mathcal{L}(f_\theta(X_1), Y) = \mathcal{L}_{\text{CE}}(\mathbf{p}_1, Y) + \lambda \cdot \mathcal{L}_{\text{contrast}}, \quad (8)$$

where \mathcal{L}_{CE} is the standard cross-entropy loss for semantic segmentation, \mathbf{p}_1 represents the predicted logits from the first view, Y is the ground truth labels, and λ is a balancing coefficient.

The contrastive loss encourages alignment between features from the weather-augmented view and the corresponding class prototypes:

$$\mathcal{L}_{\text{contrast}} = -\frac{1}{N} \sum_{i=1}^N \log \frac{\exp(\mathbf{z}_{2,i} \cdot \mathcal{M}_{y_i} / \tau)}{\sum_{c=1}^C \exp(\mathbf{z}_{2,i} \cdot \mathcal{M}_c / \tau)}, \quad (9)$$

where N is the number of points, $\mathbf{z}_{2,i}$ is the feature embedding of point i from the weather-augmented view, \mathcal{M}_{y_i} is the prototype for the ground truth class of point i , τ is a temperature parameter, and C is the total number of classes.

This contrastive approach encourages the network to learn weather-invariant representations by aligning features from the weather-augmented view (\mathbf{z}_2) with the corresponding class prototypes (\mathcal{M}_{y_i}) derived from the clean view. By minimizing the distance between features from the weather-augmented view and the corresponding class prototypes, the network learns to extract consistent semantic information regardless of weather conditions. This effectively bridges the domain gap between clear and adverse weather scenarios.

Within this contrastive learning framework, the role of physically-inspired weather modeling extends beyond data diversification and directly shapes the nature of feature alignment across views. The physical simulations introduce structured signal variations that are consistent with real atmospheric propagation processes, including range-dependent attenuation, stochastic scattering, and geometry-dependent point dropout. These variations systematically perturb the input at the signal level while preserving the underlying semantic structure of the scene. As a result, corresponding points across the clean and weather-augmented views differ in intensity, local density, and return completeness, yet remain associated with the same semantic class.

When optimized under the proposed contrastive objective, the network is encouraged to suppress weather-induced signal fluctuations and focus on invariant geometric and structural cues that remain stable across physically plausible sensing conditions. Features extracted from the weather-augmented view are pulled toward class-wise prototypes derived from the clean view, implicitly enforcing invariance to physically grounded degradation patterns. This mechanism establishes a direct link between physical modeling and feature learning,

whereby atmospheric signal variations act as structured perturbations that guide the network toward representations that generalize across unseen weather conditions.

From a domain generalization perspective, this coupling between physical simulation and contrastive alignment is critical. The proposed physically-inspired transformations expose the model to signal variations governed by the same physical laws encountered at deployment, ensuring that the augmentation distribution aligns with real-world degradation mechanisms. Consequently, the learned representations capture semantic information that is robust to weather-induced domain shifts, supporting reliable transfer from clear-weather or synthetic source domains to adverse-weather target environments without requiring access to target-domain labels.

3.4. Physically-inspired weather simulation

The simultaneous adoption of phenomenological modeling and Mie-based scattering simulation is motivated by the multi-faceted nature of weather-induced degradation in LiDAR point clouds. In adverse weather, the domain shift is driven by interrelated mechanisms that affect both the spatial completeness of returns and the signal characteristics of LiDAR measurements. A robust augmentation strategy therefore benefits from covering structured geometric effects as well as physically grounded intensity attenuation and scattering variability, so that the generated training samples remain consistent with real sensing processes.

Phenomenological modeling provides an efficient way to reproduce macroscopic patterns that dominate the geometric organization of point clouds, including anisotropic occlusion in rain, height-dependent density variation in snow, and range-limited visibility in fog. In parallel, the Mie-Monte Carlo component characterizes wavelength-dependent extinction and stochastic particle interactions along the beam path, enriching the training distribution with intensity attenuation and return variability consistent with atmospheric propagation. Together, these two augmentation paths expose the network to complementary perturbation modes that jointly shape adverse-weather LiDAR observations, supporting more stable representation learning and improved transferability when deployment environments vary across regions and acquisition campaigns.

3.4.1. Phenomenological weather simulation

The first method implements phenomenological models of LiDAR-weather interactions, motivated by atmospheric physics and prior studies in meteorology (Ulbrich, 1983; Bohren and Huffman, 2008; Seinfeld and Pandis, 2016). For a point cloud \mathcal{X} with points (x, y, z, i) where i represents intensity, the approach models weather-specific effects:

Rain simulation: Rain affects LiDAR through angle-dependent occlusion and distance-dependent intensity attenuation. The rain-induced occlusion is modeled as:

$$\mathcal{P}_{\text{rain}} = \{p \in \mathcal{X} \mid \gamma_p > p_{\text{base}} + \alpha \cdot |\cos(\phi_p - \phi_{\text{rain}})|\}, \quad (10)$$

where $\mathcal{P}_{\text{rain}}$ represents the subset of points that are retained after applying the rain simulation, \mathcal{X} is the original point cloud, $\gamma_p \sim \mathcal{U}(0, 1)$ is a random variable for each point, $\phi_p = \arctan 2(z_p, \sqrt{x_p^2 + y_p^2})$ is the vertical angle of point p , $\phi_{\text{rain}} \sim \mathcal{U}(-5^\circ, 5^\circ)$ is the rain inclination angle, $p_{\text{base}} = 0.1$ is a baseline probability, and $\alpha = 0.3$ controls the angular dependency strength. This formulation captures the anisotropic nature of rain occlusion. Raindrops are not perfectly spherical but become elongated as they fall, creating direction-dependent scattering patterns (Thurai and Bringi, 2005). An illustrative example is shown in Fig. 3. Notably, our approach leverages the implicit geometric relationship within point cloud data and probabilistic modeling, avoiding the computationally expensive explicit simulation of intersections between

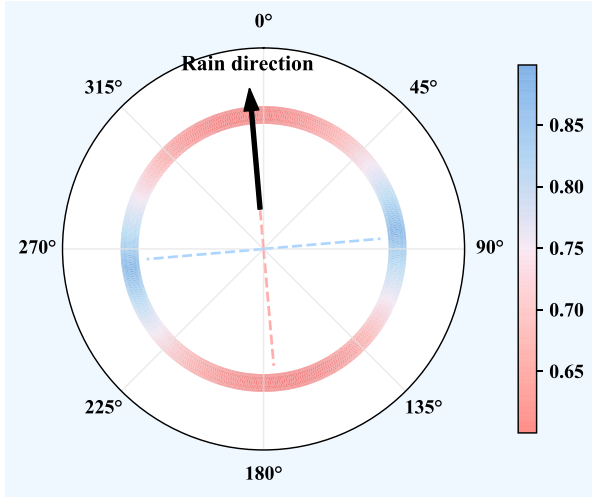


Fig. 3. Visualization of the angle-dependent rain occlusion model. The polar plot shows point retention probability as a function of angle relative to rainfall direction (black arrow). Blue regions (perpendicular to rainfall) indicate high retention rates (~90%), while pink regions (parallel to rainfall) show low retention rates (~60%). The plot illustrates a specific instance with rain inclination angle of -5° .

each LiDAR beam and individual raindrops, while still achieving physically accurate occlusion patterns. The intensity attenuation follows a modified Beer–Lambert law (Middleton, 1952; Ishimaru et al., 1978):

$$I_{\text{rain}}(d) = I_0 \cdot \exp(-\beta_{\text{rain}} \cdot d^{\gamma_{\text{rain}}}), \quad (11)$$

where I_0 is the original intensity, $d = \|\mathbf{x}\|_2$ is the distance from sensor origin, $\beta_{\text{rain}} = 0.02$ is the rain-specific attenuation coefficient, and $\gamma_{\text{rain}} = 1.2$ accounts for the non-linear attenuation in rain. The non-linear exponent models the complex interaction between raindrops and LiDAR signals, including multiple scattering and absorption effects that compound with distance. This intensity modification trains the network to interpret attenuated signals correctly, improving segmentation accuracy in rainy conditions.

Snow simulation: Snow creates a complex pattern of height-dependent occlusion, intensity attenuation, and scattered returns. The snow density factor varies with height:

$$\rho_{\text{snow}}(z) = \text{clip}\left(1 - \frac{z - z_{\min}}{z_{\max} - z_{\min} + \epsilon}, \rho_{\min}, 1.0\right), \quad (12)$$

where z is the height coordinate of a point, z_{\min} and z_{\max} are the minimum and maximum z -coordinates in the point cloud, $\rho_{\min} = 0.5$ sets a minimum density threshold, and $\epsilon = 10^{-6}$ prevents division by zero. This formulation captures the observation that snow concentration typically decreases with height, resulting in height-dependent occlusion patterns (Pruppacher et al., 1998). Fig. 4 visualizes this height-dependent density model, where concentric rings represent different height levels (labeled in meters) with corresponding dropout rates (shown as percentages). The color gradient from dark blue to light blue represents decreasing snow density with increasing height, leading to fewer dropped points (red crosses) at higher elevations. Uniformly distributed snowflake points (white asterisks) simulate noise that challenges the network to differentiate between meaningful structural information and weather-induced artifacts. By modeling this vertical density gradient, we prepare the network to handle the non-uniform occlusion patterns characteristic of snowy conditions.

The point retention probability is:

$$\mathcal{P}_{\text{snow}} = \{p \in \mathcal{X} \mid \gamma_p > \beta_d \cdot \rho_{\text{snow}}(z_p)\}, \quad (13)$$

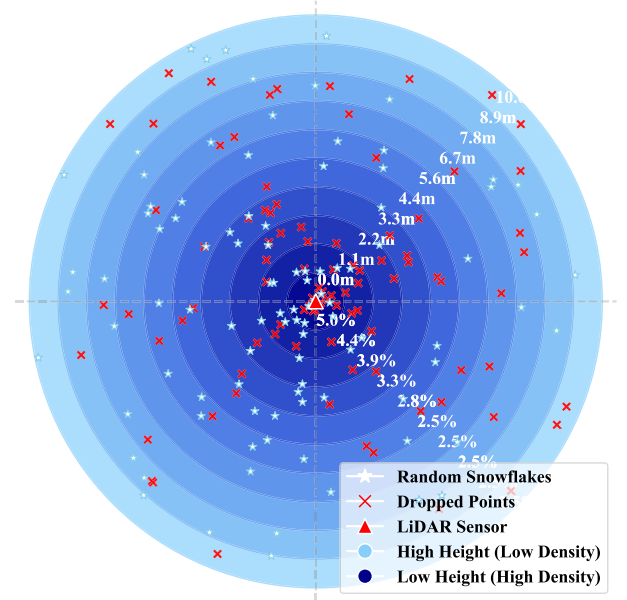


Fig. 4. Visualization of our height-dependent snow density model. Concentric rings represent height levels with corresponding dropout rates. More dropped points (red crosses) appear at lower heights due to higher snow density, while added snowflakes (white asterisks) are uniformly distributed throughout the scene.

where $\mathcal{P}_{\text{snow}}$ is the set of points retained after snow simulation, \mathcal{X} is the original point cloud, $\beta_d = 0.05$ is the snow-specific dropout coefficient, $\gamma_p \sim \mathcal{U}(0, 1)$ is a random variable, and z_p is the height coordinate of point p . The intensity attenuation is:

$$I_{\text{snow}}(d) = I_0 \cdot \exp(-\beta_{\text{snow}} \cdot d^{\gamma_{\text{snow}}}), \quad (14)$$

where I_0 is the original intensity, d is the distance from sensor origin, $\beta_{\text{snow}} = 0.03$ and $\gamma_{\text{snow}} = 1.1$ control the snow-specific attenuation pattern. Additionally, we generate simulated snowflakes randomly distributed throughout the scene volume:

$$N_{\text{snow}} = \lfloor \rho_f \cdot |\mathcal{X}| \rfloor, \quad (15)$$

where N_{snow} is the number of simulated snowflakes to add, $\rho_f = 0.1$ controls the density of added snowflakes relative to the original point cloud size, and $|\mathcal{X}|$ is the cardinality of the original point cloud. These snowflakes are positioned according to:

$$\mathcal{X}_{\text{flakes}} = \{(x, y, z, i) \mid x, y, z \sim \mathcal{U}(\mathcal{X}_{\min}, \mathcal{X}_{\max})\}, \quad (16)$$

where $\mathcal{X}_{\text{flakes}}$ is the set of simulated snowflake points, \mathcal{X}_{\min} and \mathcal{X}_{\max} define the bounding box of the original point cloud, and $i \sim \mathcal{U}(0.05, 0.4)$ is the intensity value assigned to each simulated snowflake. These simulated snowflakes introduce noise points that challenge the network to differentiate between meaningful structural information and snow-induced artifacts, enhancing robustness to such disturbances in real-world scenarios.

Fog simulation: Fog primarily affects LiDAR through exponential attenuation and range limitation (Bohren and Huffman, 2008; Bijelic et al., 2020). The intensity attenuation is:

$$I_{\text{fog}}(d) = I_0 \cdot \exp(-\beta_{\text{fog}} \cdot d^{\gamma_{\text{fog}}}), \quad (17)$$

where I_0 is the original intensity, d is the distance from sensor origin, $\gamma_{\text{fog}} = 1.5$ accounts for multiple scattering effects in fog, and β_{fog} varies with fog density ($\beta_{\text{light}} = 0.03$ for light fog, $\beta_{\text{dense}} = 0.15$ for dense fog). This higher exponent value compared to rain and snow reflects

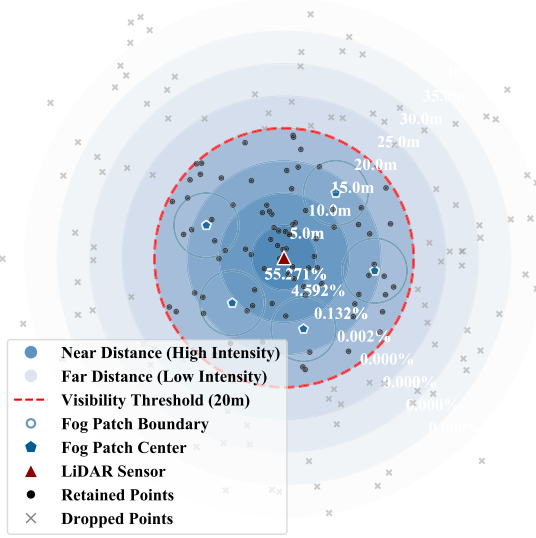


Fig. 5. Visualization of dense fog simulation showing intensity attenuation with distance (percentages) and visibility threshold at 20 m (red dashed line). Local fog patches (blue circles) create spatial heterogeneity, while points beyond the threshold are discarded (gray crosses).

the increased scatter and absorption that occurs in foggy environments due to the high density of small water droplets. For dense fog, we also limit visibility:

$$\mathcal{P}_{\text{vis}} = \{p \in \mathcal{X} \mid d(p) < V_{\text{max}}\}, \quad (18)$$

where \mathcal{P}_{vis} is the set of points retained after applying visibility constraints, \mathcal{X} is the original point cloud, $d(p) = \|\mathbf{x}_p\|_2$ is the distance from the sensor to point p , and $V_{\text{max}} \sim \mathcal{U}(15, 30)$ meters defines the visibility range. This visibility threshold accurately models the hard cutoff in detection range that occurs in dense fog conditions, forcing the network to make segmentation decisions with incomplete scene information. This capability is critical for autonomous systems operating in foggy environments.

Fig. 5 illustrates our dense fog simulation model, depicting how LiDAR intensity decreases exponentially with distance from the sensor. Points within the central region maintain high retention rates, which rapidly diminish to near-zero at the visibility threshold. The model also incorporates localized fog patches of varying density, creating realistic heterogeneous fog conditions that challenge perception systems. This physically-grounded approach helps the network maintain effective feature representations even under substantial information loss caused by dense fog. Additionally, the complete discarding of points beyond the visibility threshold forces the network to learn to make reliable predictions with incomplete scene information, a critical requirement for real-world autonomous driving systems.

3.4.2. Mie-Monte Carlo scattering simulation

The Mie-Monte Carlo scattering simulation is grounded in Mie scattering theory, which describes how electromagnetic radiation interacts with spherical particles (Hulst and van de Hulst, 1981; Bohren and Huffman, 2008). This approach models the complex interactions between LiDAR beams and atmospheric particles with theoretical accuracy, providing detailed characterization of the electromagnetic interactions that occur in various weather conditions.

Basic Mie Scattering: We derive physically accurate extinction coefficients by modeling particle size distributions for different weather conditions using the gamma distribution (Ulbrich, 1983):

$$n(D) = N_0 \cdot \frac{D^{\alpha-1}}{\Gamma(\alpha)\beta^\alpha} \exp\left(-\frac{D}{\beta}\right), \quad (19)$$

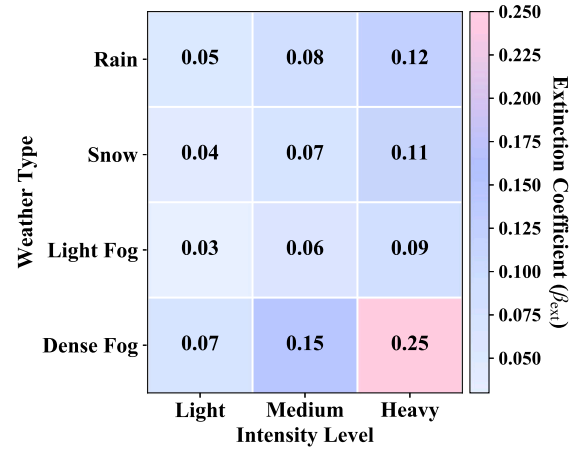


Fig. 6. Precomputed extinction coefficients lookup table (β_{ext}) for different weather types and intensity levels based on Mie scattering theory. Higher values indicate stronger attenuation of LiDAR signals.

where $n(D)$ is the particle number density per unit diameter, D is the particle diameter in millimeters, N_0 is the number concentration (particles per cubic meter), α is the shape parameter, β is the scale parameter (in mm), and $\Gamma(\alpha)$ is the gamma function. These parameters are calibrated for each weather type and intensity level based on meteorological data.

The total extinction coefficient is calculated by integrating over the particle size distribution from $D_{\text{min}} = 0.01$ mm to $D_{\text{max}} = 10$ mm:

$$\beta_{\text{ext}} = \int_{D_{\text{min}}}^{D_{\text{max}}} Q_{\text{ext}}(D, \lambda, m) \cdot \frac{\pi D^2}{4} \cdot n(D) dD, \quad (20)$$

where β_{ext} is the extinction coefficient, $Q_{\text{ext}}(D, \lambda, m)$ is the Mie extinction efficiency (the ratio of the extinction cross-section to the geometric cross-section) for a particle of diameter D with complex refractive index m (1.33+0i for water, 1.31+0i for ice) at LiDAR wavelength $\lambda = 905$ nm. The Mie extinction efficiency is calculated from electromagnetic theory as:

$$Q_{\text{ext}}(D, \lambda, m) = \frac{2}{x^2} \sum_{n=1}^{\infty} (2n+1) \Re\{a_n + b_n\}, \quad (21)$$

where $x = \pi D/\lambda$ is the size parameter, \Re denotes the real part of a complex number, and a_n and b_n are the complex Mie scattering coefficients derived from solving Maxwell's equations for spherical particles (Hulst and van de Hulst, 1981). These coefficients depend on the size parameter x and the complex refractive index m , representing the scattered electromagnetic field components. For computational efficiency, we precompute these extinction coefficients for various weather types and intensity levels and store them in a lookup table, as shown in Fig. 6.

This physically-grounded approach enables accurate modeling of how different weather conditions affect LiDAR signal propagation. As shown in the lookup table, dense fog exhibits the strongest attenuation effect with extinction coefficients ranging from 0.07 to 0.25, while light fog shows the mildest effect (0.03 to 0.09). Rain and snow occupy intermediate positions with rain generally causing slightly stronger attenuation than snow at equivalent intensity levels. These precomputed values allow our augmentation pipeline to efficiently simulate realistic weather effects without solving complex scattering equations at runtime.

Monte Carlo scattering: The basic attenuation follows the Beer-Lambert law (Ishimaru et al., 1978; Bohren and Huffman, 2008):

$$I_{\text{att}}(d) = I_0 \cdot \exp(-\beta_{\text{ext}} \cdot d), \quad (22)$$

where $I_{\text{att}}(d)$ is the attenuated intensity at distance d , I_0 is the original intensity, and β_{ext} is the extinction coefficient for the specific weather

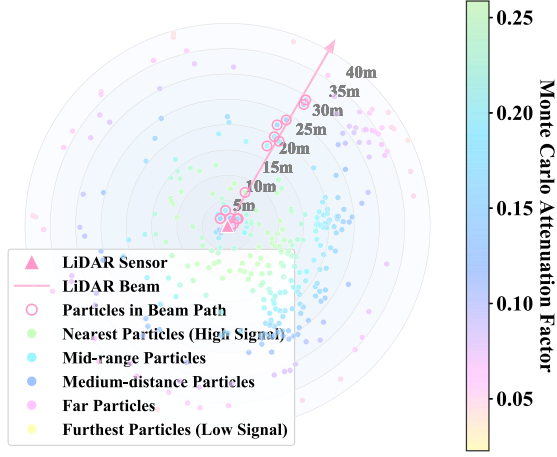


Fig. 7. Visualization of Monte Carlo scattering effects on LiDAR signal propagation. The color gradient represents attenuation factors: green (center) indicates high signal return from nearby particles, transitioning through cyan and blue to purple and yellow (perimeter) showing diminishing returns from distant particles. Pink markers highlight particles directly in the beam path.

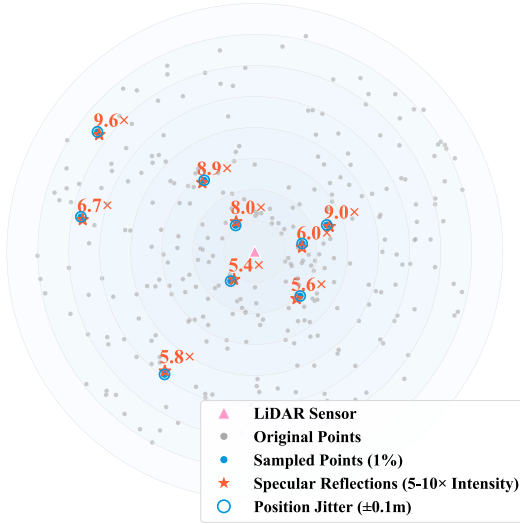


Fig. 8. Visualization of rain specular reflection simulation. Blue circles represent randomly sampled points (1% of original points), while red stars show the corresponding specular reflections with position jitter and intensity amplification (5 – 10 \times). This augmentation replicates the high-intensity reflections from raindrops that can affect LiDAR perception.

condition. We enhance this with a Monte Carlo approach to capture the stochastic nature of particle distribution (Grabner and Kvicera, 2013):

$$I_{MC}(d) = I_{att}(d) \cdot \mathcal{F}_{MC}(d, \beta_{ext}), \quad (23)$$

where $I_{MC}(d)$ is the intensity after applying Monte Carlo adjustments, and $\mathcal{F}_{MC}(d, \beta_{ext})$ is the Monte Carlo factor derived from simulations of particle interactions along the beam path:

$$\mathcal{F}_{MC}(d, \beta_{ext}) = \text{CDF}(d | \lambda). \quad (24)$$

Here, $\text{CDF}(d | \lambda)$ is the cumulative distribution function of an exponential distribution with mean free path $\lambda = 1/\beta_{ext}$, representing the probability distribution of scattering events along the beam path. This approach improves upon deterministic models by incorporating the statistical variations in particle encounters, better representing the stochastic nature of real-world weather phenomena.

Fig. 7 illustrates the Monte Carlo scattering simulation, where concentric rings represent distances from 5 m to 40 m from the LiDAR sensor (pink triangle). The colored particles demonstrate how the attenuation factor varies with distance, with particles in high-density clusters showing the statistical nature of light-particle interactions. This visualization captures the stochastic behavior that deterministic models fail to represent, particularly the random distribution of particles in adverse weather and their varying impact on signal propagation.

We also incorporate material-specific reflection properties:

$$I_{final}(p) = I_{MC}(d(p)) \cdot \mathcal{R}(c_p), \quad (25)$$

where $I_{final}(p)$ is the final intensity value for point p , $d(p)$ is the Euclidean distance from the sensor origin to point p , $\mathcal{R}(c_p)$ is the reflectivity coefficient for semantic class c_p , calibrated from empirical measurements. This class-aware reflection modeling acknowledges that different surface materials interact with LiDAR signals differently under adverse weather conditions (Kaasalainen et al., 2005). By incorporating this physical insight, the network can learn the complex relationship between object material properties and their appearance in various weather conditions, enhancing segmentation accuracy for all semantic classes.

Weather-specific physical effects: For each weather type, we incorporate unique physical phenomena. For rain, we model specular reflections:

$$\mathcal{X}_{rain} = \mathcal{X} \cup \{(\mathbf{x} + \delta, i \cdot \alpha_{spec}) \mid (\mathbf{x}, i) \in \mathcal{X}_{sampled}\}, \quad (26)$$

where \mathcal{X}_{rain} is the point cloud after adding rain effects, \mathcal{X} is the original point cloud, $\mathcal{X}_{sampled}$ is a random subset of points (approximately 1% of the original point cloud), \mathbf{x} is the 3D position vector of a point, $\delta \sim \mathcal{U}(-0.1, 0.1)$ is position jitter in meters, i is the original intensity value of a point, and $\alpha_{spec} \sim \mathcal{U}(5, 10)$ is intensity amplification. These specular reflections replicate the bright returns that occur when LiDAR pulses directly reflect off water droplets, a phenomenon that can create false positives in object segmentation systems (Hasirlioglu et al., 2016; Bijelic et al., 2018). By training with these artifacts present, the network learns to distinguish between actual objects and rain-induced specular reflections.

Fig. 8 illustrates our approach to simulating rain specular reflection effects. The visualization shows the original point cloud (gray dots), from which 1% of points are randomly sampled (blue circles). For each sampled point, a corresponding reflection point (red star) is generated with a position offset of approximately ± 0.1 m. Additionally, each reflection point's intensity is amplified by a factor of 5 – 10 \times , as indicated by the numerical values adjacent to the stars. This process accurately simulates the phenomenon where raindrops act as small reflective surfaces that create secondary, high-intensity returns in LiDAR scans. These artificial reflection points challenge perception systems in the same way that real-world rain does, forcing models to develop robustness to these common weather-induced artifacts.

For snow, we simulate accumulation on horizontal surfaces:

$$\mathcal{X}_{snow} = \mathcal{X} \cup \{(\mathbf{x} + [0, 0, h_{snow}], i) \mid (\mathbf{x}, i) \in \mathcal{X}_{hsurf}\}, \quad (27)$$

where \mathcal{X}_{snow} is the point cloud after adding snow effects, \mathcal{X} is the original point cloud, \mathcal{X}_{hsurf} represents points on horizontal surfaces (identified by normal vector analysis or semantic class), \mathbf{x} is the 3D position vector of a point, $h_{snow} \sim \mathcal{U}(0, 0.05)$ meters is the local snow height, and i is the intensity value of the point. This simulation accounts for the geometric changes that occur in snowy conditions, as snow accumulation alters the apparent shape and height of objects and ground surfaces. By exposing the network to these geometric alterations during training, we enable it to maintain accurate segmentation even when objects' apparent shapes change due to snow accumulation.

The snow simulation incorporates several key physical principles beyond simple point perturbations: it performs selective geometric modification by adding points above original surfaces with 10% probability, mimicking how snow accumulates on horizontal surfaces; it

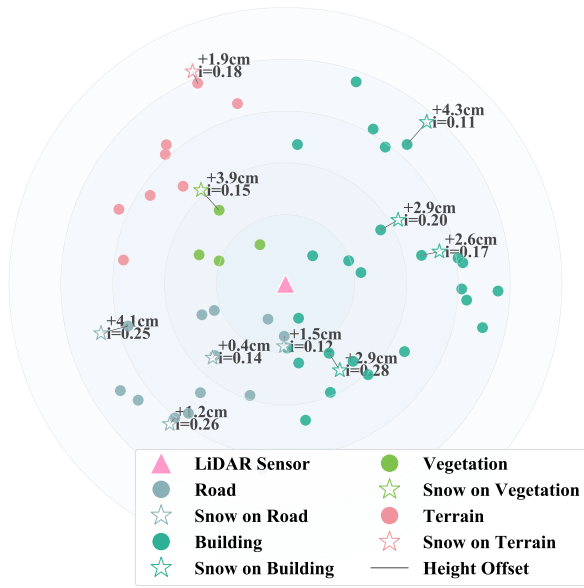


Fig. 9. Visualization of snow accumulation effects on point clouds. Representative surface classes (road, building, vegetation, etc.) are shown with corresponding snow accumulation points. The simulation uses uniform height distribution $\mathcal{U}(0, 0.05)$ m across all surface types, with snow points inheriting semantic labels from parent points while adopting lower intensity values (0.1–0.3).

applies height variations up to 5 cm based on a uniform distribution $\mathcal{U}(0, 0.05)$; it maintains semantic consistency through label inheritance, ensuring snow-covered objects retain their original classification, as illustrated in Fig. 9; and it models snow’s optical properties by assigning lower intensity values (0.1–0.3), reflecting snow’s high light absorption. This comprehensive approach maintains data structure integrity by synchronizing coordinates, intensities, and labels to create physically realistic snow-augmented point clouds.

For dense fog, we enforce visibility limitations by discarding points that exceed a fixed maximum visible distance, which approximates the signal attenuation predicted by the Beer–Lambert law. Specifically, we define the fog-affected point cloud as

$$\mathcal{X}_{\text{fog}} = \{p \in \mathcal{X} \mid d(p) < d_{\text{max}}\}, \quad (28)$$

where \mathcal{X} is the original point cloud, p is a point in the cloud, and $d(p)$ is the Euclidean distance from the sensor to p . Here, d_{max} is an empirically chosen constant (e.g., 20 meters) that approximates the maximum visible distance under dense fog conditions. This fixed threshold simulates the effect of severe fog, where signal intensity decays to a level that renders points beyond d_{max} undetectable.

For light fog, we apply the same physical attenuation principles as in dense fog but with a significantly lower extinction coefficient. In this case, we do not enforce an explicit visibility cut-off. Instead, we rely on the natural gradual attenuation of signal intensity with distance, as modeled by

$$I_{\text{fog}}(p) = I_{\text{MC}}(d(p)) \cdot \mathcal{R}(c_p), \quad (29)$$

where $I_{\text{MC}}(d(p))$ is the intensity after applying Monte Carlo attenuation to account for stochastic scattering effects, and $\mathcal{R}(c_p)$ is a class-specific reflectivity factor. This approach simulates the moderate degradation in light fog, in which objects remain visible at longer ranges but with progressively reduced clarity. The resulting partial attenuation challenges the network to handle intermediate levels of degradation, thereby bridging the gap between clear conditions and severely limited visibility.

3.4.3. Parameter calibration and sensor adaptability

The selection of parameters throughout the physically-inspired simulation framework is grounded in established atmospheric science literature. Particle size distributions are modeled using gamma distributions with weather-specific shape and scale parameters (Ulbrich, 1983). For rain, the shape parameter is set to 2 with scale parameters ranging from 0.5 to 2.0 mm depending on intensity, producing particle distributions characteristic of stratiform and convective precipitation. Snow uses a shape parameter of 3 with smaller scale values reflecting the different microphysical properties of ice crystals. Fog employs higher shape parameters (4–5) with scale values below 1.0 mm, capturing the narrow size distributions of small suspended droplets. The number concentration scaling factors increase with weather intensity, following empirically observed relationships between precipitation rate and particle density (Ulbrich, 1983). The complex refractive indices (1.33+0i for liquid water, 1.31+0i for ice) are well-established optical constants at near-infrared wavelengths (Hulst and van de Hulst, 1981). The phenomenological attenuation coefficients follow consistent physical principles: rain coefficients ($\beta_{\text{rain}} = 0.02$) correspond to moderate rainfall conditions (Middleton, 1952), fog coefficients span from $\beta_{\text{light}} = 0.03$ for light fog to $\beta_{\text{dense}} = 0.15$ for dense fog (Bohren and Huffman, 2008; Bijelic et al., 2020), and snow coefficients ($\beta_{\text{snow}} = 0.03$) reflect the lower extinction efficiency of ice crystals compared to liquid water droplets. The non-linear exponents ($\gamma_{\text{rain}} = 1.2$, $\gamma_{\text{snow}} = 1.1$, $\gamma_{\text{fog}} = 1.5$) account for multiple scattering effects, with fog exhibiting the strongest non-linearity due to the high density of small particles (Ishimaru et al., 1978).

The wavelength of 905 nm reflects the predominant configuration in current automotive LiDAR systems. The wavelength dependence of Mie scattering implies that extinction coefficients would differ for systems operating at alternative wavelengths such as 1550 nm, where increased water absorption and reduced scattering efficiency for small particles lead to different attenuation characteristics, particularly in fog. The modular structure of the simulation framework accommodates such adaptations by allowing the lookup tables to be regenerated with updated refractive indices and size parameters. Sensor-specific characteristics including beam divergence, pulse repetition rate, and detector sensitivity can similarly be incorporated through adjustments to the attenuation and noise models, enhancing applicability across diverse LiDAR configurations.

To ensure computational efficiency in large-scale training scenarios, the framework employs a decoupled design in which computationally intensive components are not executed online during network optimization. Wavelength-dependent extinction coefficients derived from Mie theory are precomputed offline and stored in lookup tables, while stochastic scattering effects are approximated through lightweight sampling operations at the point level. Importantly, all simulations are confined to the training stage; inference requires only a standard forward pass without additional physical modeling. This separation ensures that the framework scales naturally to large datasets and maintains a favorable balance between physical fidelity and computational efficiency.

4. Experiments

4.1. Datasets

Our experiments focus on domain generalization for point cloud semantic segmentation from clear-weather or synthetic domains to adverse-weather conditions. We utilize three prominent LiDAR datasets: SemanticKITTI (Behley et al., 2019) and SynLiDAR (Xiao et al., 2022b) as source domains, and SemanticSTF (Xiao et al., 2023) as the unseen target domain.

The SemanticKITTI dataset serves as our real-world source domain, comprising point clouds captured with a Velodyne HDL-64E sensor

Table 1

3D semantic segmentation results (%) with SemanticKITTI as source and SemanticSTF as target. Per-class IoU for 19 semantic categories across all weather conditions; per-weather mIoU for those 19 categories under four weather scenarios; and overall all-weather mIoU for the same 19 categories across all weather conditions are reported.

Method	Car	Bicycle	Motorcycle	Truck	Other-Vehicle	Person	Bicyclist	Motorcyclist	Road	Parking	Sidewalk	Other-Ground	Building	Fence	Vegetation	Trunk	Terrain	Pole	Traffic-Sign	Dense-Fog	Light-Fog	Rain	Snow	All-Weather	
	Per-Class IoU (%)																			Per-Weather mIoU (%)				All-Weather mIoU (%)	
Baseline	55.9	0.0	0.2	1.9	10.9	10.3	6.0	0.0	61.2	10.9	32.0	0.0	67.9	41.6	49.8	27.9	40.8	29.6	17.5	29.5	26.0	28.4	21.4	24.4	24.4
Dropout (Srivastava et al., 2014)	62.1	0.0	15.5	3.0	11.5	5.4	2.0	0.0	58.4	12.8	26.7	1.1	72.1	43.6	52.9	34.2	43.5	28.4	15.5	29.3	25.6	29.4	24.8	24.8	25.7
Perturbation	74.4	0.0	0.0	23.3	0.6	19.7	0.0	0.0	60.3	10.8	33.9	0.7	72.0	45.2	58.7	17.5	42.4	22.1	9.7	26.3	27.8	30.0	24.5	24.5	25.9
PolarMix (Xiao et al., 2022a)	57.8	1.8	3.8	16.7	3.7	26.5	0.0	2.0	65.7	2.9	32.5	0.3	71.0	48.7	53.8	20.5	45.4	25.9	15.8	29.7	25.0	28.6	25.6	26.0	26.0
MMD (Li et al., 2018)	63.6	0.0	2.6	0.1	11.4	28.1	0.0	0.0	67.0	14.1	37.9	0.3	67.3	41.2	57.1	27.4	47.9	28.2	16.2	30.4	28.1	32.8	25.2	26.9	26.9
PCL (Yao et al., 2022)	65.9	0.0	0.0	17.7	0.4	8.4	0.0	0.0	59.6	12.0	35.0	1.6	74.0	47.5	60.7	15.8	48.9	26.1	27.5	28.9	27.6	30.1	24.6	26.4	26.4
PointDR (Xiao et al., 2023)	67.3	0.0	4.5	19.6	9.0	18.8	2.7	0.0	62.6	12.9	38.1	0.6	73.3	43.8	56.4	32.2	45.7	28.7	27.4	31.3	29.7	31.9	26.2	28.6	28.6
UniMix (Zhao et al., 2024)	82.7	6.6	8.6	4.5	15.1	35.5	15.5	37.7	55.8	10.2	36.2	1.3	72.8	40.1	49.1	33.4	34.9	23.5	33.5	34.8	30.2	34.9	30.9	31.4	31.4
WADG (Du et al., 2024)	72.0	0.0	32.9	37.0	1.9	37.7	6.8	52.9	59.9	10.7	31.8	2.2	76.0	48.8	62.7	34.0	49.3	23.6	20.4	39.5	32.5	31.7	29.4	34.8	34.8
Ours	65.9	0.8	51.2	38.1	10.6	45.3	5.5	59.4	65.3	19.0	35.7	4.0	78.5	50.4	56.8	40.0	47.7	29.0	35.2	40.2	34.7	37.7	33.4	38.9	

under normal weather conditions around Karlsruhe, Germany. It provides point-wise annotations for 28 classes (typically merged into 19 evaluation classes) across 43,552 scans.

The SynLiDAR dataset functions as our synthetic source domain, containing 198,396 scans with over 19 billion points collected from virtual environments. It provides dense point-wise annotations for 32 semantic classes, making it the largest synthetic dataset available for LiDAR semantic segmentation. The synthetic nature of SynLiDAR allows us to investigate the generalization capabilities from synthetic to real-world data under adverse conditions.

The SemanticSTF dataset serves as our target domain for evaluating generalization performance. It provides point-wise annotations across 2076 scans captured with a Velodyne HDL-64 S3D LiDAR sensor. Crucially, these scans represent various adverse weather conditions: 694 in snow, 637 in dense fog, 631 in light fog, and 114 in rain. The data was collected across multiple European countries (Germany, Sweden, Denmark, and Finland), ensuring geographical diversity.

4.2. Experimental setup

To ensure consistent evaluation, we align the semantic categories between datasets. For transfer from SemanticKITTI to SemanticSTF, we merge moving and non-moving classes into 19 base categories for training, and map the “invalid” category in SemanticSTF to “ignored” during evaluation. For transfer from SynLiDAR to SemanticSTF, we use a mapping configuration to align SynLiDAR’s 32 original classes to a common evaluation set of 19 categories, with additional classes mapped to “ignored”. This alignment ensures fair comparison across different generalization approaches (Xiao et al., 2023; Du et al., 2024).

We adhere to a strict domain generalization setting where models are trained exclusively on the source domain (either SemanticKITTI or SynLiDAR) and evaluated directly on SemanticSTF without any adaptation or fine-tuning using target data. This protocol rigorously tests our model’s ability to generalize semantic understanding from clear/synthetic conditions to challenging, unseen adverse weather scenarios.

For network architecture, we adopt MinkowskiNet (Choy et al., 2019) as the segmentation backbone. This choice is consistent with PointDR (Xiao et al., 2023), UniMix (Zhao et al., 2024), and WADG (Du et al., 2024), all of which employ MinkowskiNet for feature extraction in their original implementations. This architectural consistency ensures that performance differences can be attributed to the respective methodological contributions of each approach. The evaluation results presented in Section 4.4 demonstrate that PhyDAWS achieves substantial improvements over all comparative methods across diverse weather conditions and semantic categories, with improvement patterns corresponding to the atmospheric degradation mechanisms modeled by each simulation component. These observations suggest that the performance gains originate from the physically-inspired augmentation framework.

4.3. Training and inference procedure

The training procedure implements the domain generalization strategy by integrating physically-inspired weather simulations with contrastive learning. Weather types are randomly selected for each sample with equal probability across rain, snow, light fog, and dense fog (25% each). This randomization ensures the network encounters diverse atmospheric conditions during training.

During inference, only a single forward pass is required to produce semantic segmentation predictions, with final class predictions determined by the argmax of the logit scores. The method achieves effective domain generalization without requiring any labeled data from the target domain or adaptation during inference, making it particularly valuable for autonomous systems operating across unpredictable weather conditions.

4.4. Evaluation and visualization results

4.4.1. Evaluation results: from SemanticKITTI to SemanticSTF

Table 1 presents the quantitative results of cross-domain semantic segmentation performance from SemanticKITTI (source domain) to SemanticSTF (target domain). The proposed PhyDAWS method achieves state-of-the-art performance with an overall mIoU of 38.9%, significantly outperforming previous approaches. Compared to recent state-of-the-art methods, PhyDAWS demonstrates an improvement of 4.1% over WADG (Du et al., 2024) (34.8%), 7.5% over UniMix (Zhao et al., 2024) (31.4%), and 10.3% over PointDR (Xiao et al., 2023) (28.6%). These substantial improvements validate the effectiveness of the physically-inspired data augmentation strategy.

Regarding weather-specific performance, PhyDAWS demonstrates consistent superiority across all adverse conditions. PhyDAWS achieves the highest per-weather mIoU across all four conditions, reaching 40.2% in dense fog, 34.7% in light fog, 37.7% in rain, and 33.4% in snow. Compared with representative prior methods, it outperforms PointDR by 8.9% in dense fog and WADG by 6.0% in rain, where signal attenuation and scattering effects are most pronounced. Class-wise analysis reveals clear correspondence between performance gains and the physically-inspired modeling strategies. The Mie-Monte Carlo scattering simulation contributes to detecting safety-critical road users under signal attenuation, with motorcycles achieving the most dramatic improvement (51.2% IoU, 18.3% above WADG) and person segmentation reaching 45.3% IoU. The phenomenological weather simulation enhances structural and ground surface perception through distance-dependent attenuation modeling, with buildings achieving 78.5% IoU, road reaching 65.3% IoU, and parking areas improving to 19.0% IoU (8.3% above WADG). Traffic signs achieve 35.2% IoU (14.8% above WADG) through the phenomenological simulation. These results validate that physically-inspired augmentation provides targeted improvements aligned with specific physical modeling strategies, enhancing perception capabilities most relevant to autonomous driving safety.

Table 2

3D semantic segmentation results (%) with SemanticKITTI as source and SemanticSTF as target under dense fog conditions only. Per-class IoU for 19 semantic categories in dense fog and the mIoU across these categories are reported. Dash (–) indicates categories with no instances in the dense fog scenario.

Method	Per-Class IoU (%)																			Dense-Fog mIoU (%)
	Car	Bicycle	Motorcycle	Truck	Other-Vehicle	Person	Bicyclist	Motorcyclist	Road	Parking	Sidewalk	Other-Ground	Building	Fence	Vegetation	Trunk	Terrain	Pole	Traffic-Sign	
Baseline	74.7	–	–	7.8	0.0	6.4	8.9	0.0	72.2	0.6	33.8	0.0	59.6	48.7	56.9	27.4	56.4	27.2	21.1	29.5
Dropout (Srivastava et al., 2014)	67.5	–	–	1.9	0.0	8.9	2.8	0.0	70.9	5.6	29.0	0.8	64.6	44.0	60.0	31.6	60.6	28.1	21.3	29.3
Perturbation	68.6	–	–	8.8	0.0	6.0	0.0	0.0	66.6	14.8	24.3	0.1	52.2	43.5	60.1	19.4	54.1	16.3	11.5	26.3
PolarMix (Xiao et al., 2022a)	52.3	–	–	17.2	0.0	3.6	0.0	19.3	75.2	0.0	28.7	0.6	62.4	49.5	60.5	29.0	55.4	20.8	30.7	29.7
MMD (Li et al., 2018)	75.5	–	–	0.3	0.0	4.2	0.0	0.0	75.4	11.2	33.6	0.5	64.8	51.7	64.7	26.1	62.3	23.0	23.0	30.4
PCL (Yao et al., 2022)	64.3	–	–	11.7	0.0	0.6	0.0	0.0	72.4	3.8	31.3	0.8	63.1	46.5	65.7	19.4	64.3	18.5	28.9	28.9
PointDR (Xiao et al., 2023)	69.2	–	–	7.1	0.0	2.4	6.7	0.0	73.5	8.5	33.6	0.2	65.6	47.6	63.6	31.0	60.7	24.4	38.8	31.3
WADG (Du et al., 2024)	78.8	–	–	40.6	0.0	7.0	47.7	67.3	71.6	0.3	29.8	1.4	65.0	48.0	66.5	33.0	60.8	21.3	33.0	39.5
Ours	82.3	–	–	21.6	0.0	10.0	13.6	75.1	75.1	0.1	33.1	4.9	69.6	56.6	62.3	37.0	60.1	24.9	56.8	40.2

Table 3

3D semantic segmentation results (%) with SemanticKITTI as source and SemanticSTF as target under light fog conditions only. Per-class IoU for 19 semantic categories in light fog and the mIoU across these categories are reported.

Method	Per-Class IoU (%)																			Light-Fog mIoU (%)
	Car	Bicycle	Motorcycle	Truck	Other-Vehicle	Person	Bicyclist	Motorcyclist	Road	Parking	Sidewalk	Other-Ground	Building	Fence	Vegetation	Trunk	Terrain	Pole	Traffic-Sign	
Baseline	60.0	0.0	0.0	1.3	10.9	12.3	0.0	0.0	68.6	4.5	36.0	0.0	61.5	53.1	55.6	38.0	44.7	29.2	18.2	26.0
Dropout (Srivastava et al., 2014)	63.2	0.0	0.0	3.2	10.2	5.5	0.0	0.0	63.8	4.9	29.4	0.1	62.5	53.1	58.6	42.5	46.6	27.8	14.3	25.6
Perturbation	76.6	0.0	0.0	38.2	0.0	21.9	0.0	0.0	66.6	8.8	34.6	0.1	62.4	56.1	63.2	25.3	46.2	22.4	6.5	27.8
PolarMix (Xiao et al., 2022a)	42.6	0.2	0.0	29.4	3.3	17.0	0.0	0.2	69.8	0.7	33.1	0.1	56.2	56.3	54.9	24.7	44.8	24.1	16.6	25.0
MMD (Li et al., 2018)	63.6	0.0	0.0	0.1	13.3	25.9	0.0	0.0	73.9	5.6	42.8	0.1	64.1	55.3	61.9	36.6	50.7	29.2	9.9	28.1
PCL (Yao et al., 2022)	66.3	0.0	0.0	26.7	0.2	8.7	0.0	0.0	67.8	5.0	36.7	0.4	64.3	58.0	66.1	21.2	53.1	25.5	24.6	27.6
PointDR (Xiao et al., 2023)	65.9	0.0	0.0	29.7	4.4	11.4	0.9	0.0	70.9	8.8	43.3	0.0	66.5	55.1	61.3	43.0	49.1	29.1	24.3	29.7
WADG (Du et al., 2024)	68.5	0.0	0.0	54.0	4.8	43.8	0.0	0.0	69.1	4.3	34.3	0.5	68.4	58.6	66.6	43.8	55.3	24.2	20.7	32.5
Ours	64.9	0.0	0.0	55.4	3.4	54.4	1.4	0.0	74.0	8.8	39.4	4.2	69.5	56.2	60.4	47.3	52.2	31.5	35.7	34.7

Focusing specifically on the challenging dense fog condition within the SemanticSTF target domain, Table 2 details the cross-domain segmentation performance derived from the SemanticKITTI source. In this particularly challenging scenario characterized by severely limited visibility, PhyDAWS demonstrated superior robustness, yielding a state-of-the-art mIoU of 40.2%. Compared to recent state-of-the-art methods, this represents improvements of 0.7% over WADG (39.5%), 8.9% over PointDR (31.3%), and 9.8% over MMD (30.4%). These improvements are particularly significant given that dense fog presents the most extreme visibility constraints among all weather conditions. Performance gains correlate closely with the fog-specific modeling components. Traffic signs achieve the most remarkable improvement (56.8% IoU, 23.8% above WADG) through the phenomenological simulation that preserves surface responses under severe attenuation. Cars reach the highest absolute performance (82.3% IoU), benefiting from the distance-dependent attenuation modeling. The Mie-Monte Carlo scattering simulation contributes to detecting objects under severe visibility constraints, with motorcyclists achieving 75.1% IoU. Structural classes also benefit: fences reach 56.6% IoU (8.6% above WADG) and trunks achieve 37.0% IoU through the modeling of cylindrical object responses under fog conditions. These improvements directly enhance autonomous driving safety in visibility-limited scenarios.

Table 3 presents the performance under light fog conditions, where PhyDAWS achieves a mIoU of 34.7%. This notably outperforms WADG by 2.2%, and PointDR by 5.0%. Unlike dense fog where visibility is severely restricted, light fog introduces moderate signal attenuation that partially degrades object boundaries while maintaining longer detection ranges. This characteristic makes the physically-inspired modeling of gradual attenuation patterns particularly effective. The improvements in light fog conditions reflect the complementary contributions of both simulation strategies. Traffic signs achieve the most substantial improvement (35.7% IoU, 15.0% above WADG) through the phenomenological simulation that maintains surface responses under moderate attenuation. The Mie-Monte Carlo scattering simulation enables robust segmentation of vulnerable road users under partial visibility degradation, with person segmentation reaching 54.4% IoU (10.6% above WADG). Poles benefit from the distance-dependent attenuation modeling, achieving 31.5% IoU (7.3% above WADG). The phenomenological weather simulation preserves ground surface perception through

its distance-dependent attenuation modeling, with road segmentation reaching 74.0% IoU. Trunk segmentation reaches 47.3% IoU as the gradual attenuation modeling preserves cylindrical structure features at extended ranges. These results indicate that PhyDAWS maintains effectiveness across varying fog densities, and the consistent improvements in traffic sign segmentation, pedestrian segmentation, and road segmentation directly support core autonomous driving functions under degraded visibility conditions.

Table 4 reports the segmentation performance under rain conditions, where PhyDAWS achieves an overall mIoU of 37.7%, significantly outperforming WADG by 6.0% and PointDR by 5.8%. Rain presents unique challenges for LiDAR perception due to the combination of anisotropic occlusion from falling raindrops and specular reflections from wet surfaces. These phenomena create complex signal distortions that differ fundamentally from the uniform attenuation observed in fog conditions. The angle-dependent occlusion modeling in our phenomenological rain simulation directly addresses these challenges. Person segmentation achieves the most substantial improvement (52.9% IoU, 12.3% above WADG), as the simulation accurately captures directional occlusion patterns from falling raindrops. Building segmentation reaches 82.0% IoU through the preservation of structural boundaries despite rain-induced signal degradation. The Mie-Monte Carlo scattering simulation further enhances vehicle segmentation by modeling specular reflections from wet surfaces, with cars achieving the highest absolute performance (86.9% IoU) and other-vehicles showing remarkable improvement (26.9% IoU). Trunk segmentation benefits from the accurate representation of cylindrical objects under precipitation conditions, reaching 27.5% IoU. These improvements in rain perception are particularly valuable for autonomous driving safety, as precipitation conditions challenge sensor reliability and reduce effective perception range.

Table 5 summarizes the results under snow conditions, where PhyDAWS achieves an overall mIoU of 33.4%, substantially outperforming WADG by 4.0% and PointDR by 7.2%. Snow presents distinct challenges that combine airborne particle scattering with geometric alterations from surface accumulation. These dual effects require physically-inspired modeling that addresses both signal degradation and structural changes simultaneously. The height-dependent snow density variation

Table 4

3D semantic segmentation results (%) with SemanticKITTI as source and SemanticSTF as target under rain conditions only. Per-class IoU for 19 semantic categories in rain and the mIoU across these categories are reported. Dash (–) indicates categories with no instances in the rain scenario.

Method	Car	Bicycle	Motorcycle	Truck	Other-Vehicle	Person	Bicyclist	Motorcyclist	Road	Parking	Sidewalk	Other-Ground	Building	Fence	Vegetation	Trunk	Terrain	Pole	Traffic-Sign	Rain
	Per-Class IoU (%)																			
Baseline	72.4	0.0	–	0.0	16.3	6.9	0.0	–	71.6	12.7	58.1	0.0	70.0	33.0	51.8	9.9	24.2	33.3	22.9	28.4
Dropout (Srivastava et al., 2014)	81.3	0.0	–	0.0	21.2	5.6	0.0	–	62.2	11.8	44.8	0.6	76.8	44.7	56.0	16.3	23.3	32.8	22.2	29.4
Perturbation	83.9	0.0	–	2.4	0.0	20.9	0.0	–	73.2	12.6	54.7	7.0	71.7	43.2	58.3	5.9	29.4	29.4	16.9	30.0
PolarMix (Xiao et al., 2022a)	56.7	4.0	–	9.1	1.5	29.8	0.0	–	68.2	10.9	50.2	0.5	73.2	47.2	48.3	17.8	22.3	32.3	34.1	28.6
MMD (Li et al., 2018)	83.9	0.0	–	0.0	8.9	31.6	0.0	–	77.9	17.9	60.2	0.3	69.6	39.3	58.4	14.1	32.5	34.0	30.0	32.8
PCL (Yao et al., 2022)	84.2	0.0	–	0.0	0.1	4.3	0.0	–	68.1	10.9	55.5	4.6	74.7	43.9	59.6	5.8	27.3	34.2	38.8	30.1
PointDR (Xiao et al., 2023)	78.0	0.0	–	0.0	13.8	20.0	0.0	–	72.1	14.7	60.0	1.2	76.1	36.9	58.0	18.3	24.7	36.1	32.5	31.9
WADG (Du et al., 2024)	84.3	0.0	–	0.4	0.0	40.6	0.0	–	65.1	19.2	51.3	4.2	77.3	42.7	61.3	15.7	22.8	30.9	23.5	31.7
Ours	86.9	0.0	–	0.8	26.9	52.9	0.0	–	76.0	21.4	56.7	5.3	82.0	49.3	59.1	27.5	33.3	35.9	27.7	37.7

Table 5

3D semantic segmentation results (%) with SemanticKITTI as source and SemanticSTF as target under snow conditions only. Per-class IoU for 19 semantic categories in snow and the mIoU across these categories are reported. Dash (–) indicates categories with no instances in the snow scenario.

Method	Car	Bicycle	Motorcycle	Truck	Other-Vehicle	Person	Bicyclist	Motorcyclist	Road	Parking	Sidewalk	Other-Ground	Building	Fence	Vegetation	Trunk	Terrain	Pole	Traffic-Sign	Snow
	Per-Class IoU (%)																			
Baseline	49.5	0.0	0.3	0.5	11.6	10.8	–	–	42.1	14.9	23.9	0.0	71.5	26.7	29.3	24.0	17.8	30.8	10.1	21.4
Dropout (Srivastava et al., 2014)	58.5	0.0	30.5	5.4	13.2	5.2	–	–	41.9	18.0	20.4	2.5	76.4	30.5	31.8	32.7	19.8	28.2	7.0	24.8
Perturbation	73.6	0.0	0.0	5.5	1.1	19.8	–	–	45.7	10.9	34.4	0.1	80.6	32.8	45.2	12.8	20.0	24.4	9.5	24.5
PolarMix (Xiao et al., 2022a)	66.5	3.4	9.3	3.5	5.8	32.4	–	–	55.3	3.6	30.1	0.1	77.8	36.1	34.2	12.6	25.1	29.8	10.1	25.6
MMD (Li et al., 2018)	59.4	0.0	4.7	0.0	14.7	30.5	–	–	50.8	16.9	32.8	0.2	68.4	24.4	36.6	24.1	24.1	30.0	11.4	25.2
PCL (Yao et al., 2022)	64.0	0.0	0.0	8.2	0.7	9.2	–	–	38.9	15.2	31.6	2.3	79.6	35.1	41.3	11.2	23.1	30.1	26.8	24.6
PointDR (Xiao et al., 2023)	66.2	0.0	10.4	0.0	16.7	21.3	–	–	43.0	15.2	33.0	1.7	76.8	30.3	36.1	27.6	22.2	30.0	14.1	26.2
WADG (Du et al., 2024)	70.6	0.1	39.8	8.7	0.9	39.4	–	–	39.8	13.8	27.4	3.4	80.6	36.6	46.8	30.6	26.0	23.5	12.0	29.4
Ours	61.2	1.2	65.3	11.8	11.2	45.8	–	–	45.6	23.1	30.6	3.0	82.8	37.7	39.0	38.8	24.6	28.5	17.8	33.4

modeling in our phenomenological simulation provides the foundation for improved performance. Buildings achieve the highest absolute performance (82.8% IoU) as the vertical density gradient modeling preserves structural boundaries despite varying snow concentrations at different heights. Motorcycles show the most dramatic improvement (65.3% IoU, 25.5% above WADG), benefiting from the accurate representation of how thin structures interact with height-varying particle distributions. Person segmentation reaches 45.8% IoU through modeling of snow particle effects on the person class at varying heights. The snow accumulation simulation on horizontal surfaces contributes to ground-related class improvements. Parking areas achieve 23.1% IoU (9.3% above WADG) as the accumulation modeling maintains semantic boundaries despite apparent geometric changes from snow buildup. Trunk segmentation benefits from the combined effects of height-dependent attenuation and accumulation modeling, reaching 38.8% IoU. These improvements support more reliable perception during winter driving scenarios, where snow accumulation fundamentally alters both object appearance and scene geometry.

4.4.2. Visualization results: from SemanticKITTI to SemanticSTF

Figs. 10, 11, 12, and 13 provide qualitative comparisons of semantic segmentation results on the SemanticSTF target domain, using models trained solely on SemanticKITTI. These figures showcase performance under dense fog, light fog, rain, and snow conditions, respectively. Each figure compares the ground truth (GT) with predictions from PointDR (Xiao et al., 2023), WADG (Du et al., 2024), and our proposed PhyDAWS method. Additionally, visualizations highlighting correct and incorrect predictions for each method are provided.

Specifically, in dense fog conditions (Fig. 10), PhyDAWS demonstrates enhanced capabilities in segmenting large structures like road and terrain surfaces, maintaining better structural integrity and class consistency compared to the baseline methods, which often produce more fragmented or noisy results. In light fog (Fig. 11), PhyDAWS achieves superior structural coherence and classification accuracy compared to the often fragmented or erroneous results from PointDR and WADG for building structures, as evidenced by cleaner predictions and fewer errors in the associated map. Additionally, the segmentation of the road surface is rendered more accurately and with fewer

misclassified points by PhyDAWS compared to the baseline methods. During rain (Fig. 12), PhyDAWS more accurately segments cars and road surfaces, indicating better robustness to rain-induced effects like occlusion and potentially specular reflections, areas where baselines show more confusion. Under snow conditions (Fig. 13), our method excels at identifying road and vegetation boundaries, effectively handling the geometric alterations and point cloud sparsity caused by snowfall and accumulation, while competing methods often misclassify large regions or lose definition.

These qualitative results strongly support the quantitative findings, illustrating how the physically-inspired data augmentation strategy employed by PhyDAWS enables the model to learn weather-invariant features. By simulating the physical interactions of LiDAR with various weather phenomena during training on clear-weather data, our model achieves significantly improved generalization and robustness when deployed in unseen adverse weather scenarios.

4.4.3. Evaluation results: from SynLiDAR to SemanticSTF

Table 6 evaluates the synthetic-to-real domain generalization from SynLiDAR to SemanticSTF, where PhyDAWS achieves an overall mIoU of 23.7%, outperforming WADG by 1.8%, UniMix by 0.3%, and PointDR by 5.2%. This transfer task presents compounded challenges: the synthetic-to-real domain gap combines with adverse weather degradation, requiring augmentation strategies that address both distribution shifts simultaneously. PhyDAWS achieves competitive performance across all four weather conditions (dense fog 22.8%, light fog 23.1%, rain 28.2%, snow 24.4%), with the best results in rain and snow, demonstrating that physically-inspired augmentation enhances transferability even when source data originates from simulation environments. The physically-inspired modeling strategies prove particularly effective for bridging the synthetic-to-real gap. Cars show the largest improvement (48.1% IoU, 14.3% above WADG) as the phenomenological simulation preserves vehicle geometries that transfer effectively across domain differences. Buildings achieve the highest absolute performance (71.1% IoU, 6.8% above WADG), as the weather simulation preserves structural boundaries across domains. Person segmentation achieves 35.6% IoU (8.8% above WADG) as the Mie-Monte Carlo scattering simulation accurately models signal interactions

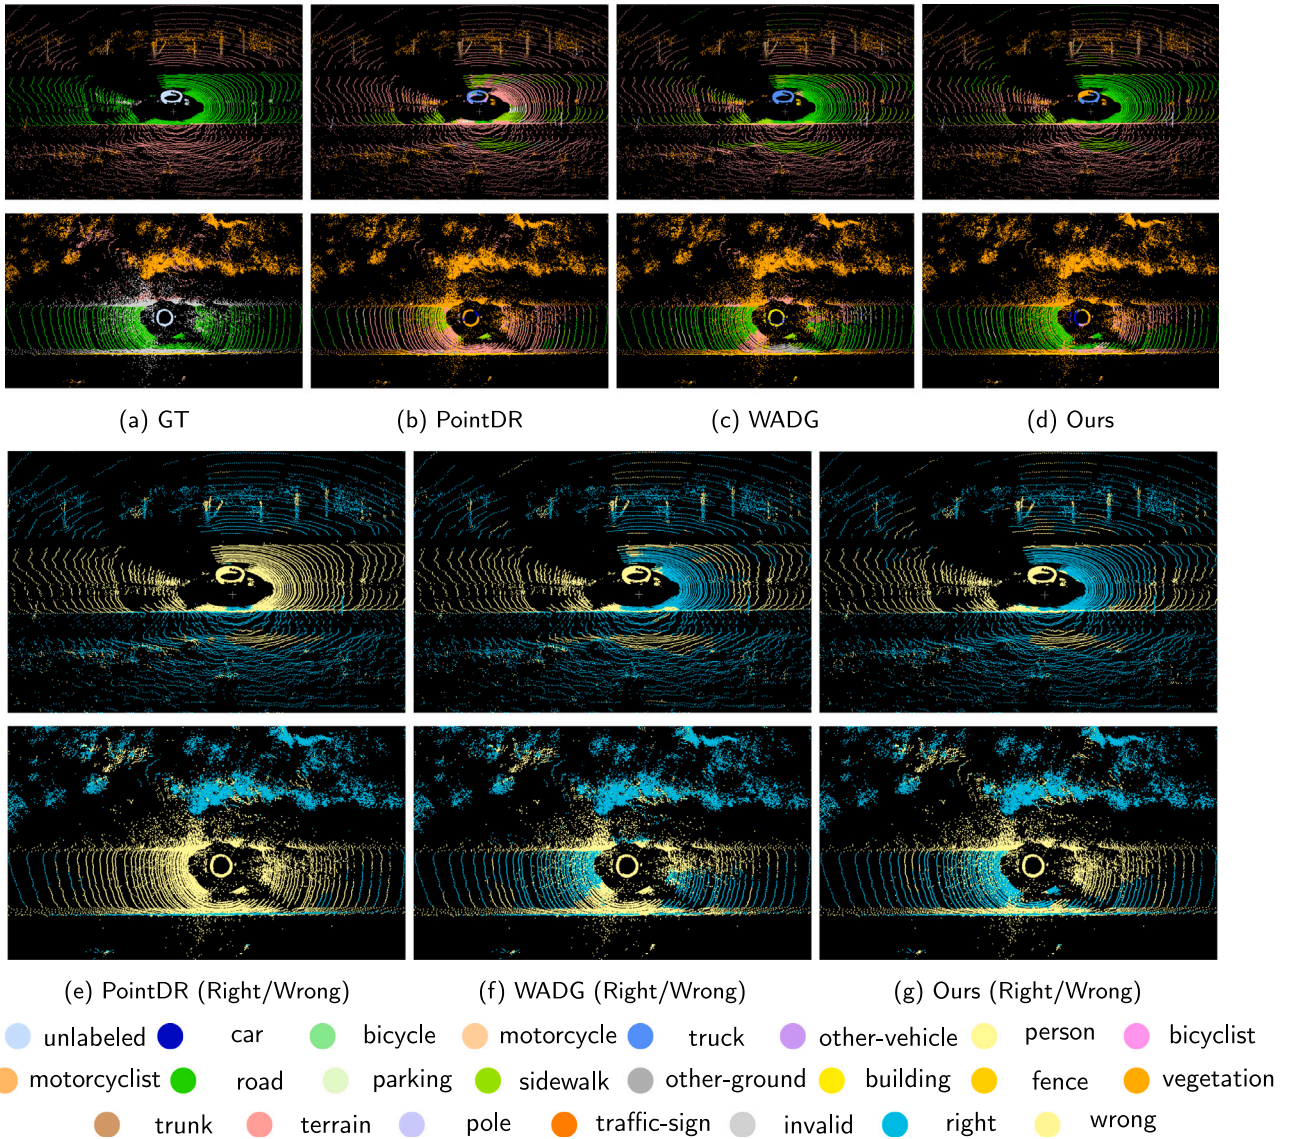


Fig. 10. Qualitative results with SemanticKITTI as source and SemanticSTF as target (dense fog). Cyan areas (Right) indicate correctly segmented regions, while yellow areas (Wrong) represent erroneously segmented regions.

Table 6

3D semantic segmentation results (%) with SynLiDAR as source and SemanticSTF as target. Per-class IoU for 19 semantic categories across all weather conditions; per-weather mIoU for those 19 categories under four weather scenarios; and overall all-weather mIoU for the same 19 categories across all weather conditions are reported.

Method	Per-Class IoU (%)																		Per-Weather mIoU (%)				All-Weather mIoU (%)	
	Car	Bicycle	Motorcycle	Truck	Other-Vehicle	Person	Bicyclist	Motorcyclist	Road	Parking	Sidewalk	Other-Ground	Building	Fence	Vegetation	Trunk	Terrain	Pole	Traffic-Sign	Dense-Fog	Light-Fog	Rain		Snow
Baseline	27.1	3.0	0.6	15.8	0.1	25.2	1.8	5.6	23.9	0.3	14.6	0.6	36.3	19.9	37.9	17.9	41.8	9.5	2.3	16.9	17.2	17.2	11.9	15.0
Dropout (Srivastava et al., 2014)	28.0	3.0	1.4	9.6	0.0	17.1	0.8	0.7	34.2	6.8	19.1	0.1	35.5	19.1	42.3	17.6	36.0	14.0	2.8	15.3	16.6	20.4	14.0	15.2
Perturbation	27.1	2.3	2.3	16.0	0.1	23.7	1.2	4.0	27.0	3.6	16.2	0.8	29.2	16.7	35.3	22.7	38.3	17.9	5.1	16.3	16.7	19.3	13.4	15.2
PolarMix (Xiao et al., 2022a)	39.2	1.1	1.2	8.3	1.5	17.8	0.8	0.7	23.3	1.3	17.5	0.4	45.2	24.8	46.2	20.1	38.7	7.6	1.9	16.1	15.5	19.2	15.6	15.7
MMD (Li et al., 2018)	25.5	2.3	2.1	13.2	0.7	22.1	1.4	7.5	30.8	0.4	17.6	0.2	30.9	19.7	37.6	19.3	43.5	9.9	2.6	17.3	16.3	20.0	12.7	15.1
PCL (Yao et al., 2022)	30.9	0.8	1.4	10.0	0.4	23.3	4.0	7.9	28.5	1.3	17.7	1.2	39.4	18.5	40.0	16.0	38.6	12.1	2.3	17.8	16.7	19.3	14.1	15.5
PointDR (Xiao et al., 2023)	37.8	2.5	2.4	23.6	0.1	26.3	2.2	3.3	27.9	7.7	17.5	0.5	47.6	25.3	45.7	21.0	37.5	17.9	5.5	19.5	19.9	21.1	16.9	18.5
UniMix (Zhao et al., 2024)	65.4	0.1	3.9	16.9	5.3	32.3	2.0	19.3	52.1	5.0	27.3	3.0	49.4	20.3	58.5	22.7	23.2	26.9	10.4	24.3	22.9	26.1	20.9	23.4
WADG (Du et al., 2024)	33.8	1.1	2.9	17.0	0.2	26.8	1.0	4.3	53.9	5.0	20.6	2.2	64.3	27.1	53.8	27.0	37.0	28.6	8.6	21.6	23.4	27.2	21.4	21.9
Ours	48.1	3.9	2.1	18.9	0.7	35.6	3.5	5.7	38.9	9.9	20.7	0.9	71.1	32.3	55.1	27.1	40.7	23.1	11.7	22.8	23.1	28.2	24.4	23.7

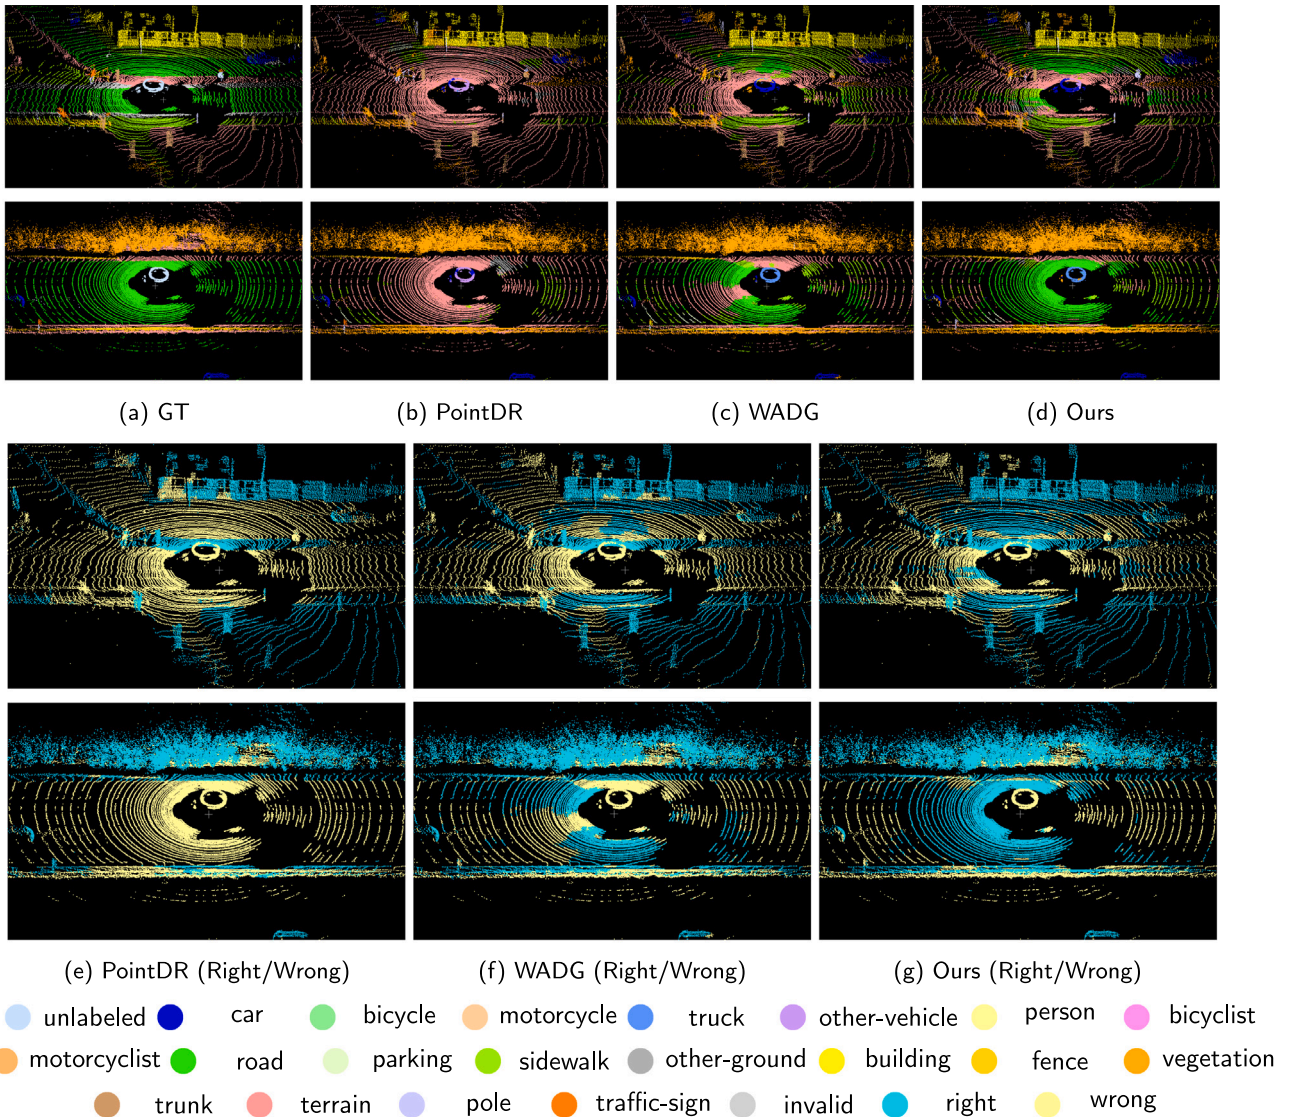


Fig. 11. Qualitative results with SemanticKITTI as source and SemanticSTF as target (light fog). Cyan areas (Right) indicate correctly segmented regions, while yellow areas (Wrong) represent erroneously segmented regions.

Table 7

3D semantic segmentation results (%) with SynLiDAR as source and SemanticSTF as target under dense fog conditions only. Per-class IoU for 19 semantic categories in dense fog and the mIoU across these categories are reported. Dash (-) indicates categories with no instances in the dense fog scenario.

Method	Car	Bicycle	Motorcycle	Truck	Other-Vehicle	Person	Bicyclist	Motorcyclist	Road	Parking	Sidewalk	Other-Ground	Building	Fence	Vegetation	Trunk	Terrain	Pole	Traffic-Sign	Dense-Fog mIoU (%)
Baseline	21.6	-	-	6.4	0.0	3.7	2.9	18.9	25.7	0.0	7.7	1.0	41.2	22.5	52.3	15.4	55.5	9.3	2.4	16.9
Dropout (Srivastava et al., 2014)	12.7	-	-	7.7	0.0	1.9	0.4	2.5	38.3	0.1	10.2	0.3	37.3	21.8	57.4	13.1	44.5	10.1	1.0	15.3
Perturbation	13.3	-	-	10.4	0.0	4.3	2.8	19.1	30.0	0.7	8.8	1.2	30.5	17.5	48.9	18.4	50.3	16.3	5.2	16.3
PolarMix (Xiao et al., 2022a)	15.8	-	-	10.6	0.0	1.5	1.7	3.5	27.7	0.0	9.9	0.3	46.2	28.9	59.2	13.5	49.5	4.4	1.7	16.1
MMD (Li et al., 2018)	26.5	-	-	12.7	0.0	2.7	4.0	22.3	30.6	0.0	9.4	0.0	31.6	21.7	52.6	13.9	54.3	8.9	2.5	17.3
PCL (Yao et al., 2022)	22.9	-	-	20.1	0.0	2.2	6.2	28.3	29.0	0.0	9.2	2.6	37.9	22.9	54.5	11.4	45.9	8.5	1.1	17.8
PointDR (Xiao et al., 2023)	42.5	-	-	16.6	0.0	2.4	3.2	12.2	31.9	0.2	9.0	0.8	42.8	27.1	59.8	18.3	44.0	15.4	5.7	19.5
WADG (Du et al., 2024)	24.4	-	-	13.1	0.0	3.5	0.0	19.9	52.0	0.1	10.5	0.7	57.3	31.7	61.7	20.6	42.0	21.5	7.8	21.6
Ours	32.5	-	-	13.8	0.0	3.4	7.5	20.0	39.7	0.0	11.1	1.8	61.6	36.8	62.7	21.7	45.6	18.5	11.5	22.8

under attenuation. Fences reach 32.3% IoU through the class-aware reflection modeling that captures boundary characteristics regardless of rendering differences. These results indicate that physically-inspired augmentation not only improves weather robustness but also enhances the practical deployment of simulation-trained models, reducing the cost and safety risks associated with real-world adverse weather data collection.

Table 7 provides the dense fog results for the SynLiDAR to SemanticSTF transfer, where PhyDAWS achieves a mIoU of 22.8%, surpassing WADG by 1.2% and PointDR by 3.3%. This scenario combines the synthetic-to-real distribution shift with severe visibility constraints from dense fog, representing one of the most demanding evaluation settings. The distance-dependent exponential attenuation modeling in our fog simulation proves essential for this transfer task, as it captures visibility degradation patterns that generalize across synthetic and

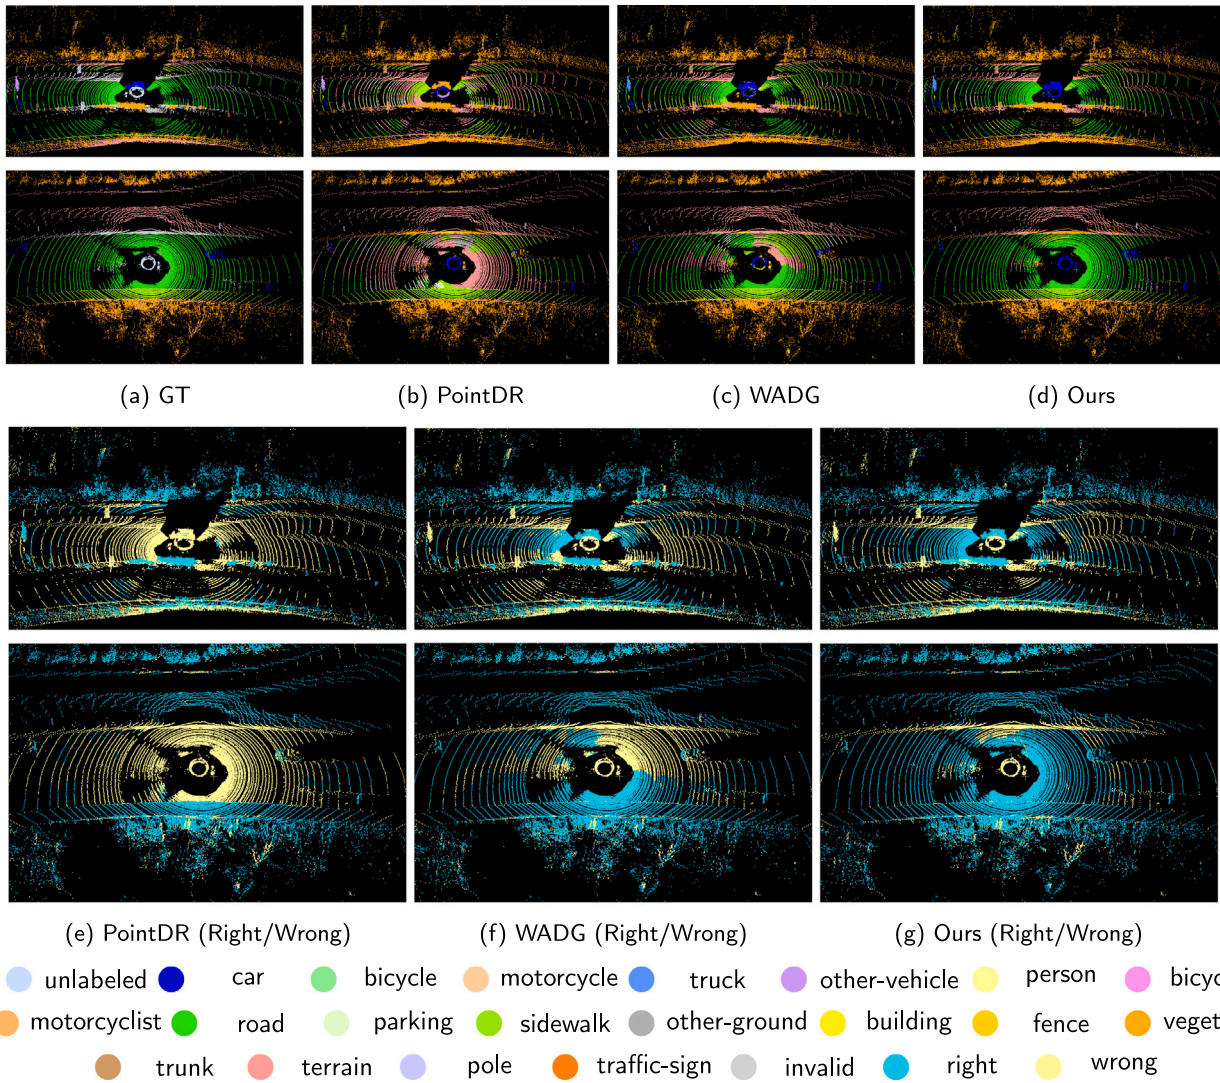


Fig. 12. Qualitative results with SemanticKITTI as source and SemanticSTF as target (rain). Cyan areas (Right) indicate correctly segmented regions, while yellow areas (Wrong) represent erroneously segmented regions.

Table 8

3D semantic segmentation results (%) with SynLiDAR as source and SemanticSTF as target under light fog conditions only. Per-class IoU for 19 semantic categories in light fog and the mIoU across these categories are reported.

Method	Per-Class IoU (%)																			Light-Fog mIoU (%)
	Car	Bicycle	Motorcycle	Truck	Other-Vehicle	Person	Bicyclist	Motorcyclist	Road	Parking	Sidewalk	Other-Ground	Building	Fence	Vegetation	Trunk	Terrain	Pole	Traffic-Sign	
Baseline	32.0	4.2	0.5	27.3	0.2	14.0	6.2	0.0	31.0	0.0	12.6	0.9	38.7	24.8	51.5	26.7	46.4	8.5	1.3	17.2
Dropout (Srivastava et al., 2014)	22.5	3.0	0.9	16.0	0.1	10.0	5.2	0.2	40.3	1.3	18.1	0.0	38.9	22.1	57.6	23.5	38.5	13.8	3.7	16.6
Perturbation	31.1	1.9	1.6	21.5	0.0	12.5	2.6	0.0	33.2	1.6	14.3	1.1	34.3	20.1	48.7	29.8	42.0	16.7	4.5	16.7
PolarMix (Xiao et al., 2022a)	27.3	0.3	0.4	8.9	1.4	8.2	1.2	0.0	29.0	0.2	15.5	0.7	39.9	27.4	57.3	28.8	40.9	5.8	1.5	15.5
MMD (Li et al., 2018)	31.0	2.1	0.5	16.0	0.0	10.5	1.7	0.0	37.7	0.3	16.3	0.6	29.2	24.9	51.8	29.6	47.8	8.3	1.8	16.3
PCL (Yao et al., 2022)	31.7	0.7	0.8	10.1	0.1	10.2	21.6	0.0	33.9	0.6	16.1	0.1	37.8	22.2	52.5	23.8	42.6	11.3	2.2	16.7
PointDR (Xiao et al., 2023)	44.7	1.7	1.0	33.9	0.3	12.9	4.7	0.0	36.0	0.9	15.8	0.7	44.4	30.3	60.0	28.3	42.4	15.1	5.7	19.9
WADG (Du et al., 2024)	39.8	0.4	1.2	33.0	0.1	17.8	5.3	0.0	59.7	3.8	21.3	5.2	52.7	31.2	63.2	33.4	40.1	27.1	8.6	23.4
Ours	49.5	3.0	0.0	30.0	0.1	17.4	7.1	0.0	44.6	0.9	18.0	0.4	60.3	36.0	63.1	32.3	44.5	20.5	11.9	23.1

real domains. Cars show the largest improvement (32.5% IoU, 8.1% above WADG) as the attenuation modeling preserves vehicle geometries under severe visibility constraints. Buildings achieve strong performance (61.6% IoU, 4.3% above WADG) as the fog simulation preserves structural boundaries despite severe signal attenuation. Fences reach 36.8% IoU (5.1% above WADG) through the class-aware reflection modeling that maintains boundary characteristics across domain differences. Traffic signs achieve 11.5% IoU (3.7% above WADG) as the physically-inspired attenuation modeling preserves surface responses.

These improvements demonstrate that physically-inspired augmentation effectively addresses the compounded challenges of synthetic-to-real transfer under dense fog, enabling reliable infrastructure segmentation and obstacle segmentation in severely degraded visibility conditions.

Table 8 shows the light fog results for the SynLiDAR to SemanticSTF transfer, where PhyDAWS achieves an mIoU of 23.1%. While comparable to WADG (23.4%), this represents a significant 3.2% gain over PointDR. Light fog in the synthetic-to-real setting introduces moderate




Fig. 13. Qualitative results with SemanticKITTI as source and SemanticSTF as target (snow). Cyan areas (Right) indicate correctly segmented regions, while yellow areas (Wrong) represent erroneously segmented regions.

Table 9

3D semantic segmentation results (%) with SynLiDAR as source and SemanticSTF as target under rain conditions only. Per-class IoU for 19 semantic categories in rain and the mIoU across these categories are reported. Dash (-) indicates categories with no instances in the rain scenario.

Method	Per-Class IoU (%)																			Rain mIoU (%)
	Car	Bicycle	Motorcycle	Truck	Other-Vehicle	Person	Bicyclist	Motorcyclist	Road	Parking	Sidewalk	Other-Ground	Building	Fence	Vegetation	Trunk	Terrain	Pole	Traffic-Sign	
Baseline	45.8	4.5	-	6.8	0.4	38.9	0.0	-	32.0	0.0	24.3	0.0	43.0	8.0	33.8	11.3	23.9	11.5	7.7	17.2
Dropout (Srivastava et al., 2014)	47.0	7.6	-	7.7	0.0	34.0	0.0	-	47.3	6.9	34.6	0.0	39.8	11.5	37.5	13.8	29.6	21.6	8.6	20.4
Perturbation	57.5	5.3	-	18.2	0.0	36.3	0.1	-	37.1	1.5	26.9	0.3	34.9	10.4	32.6	12.2	20.5	23.2	10.4	19.3
PolarMix (Xiao et al., 2022a)	59.6	1.5	-	6.0	5.2	24.6	1.0	-	31.4	0.1	30.4	0.0	55.5	12.2	44.6	13.1	25.0	11.0	4.7	19.2
MMD (Li et al., 2018)	49.5	4.8	-	20.0	4.7	37.6	0.0	-	43.7	0.0	32.4	0.0	42.1	11.3	34.4	12.3	25.1	13.4	8.1	20.0
PCL (Yao et al., 2022)	51.3	0.9	-	4.3	2.1	35.6	0.0	-	41.4	0.0	32.0	0.0	54.8	9.7	37.1	11.4	24.2	16.6	6.3	19.3
PointDR (Xiao et al., 2023)	42.2	3.3	-	21.9	0.0	30.4	1.7	-	35.8	3.2	31.9	0.0	54.0	14.4	40.7	12.5	31.9	23.6	11.8	21.1
WADG (Du et al., 2024)	61.8	1.5	-	22.0	0.0	38.4	0.0	-	65.1	1.2	31.2	4.1	68.7	20.5	54.3	14.5	29.7	32.3	17.2	27.2
Ours	63.4	1.1	-	12.7	1.3	45.8	0.0	-	50.1	7.6	35.9	0.2	76.9	24.9	52.8	25.9	34.5	32.9	13.9	28.2


signal degradation that tests whether physically-inspired augmentation can bridge domain gaps under partial visibility reduction. The distance-dependent attenuation modeling captures gradual signal decay patterns that transfer effectively from synthetic training environments. Buildings achieve robust performance (60.3% IoU, 7.6% above WADG) as the fog simulation preserves large structural boundaries across domain differences. Fences reach 36.0% IoU through the class-aware reflection

modeling that maintains boundary characteristics despite the synthetic-to-real gap. Cars achieve 49.5% IoU (9.7% above WADG) as the attenuation modeling preserves vehicle geometries at varying detection ranges. Traffic signs reach 11.9% IoU through the phenomenological simulation that maintains surface responses under moderate attenuation. Terrain benefits from the fog simulation with 44.5% IoU (4.4% above WADG), as the distance-dependent modeling preserves ground surface characteristics. These results confirm that physically-inspired


Table 10

3D semantic segmentation results (%) with SynLiDAR as source and SemanticSTF as target under snow conditions only. Per-class IoU for 19 semantic categories in snow and the mIoU across these categories are reported. Dash (–) indicates categories with no instances in the snow scenario.

Method	Car	Bicycle	Motorcycle	Truck	Other-Vehicle	Person	Bicyclist	Motorcyclist	Road	Parking	Sidewalk	Other-Ground	Building	Fence	Vegetation	Trunk	Terrain	Pole	Traffic-Sign	Snow
	Per-Class IoU (%)																			
Baseline	24.6	2.7	1.5	2.4	0.0	32.2	–	–	12.9	0.4	18.3	0.0	33.3	13.8	15.7	14.9	18.1	10.1	1.9	11.9
Dropout (Srivastava et al., 2014)	35.9	2.8	3.7	3.0	0.0	21.9	–	–	20.9	10.0	22.8	0.0	33.2	14.8	17.1	16.8	16.5	15.7	2.6	14.0
Perturbation	27.1	2.4	6.8	6.8	0.2	31.0	–	–	15.4	4.8	19.7	0.0	26.3	12.4	14.0	22.0	16.4	19.0	4.1	13.4
PolarMix (Xiao et al., 2022a)	53.4	2.3	4.1	6.0	1.2	27.9	–	–	11.7	1.9	21.5	0.3	45.2	20.8	21.7	18.8	16.5	10.5	1.7	15.6
MMD (Li et al., 2018)	20.8	2.7	6.0	4.8	0.2	31.3	–	–	20.1	0.5	21.0	0.1	29.6	12.2	15.0	16.6	21.8	11.3	2.4	12.7
PCL (Yao et al., 2022)	30.7	1.1	4.4	6.2	0.3	34.6	–	–	19.1	1.7	22.0	0.3	37.8	12.6	16.4	14.2	19.9	14.7	3.0	14.1
PointDR (Xiao et al., 2023)	34.2	4.0	7.4	7.5	0.1	36.2	–	–	13.8	12.0	22.7	0.0	48.8	19.9	19.9	18.9	17.0	20.7	3.4	16.9
WADG (Du et al., 2024)	30.7	1.6	7.9	2.6	0.4	32.2	–	–	46.5	6.1	25.1	0.9	69.9	19.1	30.2	28.0	21.6	33.9	7.5	21.4
Ours	49.5	7.3	12.7	4.7	1.1	46.5	–	–	29.1	14.8	27.8	0.3	76.1	24.9	32.8	27.0	22.6	26.6	11.1	24.4



(a) GT (b) PointDR (c) WADG (d) Ours



(e) PointDR (Right/Wrong) (f) WADG (Right/Wrong) (g) Ours (Right/Wrong)



Fig. 14. Qualitative results with SynLiDAR as source and SemanticSTF as target (dense fog). Cyan areas (Right) indicate correctly segmented regions, while yellow areas (Wrong) represent erroneously segmented regions.

augmentation maintains effectiveness across the spectrum from severe to moderate visibility degradation in synthetic-to-real transfer scenarios.

Table 9 details the rain results for the SynLiDAR to SemanticSTF transfer, where PhyDAWS achieves an overall mIoU of 28.2%, outperforming WADG by 1.0% and PointDR by 7.1%. The most pronounced

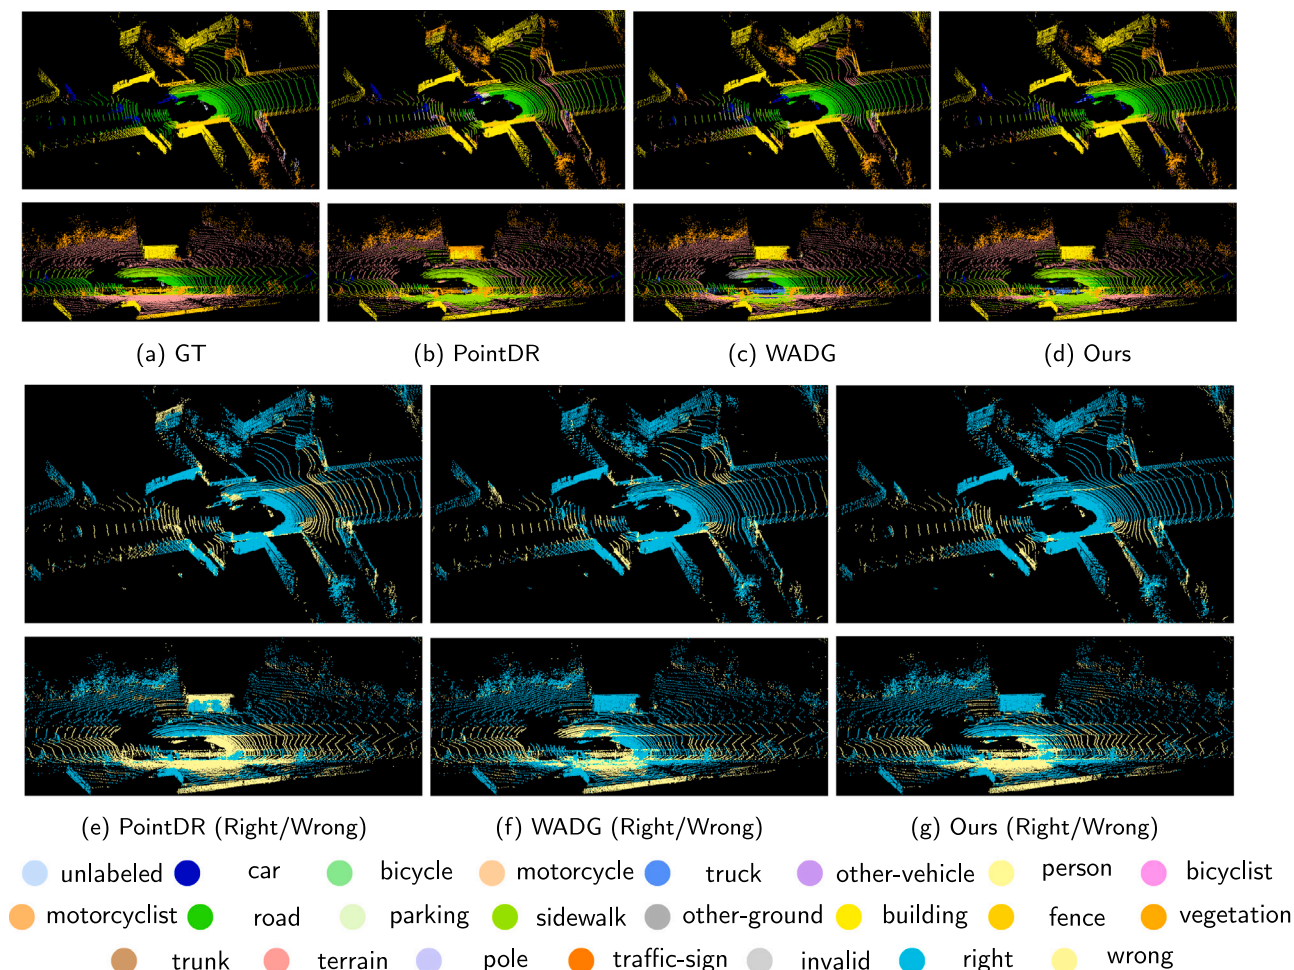


Fig. 15. Qualitative results with SynLiDAR as source and SemanticSTF as target (light fog). Cyan areas (Right) indicate correctly segmented regions, while yellow areas (Wrong) represent erroneously segmented regions.

improvements appear in classes affected by rain-specific signal distortions. Trunk segmentation benefits substantially from the rain intensity attenuation modeling, achieving 25.9% IoU with an 11.4% gain over WADG, as the simulation captures how precipitation affects cylindrical structures at varying distances. Buildings reach 76.9% IoU (8.2% above WADG) through the phenomenological simulation that preserves structural boundaries despite rain-induced signal degradation. The angle-dependent occlusion modeling proves effective for person segmentation (45.8% IoU, 7.4% above WADG), accurately segmenting persons partially obscured by falling raindrops. Parking areas show notable relative improvement (7.6% IoU, 6.4% above WADG) as the rain simulation models ground surface responses under wet conditions. These gains in the synthetic-to-real rain scenario indicate that physically-inspired augmentation captures precipitation effects that generalize across the domain gap between simulated and real-world environments.

Table 10 reports the snow results for the SynLiDAR to SemanticSTF transfer, where PhyDAWS achieves 24.4% mIoU, exceeding WADG by 3.0% and PointDR by 7.5%. Snow in the synthetic-to-real setting presents dual challenges: airborne particle scattering combined with surface accumulation that alters object geometries. These effects compound the domain gap as snow appearance varies significantly between synthetic rendering and real-world capture. The height-dependent snow density variation modeling proves particularly effective for this transfer task, as vertical density gradients exhibit consistent physical behavior across domains. Buildings achieve the highest absolute performance

(76.1% IoU, 6.2% above WADG) through the preservation of structural boundaries despite varying snow concentrations at different heights. Person segmentation reaches 46.5% IoU (14.3% above WADG) as the snow simulation accurately segments persons partially affected by accumulation. Cars achieve 49.5% IoU (18.8% above WADG) through the modeling of vehicle surfaces under snow conditions. The snow accumulation simulation on horizontal surfaces enhances ground-related perception, with parking areas achieving 14.8% IoU (8.7% above WADG). These improvements enable reliable perception during winter driving scenarios, which represent a significant proportion of weather-related incidents in northern climates.

4.4.4. Visualization results: from SynLiDAR to SemanticSTF

Figs. 14, 15, 16, and 17 present qualitative results for the more challenging domain generalization task from the synthetic SynLiDAR source domain to the real-world adverse weather SemanticSTF target domain. The figures illustrate performance under dense fog, light fog, rain, and snow, respectively, comparing the GT against predictions from PointDR (Xiao et al., 2023), WADG (Du et al., 2024), and our PhyDAWS method, along with Right/Wrong prediction visualizations.

Bridging the synthetic-to-real gap, especially under adverse weather, poses significant challenges, as reflected in the overall lower visual quality compared to the SemanticKITTI-to-SemanticSTF transfer. However, even in this demanding setting, PhyDAWS consistently demonstrates qualitative advantages over the baseline methods. Our approach

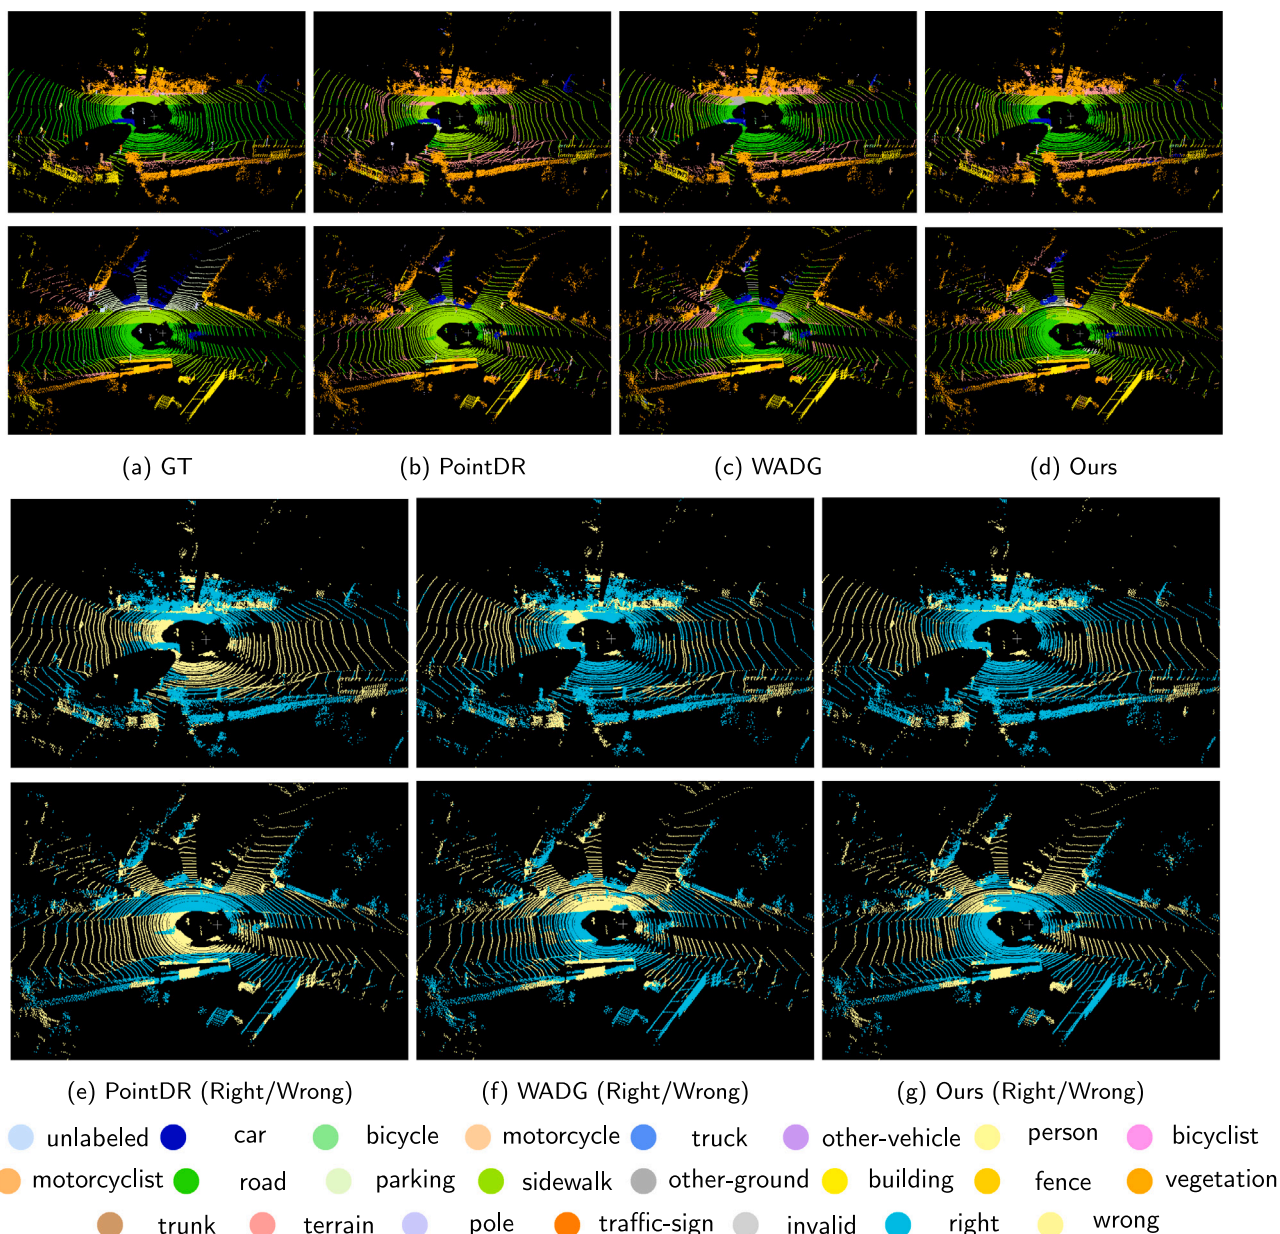


Fig. 16. Qualitative results with SynLiDAR as source and SemanticSTF as target (rain). Cyan areas (Right) indicate correctly segmented regions, while yellow areas (Wrong) represent erroneously segmented regions.

generally produces more coherent and structurally accurate segmentation maps, better preserving the shapes of objects and reducing large-scale misclassifications. These visualizations highlight the effectiveness of the physically-inspired augmentation in PhyDAWS, not only for real-to-real transfer but also for the challenging synthetic-to-real scenario under adverse weather. By grounding the augmentation in physical principles, our method learns representations that are more robust to both the synthetic-to-real domain shift and the variations introduced by different weather conditions, leading to qualitatively superior segmentation results compared to existing state-of-the-art approaches.

4.5. Ablation study

4.5.1. Ablation study on weather simulation

To assess the individual contributions of simulating distinct meteorological conditions within the proposed PhyDAWS framework, a systematic ablation study was conducted by selectively excluding the

simulation components corresponding to one specific weather type at a time. For each ablation setting, both the Phenomenological Weather Simulation (PWS) and the Mie-Monte Carlo Scattering Simulation (MSS) associated with that weather condition were removed, while the simulations for the remaining three conditions were retained. The evaluation was carried out on the SemanticKITTI to SemanticSTF domain generalization task, and the quantitative results are summarized in Table 11. When all weather simulations are enabled, PhyDAWS achieves an overall mIoU of 38.9% across target-domain conditions. Removing any single weather type results in consistent performance reduction, with the overall mIoU decreasing to 36.3% without dense fog, 36.8% without light fog, and 36.2% without rain or snow simulation. The narrow degradation range of 2.1% to 2.7% indicates that each weather-specific simulation contributes positively to generalization, and that robust performance benefits from comprehensive exposure to diverse atmospheric phenomena during training.

Condition-specific analysis reveals non-uniform interactions among weather simulations that reflect their underlying physical relationships.

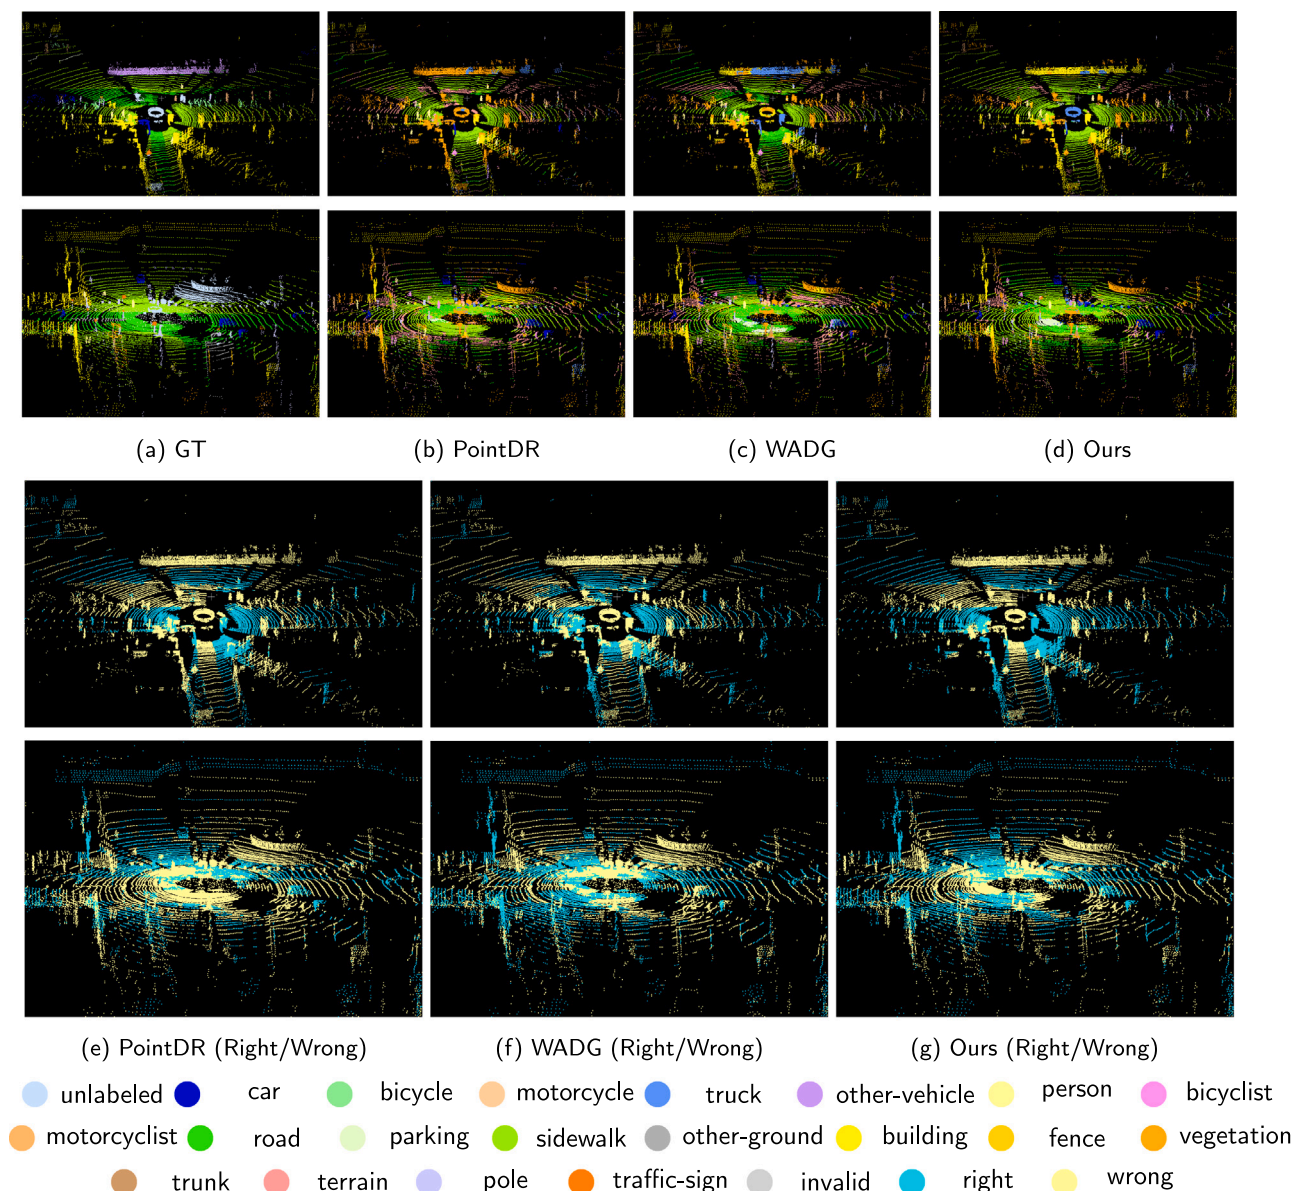


Fig. 17. Qualitative results with SynLiDAR as source and SemanticSTF as target (snow). Cyan areas (Right) indicate correctly segmented regions, while yellow areas (Wrong) represent erroneously segmented regions.

Snow exclusion produces the most pronounced matched-condition degradation, with mIoU in snowy scenes dropping from 33.4% to 29.8%, because snow introduces coupled distortions in point density and local geometric structure that other simulations cannot adequately replicate. This effect is particularly evident for geometrically complex objects, where motorcycle segmentation IoU collapses from 65.3% to 35.1% upon snow simulation removal. In contrast, fog and rain conditions exhibit greater mutual transferability: excluding dense fog simulation slightly improves dense fog performance (40.7% versus 40.2%), and peak scores in light fog (35.3%) and rain (37.9%) are achieved when snow simulation is excluded. These observations suggest that fog and rain share atmospheric attenuation and scattering mechanisms enabling partial compensation, while snow’s distinct physical characteristics make its simulation less interchangeable with others. Nevertheless, the full configuration provides the most reliable performance across all conditions, confirming that explicit modeling of each weather type’s sensing degradation remains essential for robust segmentation under diverse adverse meteorological scenarios.

To further examine the respective roles of different weather simulation strategies within the proposed framework, we conducted an ablation study comparing the use of Phenomenological Weather Simulation (PWS) alone, Mie–Monte Carlo Scattering Simulation (MSS) alone, and their combination in the full PhyDAWS configuration. The evaluation was performed on the SemanticKITTI to SemanticSTF domain generalization task, and the quantitative results are summarized in Table 12. When applied independently, both simulation strategies achieve comparable overall performance, with PWS reaching an all-weather mIoU of 36.1% and MSS achieving 36.3%. Their weather-specific behavior, however, differs noticeably. PWS shows relatively stronger performance in rain conditions, whereas MSS performs more favorably in dense fog scenarios. Both approaches exhibit reduced effectiveness in snow, reflecting the difficulty of fully capturing snow-induced sensing degradations with either method alone.

Combining the two simulation strategies in the full PhyDAWS framework leads to consistent improvement across all target-domain conditions, yielding an overall mIoU of 38.9%. The gains are observed across weather types, with particularly pronounced improvements under snowy conditions, where the combined approach substantially

Table 11

Ablation study on the effects of weather types with SemanticKITTI as source and SemanticSTF as target (%). PWS: Phenomenological Weather Simulation; MSS: Mie-Monte Carlo Scattering Simulation; D-Fog: Dense Fog; L-Fog: Light Fog. Each row represents experiments where both simulation methods excluded the specific weather type. “w/o” denotes “without”.

PWS	MSS	Test Weather	Car		Truck	Other-Vehicle	Person	Bicyclist	Motorcyclist	Road	Parking	Sidewalk	Other-Ground	Building	Fence	Vegetation	Trunk	Terrain	Pole	Traffic-Sign	Dense-Fog	Light-Fog	Rain	Snow	All-Weather		
			Per-Class	IoU (%)																							
w/o D-Fog	w/o D-Fog	D-Fog	80.1	-	-	21.3	0.0	5.3	37.2	64.9	74.1	0.5	34.7	3.9	72.8	56.7	60.9	39.6	60.5	24.7	55.0	40.7				36.3	
		L-Fog	64.3	0.0	0.0	38.2	4.3	43.8	18.3	0.0	72.1	8.1	38.3	8.2	70.6	54.9	59.8	46.1	52.0	29.1	28.6	33.5					
		Rain	87.0	0.0	-	0.8	7.0	49.2	0.0	-	73.8	16.4	57.3	11.1	81.5	46.6	60.9	36.0	30.0	36.3	31.4	36.8					
		Snow	65.3	1.4	45.1	6.8	3.5	37.2	-	-	37.0	19.7	27.6	5.6	83.3	36.6	38.6	37.3	21.9	26.4	24.2	30.4					
		All	68.1	0.9	37.5	25.9	3.9	37.1	11.5	51.5	61.3	16.1	34.4	6.1	79.4	49.4	56.1	40.1	46.2	27.4	36.0						
w/o L-Fog	w/o L-Fog	D-Fog	83.4	-	-	25.4	0.0	7.9	14.9	62.6	76.7	0.2	34.7	1.8	72.9	59.3	64.0	37.2	63.6	24.2	30.0	38.7	34.5				36.8
		L-Fog	72.8	0.0	0.0	36.5	8.0	57.6	0.0	0.0	73.9	9.8	43.3	2.1	66.2	60.0	64.2	44.8	55.4	28.6	32.2	37.0					
		Rain	88.9	0.0	-	0.6	11.1	49.7	0.0	-	79.8	19.1	61.3	2.4	79.5	45.1	60.2	27.0	36.3	33.7	34.9	31.6					
		Snow	71.4	0.6	58.2	15.3	2.5	38.4	-	-	43.2	22.3	30.3	5.0	79.2	33.2	39.0	36.6	22.7	24.5	14.5	36.8					
		All	74.2	0.4	49.3	26.8	4.9	39.3	8.9	46.9	65.1	18.7	37.1	2.9	76.0	50.0	58.6	38.4	49.4	26.2	26.6						
w/o Rain	w/o Rain	D-Fog	80.1	-	-	36.0	0.0	11.4	9.3	81.0	72.5	4.5	33.1	7.6	69.5	57.0	59.2	41.0	49.1	25.1	66.9	41.4	34.0				36.2
		L-Fog	70.4	0.0	0.0	34.1	14.8	46.6	1.3	0.0	71.4	5.4	39.7	5.2	67.7	58.6	58.5	48.0	49.3	29.2	41.3	37.5					
		Rain	87.3	0.0	-	4.2	21.6	46.6	0.0	-	75.6	19.1	60.0	16.1	74.6	45.2	57.0	26.5	28.3	32.7	43.0	30.4					
		Snow	69.2	1.4	21.0	22.9	12.0	51.8	-	-	43.4	18.6	28.3	3.9	71.7	25.5	34.5	34.8	30.3	22.5	25.2	36.2					
		All	71.8	0.9	6.5	29.2	17.5	47.7	4.8	64.5	62.9	14.9	35.0	5.5	71.0	45.3	53.5	39.3	45.1	25.7	46.0						
w/o Snow	w/o Snow	D-Fog	81.0	-	-	21.4	0.0	4.1	7.5	79.6	76.0	0.5	37.3	9.0	68.6	54.7	62.3	38.2	61.5	24.6	56.6	40.2	35.3				36.2
		L-Fog	74.8	0.0	0.0	44.7	14.8	46.6	1.3	0.0	73.1	5.6	44.3	10.2	68.3	56.1	61.5	45.8	53.1	30.9	38.9	37.9					
		Rain	88.5	0.0	-	0.1	21.3	31.7	0.0	-	77.9	20.8	59.9	22.4	81.1	51.1	59.6	20.2	31.8	37.2	41.3	29.8					
		Snow	72.3	1.5	35.1	11.6	9.0	27.2	-	-	42.0	13.9	27.1	8.5	78.2	38.9	34.6	34.6	22.3	27.4	22.4	36.2					
		All	74.9	0.9	17.8	30.7	11.7	27.3	3.2	63.1	64.1	11.8	36.1	10.1	75.4	50.4	55.7	37.7	47.5	28.3	41.1						
Full Model	Full Model	D-Fog	82.3	-	-	21.6	0.0	10.0	13.6	75.1	75.1	0.1	33.1	4.9	69.6	56.6	62.3	37.0	60.1	24.9	56.8	40.2	34.7				38.9
		L-Fog	64.9	0.0	0.0	55.4	3.4	54.4	1.4	0.0	74.0	8.8	39.4	4.2	69.5	56.2	60.4	47.3	52.2	31.5	35.7	37.7					
		Rain	86.9	0.0	-	0.8	26.9	52.9	0.0	-	76.0	21.4	56.7	5.3	82.0	49.3	59.1	27.5	33.3	35.9	27.7	33.4					
		Snow	61.2	1.2	65.3	11.8	11.2	45.8	-	-	45.6	23.1	30.6	3.0	82.8	37.7	39.0	38.8	24.6	28.5	17.8	38.9					
		All	65.9	0.8	51.2	38.1	10.6	45.3	5.5	59.4	65.3	19.0	35.7	4.0	78.5	50.4	56.8	40.0	47.7	29.0	35.2						

Table 12

Ablation study on different weather simulation methods with SemanticKITTI as source and SemanticSTF as target (%). PWS: Phenomenological Weather Simulation; MSS: Mie-Monte Carlo Scattering Simulation; D-Fog: Dense Fog; L-Fog: Light Fog. Each row represents performance on different target domain weather conditions.

PWS	MSS	Test Weather	Car		Truck	Other-Vehicle	Person	Bicyclist	Motorcyclist	Road	Parking	Sidewalk	Other-Ground	Building	Fence	Vegetation	Trunk	Terrain	Pole	Traffic-Sign	Dense-Fog	Light-Fog	Rain	Snow	All-Weather		
			Per-Class	IoU (%)																							
✓		D-Fog	80.1	-	-	23.5	0.0	10.1	9.5	64.9	72.2	0.1	34.5	4.9	73.5	59.8	63.4	40.2	60.8	24.1	55.2	39.8				36.1	
		L-Fog	64.2	0.0	0.0	36.2	0.2	41.9	0.6	0.0	69.8	8.2	40.4	6.3	71.0	61.9	62.4	45.3	54.5	24.7	37.9	32.9					
		Rain	86.7	0.0	-	0.2	1.8	50.1	0.0	-	74.2	18.5	57.7	9.9	80.6	55.3	60.7	27.3	36.7	35.6	37.5	37.2					
		Snow	54.8	0.7	51.1	8.7	7.4	35.6	-	-	35.6	20.4	31.7	5.0	81.9	35.7	38.5	34.9	23.4	24.0	18.0	29.8					
		All	61.5	0.5	41.2	26.6	4.3	36.4	7.0	51.5	59.7	16.5	36.9	5.5	78.7	52.8	57.8	38.5	48.2	25.0	37.5						
✓	✓	D-Fog	78.6	-	-	18.0	0.0	8.1	20.3	79.3	71.9	1.2	33.7	9.2	68.3	52.1	62.9	40.4	60.3	24.2	59.1	40.5	33.0				36.3
		L-Fog	62.6	0.0	0.0	34.6	8.6	44.9	1.7	0.0	69.6	5.0	36.6	6.7	68.1	57.4	61.6	47.4	52.1	28.4	41.4	36.2					
		Rain	80.0	0.0	-	0.0	28.7	42.2	0.0	-	72.9	19.2	55.1	7.7	80.6	44.3	55.8	29.4	27.0	32.4	39.8	27.9					
		Snow	53.0	0.0	56.3	0.3	3.3	37.9	-	-	34.0	14.9	23.5	4.7	80.3	29.4	34.9	36.7	20.6	23.4	20.7	36.3					
		All	59.5	0.1	43.8	24.2	7.7	37.5	13.4	63.0	59.1	12.4	31.9	6.7	76.4	46.4	55.8	39.9	45.4	25.6	41.2						
✓	✓	D-Fog	82.3	-	-	21.6	0.0	10.0	13.6	75.1	75.1	0.1	33.1	4.9	69.6	56.6	62.3	37.0	60.1	24.9	56.8	40.2	34.7				38.9
		L-Fog	64.9	0.0	0.0	55.4	3.4	54.4	1.4	0.0	74.0	8.8	39.4	4.2	69.5	56.2	60.4	47.3	52.2	31.5	35.7	37.7					
		Rain	86.9	0.0	-	0.8	26.9	52.9	0.0	-	76.0	21.4	56.7	5.3	82.0	49.3	59.1	27.5	33.3	35.9	27.7	33.4					
		Snow	61.2	1.2	65.3	11.8	11.2	45.8	-	-	45.6	23.1	30.6	3.0	82.8	37.7	39.0	38.8	24.6	28.5	17.8	38.9					
		All	65.9	0.8	51.2	38.1	10.6	45.3	5.5	59.4	65.3	19.0	35.7	4.0	78.5	50.4	56.8	40.0	47.7	29.0	35.2						

outperforms PWS or MSS used in isolation. Beyond the absolute performance increase, the integrated configuration also demonstrates more balanced behavior across weather conditions, indicating improved robustness to cross-weather variability. These results suggest that PWS and MSS capture complementary aspects of weather-induced distortions: PWS addresses geometry-related effects while MSS models optical scattering phenomena. Their integration enables the network to learn representations that generalize more reliably across different adverse weather conditions.

To examine whether the proposed framework maintains its effectiveness when the source domain shifts from real-world to synthetic data, we extended the ablation study to the SynLiDAR-to-SemanticSTF setting, with results presented in Table 13. This configuration introduces the additional challenge of bridging both the synthetic-to-real gap and weather-induced domain shift, resulting in lower absolute performance across all methods compared to the SemanticKITTI experiments. Nevertheless, the relative behavior among configurations remains consistent. The combined PhyDAWS framework achieves 23.7% overall mIoU, outperforming both PWS alone (22.4%) and MSS alone (21.8%), with improvements observed across all weather conditions.

The consistency of these patterns across different source domains supports the complementary relationship between the two simulation

strategies. Snow conditions again benefit most from the integrated approach, with mIoU increasing from 23.0% (PWS) and 21.6% (MSS) to 24.4%. At the class level, terrain segmentation shows the most substantial improvement, rising from 29.0% (PWS) and 24.9% (MSS) to 40.7% with the combined approach. These findings confirm that the synergy between PWS and MSS extends to synthetic source domains, supporting the general applicability of the proposed framework for diverse domain generalization scenarios.

4.5.2. Ablation study on physically-inspired components

To analyze how individual physically-inspired components contribute to the robustness of PhyDAWS, we performed a targeted ablation study by selectively disabling key modeling elements, with results summarized in Table 14. The full model achieves 38.9% overall mIoU, and removing any single component leads to clear performance degradation ranging from 2.5% to 4.8%. The magnitude and pattern of degradation vary systematically across weather conditions, reflecting the specific physical phenomena each component is designed to capture. Angle-dependent Occlusion (ADO) produces the largest impact when removed, reducing overall performance to 34.1%, followed by Class-specific Material Reflectivity (CMR) at 34.8% and Height-dependent Density Variation (HDV) at 35.2%.

Table 13

Ablation study on different weather simulation methods with SynLiDAR as source and SemanticSTF as target (%). PWS: Phenomenological Weather Simulation; MSS: Mie-Monte Carlo Scattering Simulation; D-Fog: Dense Fog; L-Fog: Light Fog. Each row represents performance on different target domain weather conditions.

PWS	MSS	Test Weather	Per-Class IoU (%)																	Per-Weather mIoU (%)	All-Weather mIoU (%)					
			Car	Bicycle	Motorcycle	Truck	Other-Vehicle	Person	Bicyclist	Motorcyclist	Road	Parking	Sidewalk	Other-Ground	Building	Fence	Vegetation	Trunk	Terrain			Pole	Traffic-Sign			
✓		D-Fog	20.9	-	-	10.2	0.0	3.1	3.6	26.1	52.3	0.2	11.3	1.8	60.9	27.2	63.1	21.4	30.7	16.9	7.1	21.0	22.7	28.1	23.0	22.4
		L-Fog	34.6	0.9	1.9	34.0	0.6	18.6	4.5	0.0	58.5	2.2	18.5	1.4	60.6	29.6	65.9	36.1	31.7	22.6	9.2					
		Rain	57.3	3.1	-	14.5	0.0	48.6	0.6	-	67.3	7.1	41.0	0.1	77.0	12.9	55.4	23.3	26.6	31.8	11.5					
		Snow	40.7	3.2	12.7	4.6	0.7	44.3	-	-	39.5	7.2	27.3	0.1	77.3	19.8	34.8	28.0	18.3	25.4	7.2					
		All	36.8	1.9	4.2	19.3	0.5	34.1	2.2	5.3	51.4	5.6	21.0	1.1	71.6	25.0	57.3	28.4	29.0	22.8	8.1					
✓		D-Fog	37.7	-	-	13.2	0.0	3.3	17.6	33.9	42.1	0.9	9.2	1.2	51.6	25.0	61.5	22.1	27.4	20.7	9.2	22.2	22.1	27.6	21.6	21.8
		L-Fog	53.8	1.4	1.0	31.1	1.1	18.3	5.6	0.0	49.5	5.5	16.8	2.1	49.9	24.8	61.1	33.9	26.4	27.0	10.8					
		Rain	66.4	0.6	-	23.0	6.2	47.7	1.4	-	56.6	7.4	37.1	0.3	72.8	10.0	50.3	21.6	22.1	33.3	13.2					
		Snow	48.3	1.8	5.5	2.9	0.4	43.1	-	-	25.7	18.0	25.8	1.8	65.2	18.8	30.0	25.0	15.3	30.5	9.3					
		All	49.6	1.3	2.1	17.7	1.2	31.4	6.1	7.4	40.6	12.5	18.8	1.5	60.6	22.1	52.8	26.6	24.9	27.2	10.0					
✓	✓	D-Fog	32.5	-	-	13.8	0.0	3.4	7.5	20.0	39.7	0.0	11.1	1.8	61.6	36.8	62.7	21.7	45.6	18.5	11.5	22.8	23.1	28.2	24.4	23.7
		L-Fog	49.5	3.0	0.0	30.0	0.1	17.4	7.1	0.0	44.6	0.9	18.0	0.4	60.3	36.0	63.1	32.3	44.5	20.5	11.9					
		Rain	63.4	1.1	-	12.7	1.3	45.8	0.0	-	50.1	7.6	35.9	0.2	76.9	24.9	52.8	25.9	34.5	32.9	13.9					
		Snow	49.5	7.3	12.7	4.7	1.1	46.5	-	-	29.1	14.8	27.8	0.3	76.1	24.9	32.8	27.0	22.6	26.6	11.1					
		All	48.1	3.9	2.1	18.9	0.7	35.6	3.5	5.7	38.9	9.9	20.7	0.9	71.1	32.3	55.1	27.1	40.7	23.1	11.7					

Table 14

Ablation study on physically-inspired components with SemanticKITTI as source and SemanticSTF as target (%). ADO: Angle-dependent Occlusion in rain; HDV: Height-dependent Density Variation in snow; DDA: Distance-dependent Attenuation in fog; MCS: Monte Carlo Sampling; CMR: Class-specific Material Reflectivity; D-Fog: Dense Fog; L-Fog: Light Fog. “w/o” denotes “without”.

Model	Test Weather	Per-Class IoU (%)																	Per-Weather mIoU (%)	All-Weather mIoU (%)					
		Car	Bicycle	Motorcycle	Truck	Other-Vehicle	Person	Bicyclist	Motorcyclist	Road	Parking	Sidewalk	Other-Ground	Building	Fence	Vegetation	Trunk	Terrain			Pole	Traffic-Sign			
w/o ADO	D-Fog	75.2	-	-	37.2	0.0	6.2	2.2	18.5	77.7	0.1	41.6	4.1	68.7	53.7	63.7	33.6	61.2	24.4	51.9	36.5	34.9	36.0	31.4	34.1
	L-Fog	82.2	0.0	0.0	48.9	12.1	43.9	0.0	0.0	75.0	6.7	43.8	1.8	66.8	55.1	64.3	41.2	52.0	31.4	37.7					
	Rain	88.3	0.0	-	2.0	0.3	41.1	0.0	-	80.4	13.3	59.2	2.8	80.9	50.6	62.0	17.2	38.7	36.0	38.8					
	Snow	76.5	0.0	12.6	8.8	11.3	38.4	-	-	52.7	14.7	28.8	2.2	84.6	40.1	44.8	31.9	28.8	32.2	25.0					
	All	78.4	0.0	10.4	32.1	8.8	37.4	1.3	14.3	68.8	12.6	37.8	2.6	78.7	50.4	60.4	34.2	50.3	30.2	39.1					
w/o HDV	D-Fog	78.3	-	-	33.6	0.0	6.2	10.6	36.5	71.9	0.1	34.8	4.3	70.8	55.2	59.6	37.7	58.3	25.8	60.5	37.9	32.1	35.2	30.4	35.2
	L-Fog	65.2	0.0	0.0	42.5	4.6	29.7	0.0	0.0	66.7	7.7	38.0	7.2	70.3	58.6	58.9	43.3	51.2	28.4	37.4					
	Rain	85.5	0.0	-	0.3	4.1	48.4	0.0	-	67.3	19.3	58.8	10.6	77.5	45.6	56.6	24.6	27.7	36.1	35.0					
	Snow	63.4	0.9	63.5	2.8	8.9	41.2	-	-	31.1	20.3	28.6	4.1	80.1	36.3	34.1	34.6	24.5	26.0	16.1					
	All	66.9	0.6	52.2	29.3	6.3	37.6	5.4	28.9	56.6	16.4	34.9	5.3	76.8	50.1	53.6	37.1	46.9	27.3	36.2					
w/o DDA	D-Fog	83.2	-	-	18.9	0.0	9.8	16.4	61.0	69.3	0.2	27.8	6.8	69.7	55.0	64.6	36.9	61.4	23.7	57.8	39.0	33.7	36.8	30.5	36.4
	L-Fog	73.7	0.0	0.0	42.1	5.3	48.3	0.0	0.0	69.6	2.7	35.6	6.8	68.1	59.4	67.0	42.7	56.9	28.8	32.9					
	Rain	88.0	0.0	-	0.0	9.0	49.2	0.0	-	77.1	19.9	57.8	11.1	78.3	46.4	60.6	19.8	34.0	36.5	37.3					
	Snow	69.3	0.3	64.7	9.9	4.8	44.5	-	-	35.4	11.9	28.2	4.7	76.6	28.7	38.4	31.4	25.2	26.1	19.2					
	All	73.0	0.2	52.8	28.2	5.3	42.9	5.5	45.6	59.2	9.4	32.3	6.1	74.2	47.1	59.7	35.2	50.6	26.9	37.7					
w/o MCS	D-Fog	77.2	-	-	48.1	0.0	2.8	35.1	61.3	74.9	1.1	34.2	5.6	71.7	57.0	62.0	38.2	61.2	24.6	54.7	41.7	34.1	35.6	29.9	36.0
	L-Fog	69.3	0.0	0.0	43.2	5.5	50.6	0.0	0.0	67.8	8.5	37.9	5.3	71.5	60.4	62.4	45.6	55.0	29.0	35.9					
	Rain	85.4	0.0	-	0.1	10.7	49.4	0.0	-	64.0	19.0	55.3	9.8	81.9	49.8	58.2	27.8	26.4	33.6	33.1					
	Snow	68.3	1.2	13.7	20.0	10.0	45.3	-	-	36.4	20.9	31.4	2.0	81.4	30.3	39.7	35.8	28.6	26.1	17.8					
	All	70.7	0.8	13.0	34.7	8.2	43.2	13.4	48.5	59.0	16.8	35.7	4.1	78.4	49.4	57.6	38.5	50.9	27.0	34.6					
w/o CMR	D-Fog	82.6	-	-	30.4	0.0	10.3	7.2	43.8	74.6	0.2	36.2	2.2	70.6	53.3	63.3	37.5	55.8	22.9	35.6	36.8	33.8	36.7	31.9	34.8
	L-Fog	78.7	0.0	0.6	34.7	4.1	50.4	0.0	0.0	72.5	7.6	45.4	4.0	66.9	57.5	63.6	43.6	50.3	28.4	33.3					
	Rain	87.8	6.0	-	0.0	13.9	35.7	0.0	-	79.3	17.9	61.6	10.1	78.0	47.9	59.8	20.2	35.1	32.8	37.0					
	Snow	76.3	9.5	37.7	12.0	9.3	35.3	-	-	47.9	11.5	32.9	2.9	78.2	35.4	39.5	36.2	26.5	26.6	24.2					
	All	78.2	7.5	22.8	26.0	6.7	35.7	4.5	34.6	65.5	10.4	39.3	3.3	75.1	48.9	58.4	37.7	47.1	26.6	32.1					
Full Model	D-Fog	82.3	-	-	21.6	0.0	10.0	13.6	75.1	75.1	0.1	33.1	4.9	69.6	56.6	62.3	37.0	60.1	24.9	56.8	40.2	34.7	37.7	33.4	38.9
	L-Fog	64.9	0.0	0.0	55.4	3.4	54.4	1.4	0.0	74.0	8.8	39.4	4.2	69.5	56.2	60.4	47.3	52.2	31.5	35.7					
	Rain	86.9	0.0	-	0.8	26.9	52.9	0.0	-	76.0	21.4	56.7	5.3	82.0	49.3	59.1	27.5	33.3	35.9	27.7					
	Snow	61.2	1.2	65.3	11.8	11.2	45.8	-	-	45.6	23.1	30.6	3.0	82.8	37.7	39.0	38.8	24.6	28.5	17.8					
	All	65.9	0.8	51.2	38.1	10.6	45.3	5.5	59.4	65.3	19.0	35.7	4.0	78.5	50.4	56.8	40.0	47.7	29.0	35.2					

The correspondence between components and their target weather conditions is evident in the condition-specific results. Removing HDV causes substantial degradation in snow scenes (30.4% compared to 33.4% for the full model), confirming the importance of vertical particle distribution in shaping snow-affected point cloud structure. Similarly, removing ADO disproportionately affects rain performance (36.0% versus 37.7%), where anisotropic occlusion from falling raindrops dominates the sensing degradation. Removing Distance-dependent Attenuation (DDA) impacts fog and snow conditions, particularly affecting boundary delineation for distant objects. An interesting pattern emerges with Monte Carlo Sampling (MCS), whose removal slightly improves dense fog performance (41.7% versus 40.2%) but substantially degrades snow performance (29.9% versus 33.4%) and causes motorcycle segmentation to collapse from 51.2% to 13.0%. This indicates that stochastic modeling is critical when particle distribution exhibits high spatial variability, particularly for objects with complex geometric profiles. These findings confirm that each component addresses distinct degradation mechanisms, and their collective integration enables robust segmentation across diverse adverse weather conditions.

5. Discussion

5.1. Physical consistency of weather-aware LiDAR augmentation

The experimental results demonstrate that embedding physically grounded weather modeling into the data augmentation process can substantially enhance domain generalization for LiDAR-based semantic segmentation under adverse atmospheric conditions. The proposed PhyDAWS framework explicitly models the interaction between emitted laser pulses and atmospheric particles, thereby preserving the physical coherence between signal attenuation, scattering behavior, and the resulting degradation patterns observed in point cloud data. This aspect is particularly important in Earth observation contexts, where LiDAR serves as an active remote sensing modality whose measurements are directly governed by physical signal propagation through the atmosphere (Zhao et al., 2024; Park et al., 2024). The consistent performance gains observed across fog, rain, and snow scenarios suggest that the improvements primarily stem from exposure to physically plausible perturbations that genuinely reflect real sensing mechanisms.

By incorporating first-principles atmospheric modeling into the learning pipeline, PhyDAWS aligns the augmented training data with the underlying physics of active remote sensing. This alignment ensures that the simulated degradations remain consistent with realistic acquisition conditions, enabling the network to learn representations that are robust to weather-induced signal distortions. From an Earth observation perspective, maintaining such physical fidelity supports reliable deployment across large spatial extents and extended temporal horizons, where atmospheric conditions can vary substantially (Yan et al., 2024; Yang et al., 2025). The results indicate that physically consistent augmentation provides a principled pathway for improving robustness by ensuring that the training distribution encompasses degradation patterns grounded in actual sensing physics.

5.2. Generalization behavior under adverse atmospheric conditions

From a generalization standpoint, physically-inspired augmentation promotes more stable feature learning under weather-induced domain shifts. Performance improvements are observed under moderate atmospheric degradation as well as in more challenging conditions characterized by reduced point density and spatially uneven signal returns. This behavior suggests that the learned representations become less sensitive to weather-specific sampling artifacts and increasingly rely on intrinsic geometric and structural cues present in the scene. As a result, the network exhibits improved robustness when encountering previously unseen weather conditions during inference.

PhyDAWS addresses environmental variability by integrating physical knowledge directly into the training process. This integration enables the network to internalize weather-induced variations as part of the learning objective, reducing reliance on explicit compensation at deployment time. Such a strategy is particularly relevant for Earth observation applications involving heterogeneous datasets and diverse acquisition settings, where consistent performance across varying conditions is essential (Wang et al., 2025; Xiao et al., 2024). The observed generalization behavior across multiple datasets suggests that physically grounded augmentation can effectively enhance transferability without requiring explicit adaptation at test time.

In addition to the overall generalization gains, it is important to consider how the strength of physically-inspired weather augmentation interacts with the contrastive learning objective in shaping weather-invariant representations. In PhyDAWS, weather simulation introduces physically plausible degradations, including attenuation-driven intensity decay, scattering-induced return variability, and geometry-dependent point dropout, while preserving the semantic identity of the underlying scene. As a result, paired views remain structurally related, allowing the contrastive objective to interpret weather-induced signal variations as nuisance factors that do not alter the underlying class membership. When the augmentation strength is appropriately calibrated, the discrepancy between the clean and weather-augmented views provides a sufficiently challenging invariance signal, encouraging the network to suppress weather-specific fluctuations and emphasize geometric and structural cues that remain stable across conditions.

At the same time, the strength of physical augmentation cannot be increased without bound. Excessively severe degradations may reduce the amount of shared informative content between paired views, thereby weakening the contrastive signal and limiting effective feature alignment. The parameter ranges adopted in the proposed framework are therefore selected to balance diversity and correspondence: they are strong enough to expose the network to adverse-weather sensing patterns, yet constrained to preserve sufficient cross-view coherence for meaningful contrastive learning. This coupling between physically grounded perturbation strength and contrastive alignment serves as a stable mechanism for absorbing atmospheric variability into the feature space, supporting the consistent generalization behavior observed across different weather conditions and datasets.

The stability of this generalization behavior under variations in key physical parameters warrants further consideration in the context of operational deployment. Within the proposed framework, extinction coefficients span physically meaningful ranges calibrated to distinct meteorological regimes: fog attenuation coefficients range from $\beta_{\text{light}} = 0.03$ to $\beta_{\text{dense}} = 0.15$, while rain and snow coefficients are set at 0.02 and 0.03 respectively, each accompanied by non-linear exponents that modulate how attenuation compounds with distance. Particle number concentrations similarly vary across intensity levels, with rain conditions characterized by N_0 ranging from 5000 to 15,000 particles/m³ depending on precipitation intensity, and fog employing substantially higher concentrations parameterized through weather-specific gamma distributions calibrated from established meteorological data. Moderate deviations from these calibrated values alter the severity of augmentation, and are expected to produce more gradual performance changes than those recorded under component-level ablation. The ablation results in Table 14 provide empirical grounding for this reasoning: complete removal of individual physically-inspired components produces overall mIoU reductions ranging from 2.5% to 4.8%, implying that partial parameter variation would remain within a substantially narrower margin. The stochastic Monte Carlo scattering factor further contributes to this stability, as weather types and intensity levels are randomly assigned across training samples, exposing the network to a continuous distribution of effective attenuation levels around the calibrated parameter values, providing implicit coverage of the surrounding parameter space. The contrastive learning objective reinforces this behavior by encouraging semantic alignment at the class level, such that the learned representations capture stable structural and geometric cues that persist across a range of attenuation magnitudes, supporting consistent segmentation performance under realistic parameter variation in adverse atmospheric conditions.

5.3. Implications for cross-regional deployment

The domain generalization capability demonstrated by PhyDAWS carries practical implications for deploying autonomous perception systems across geographically diverse regions. The physically-inspired augmentation strategy enables models trained on accessible clear-weather datasets to generalize effectively to previously unseen adverse conditions, reducing the need for extensive region-specific data collection campaigns under various atmospheric conditions. The cross-regional applicability is supported by the physical basis of the augmentation mechanisms, as atmospheric phenomena such as Mie scattering, Beer-Lambert attenuation, and particle size distributions follow consistent physical laws regardless of geographic location.

The experimental results on the SemanticSTF dataset provide supporting evidence for this cross-regional generalization. SemanticSTF comprises data collected across multiple European countries including Germany, Sweden, Denmark, and Finland, representing diverse climatic characteristics and acquisition settings. The consistent performance improvements observed across these geographically distributed test sites indicate that the learned weather-invariant representations transfer effectively across regional variations. The successful synthetic-to-real transfer from SynLiDAR to SemanticSTF demonstrates that PhyDAWS can bridge distribution shifts between simulated and real-world data, a capability valuable for scenarios where real-world adverse weather data is scarce or unavailable.

5.4. Methodological limitations

Despite the demonstrated benefits, several limitations of the current PhyDAWS implementation warrant discussion. The physically-inspired simulations, employed in the framework, inevitably involve simplifying assumptions, and understanding the scope of these simplifications helps contextualize the applicability of the approach.

For example, the use of Mie scattering formulations assumes spherical particles, which may not fully capture the complexity of atmospheric constituents under certain conditions, particularly in snow or mixed-phase environments. This assumption aligns well with fog and rain conditions where liquid droplets predominate, but represents a simplification for snow scenarios involving ice crystals with more varied geometries. The experimental results across all weather conditions suggest that, for the purpose of domain generalization, the augmentation strategy remains effective because the contrastive learning framework encourages the network to focus on weather-invariant structural cues. The simulation aims to expose the network to physically plausible degradation patterns that promote robust feature learning, and the simplifying assumptions primarily affect the severity and spatial distribution of signal degradation while preserving the underlying geometric structure of the observed scenes.

Similarly, the visibility modeling in dense fog employs a threshold-based approach that captures the primary effect of severe signal attenuation while remaining computationally efficient. Real fog environments exhibit continuous attenuation gradients and spatial heterogeneity, which the current model addresses partially through local fog patches as illustrated in Fig. 5. The consistent performance improvements observed in fog conditions indicate that this level of approximation is sufficient for bridging the domain gap in the context of semantic segmentation, though applications requiring finer-grained atmospheric characterization may benefit from more detailed modeling.

The potential for bias introduced by these simplifying assumptions is most consequential in snow scenarios, where the spherical particle approximation for ice crystal scattering introduces systematic errors in extinction coefficient estimates that propagate into the intensity attenuation applied during augmentation. The comparatively lower snow performance of the full model, at 33.4% mIoU relative to 40.2% in dense fog and 37.7% in rain, reflects this reduced physical accuracy under non-spherical particle conditions, though point cloud sparsity inherent to snowy acquisitions also contributes to this performance gap. For the visibility threshold approximation, the residual error grows with fog density beyond the modeled range, as the hard cutoff cannot fully reproduce the spatially continuous signal degradation characteristic of extreme fog environments, a limitation that the local fog patches only partially address. The contrastive learning framework provides partial mitigation across both cases by directing the network toward weather-invariant geometric and structural features, though its effectiveness diminishes in scenarios where extreme atmospheric conditions fundamentally alter the geometric completeness of the point cloud, which constitutes the practical boundary of the current framework's robustness guarantees.

Secondary effects such as sensor-specific noise characteristics and higher-order scattering interactions are not explicitly modeled. These factors may influence the fidelity of the simulated degradations and, consequently, the extent to which the learned representations generalize across different sensing systems. The effectiveness of the current approach reflects a pragmatic balance between computational tractability and physical realism sufficient for the domain generalization objective. While the contrastive learning component proved beneficial, its optimal integration with network architecture remains an area for potential exploration.

5.5. Framework extensibility and future directions

Beyond adverse weather scenarios, the physically-inspired augmentation strategy adopted in PhyDAWS reflects a broader perspective on domain shift in LiDAR-based Earth observation. The proposed framework models perturbations arising from physical sensing processes, introducing structured variations in point density, intensity distribution, and local geometric completeness while preserving the semantic organization of the scene. Although the present study focuses on atmospheric

effects as a representative and practically important source of domain variation, similar measurement characteristics are also observed in other acquisition-related discrepancies commonly encountered in real-world LiDAR data.

In operational Earth observation settings, differences in sensing platforms, calibration conditions, and data quality can introduce systematic variations in LiDAR measurements, manifesting as intensity inconsistencies, range-dependent sparsity, or non-uniform spatial sampling. These factors affect point clouds through mechanisms that share conceptual similarities with the physically grounded degradations considered in this work, even though they are not explicitly parameterized in the current implementation. From this perspective, physically-inspired augmentation provides a general strategy for exposing learning models to realistic sensing-induced variability, encouraging representations that capture stable geometric and structural cues present across diverse acquisition conditions. This alignment suggests potential relevance of the proposed framework to a wider class of domain shifts governed by physical sensing processes, extending beyond the specific weather conditions evaluated in this study.

The modular design of PhyDAWS provides a flexible foundation for future extensions. The physically-inspired augmentation framework can be adapted to incorporate more advanced atmospheric models or additional Earth observation data sources as they become available, such as multi-return LiDAR measurements or complementary active sensing modalities (Huang et al., 2025b,a). This extensibility supports broader transferability across datasets and application contexts, helping to address data gaps that commonly arise in large-scale Earth observation workflows. As such, PhyDAWS represents a practical approach for improving robustness under adverse sensing conditions, while offering a scalable framework that can evolve alongside emerging sensing technologies and expanding Earth observation data infrastructures.

Based on these findings, several directions for future research naturally arise. Further refinement of the physical models, potentially incorporating techniques for non-spherical particle scattering or more detailed atmospheric modeling, could enhance simulation accuracy (La Luna et al., 2025; Raisuddin et al., 2026). Extending the PhyDAWS augmentation principles to other 3D perception tasks within Earth observation, such as object detection and tracking under challenging acquisition conditions, represents a promising direction (Wu et al., 2025; Chae et al., 2025). Investigating the combination of PhyDAWS with other domain generalization strategies (e.g., domain randomization, adversarial alignment) could yield further improvements in robustness. Exploring adaptive simulation techniques, where weather parameters are adjusted dynamically during training, might offer additional benefits. Developing theoretical analyses that relate the fidelity of physical simulation to downstream generalization performance could provide principled guidance for the design of future augmentation strategies.

6. Conclusion

The reliable perception of complex 3D environments under diverse meteorological conditions presents a significant and persistent challenge for autonomous driving systems. This issue constitutes a key barrier to their widespread and safe deployment. Conventional deep learning models for LiDAR point cloud semantic segmentation, often trained on extensive datasets gathered in clear weather, typically experience marked performance degradation when operating in rain, snow, or fog. These atmospheric phenomena fundamentally alter LiDAR measurements via complex physical mechanisms, including scattering, attenuation, and occlusion. This creates a significant domain gap between training and deployment conditions. This paper addressed this challenge by focusing on domain generalization. The goal is to enable models trained exclusively on source domain data (such as clear weather or synthetic datasets) to perform effectively in unseen target

domains characterized by adverse weather, without requiring access to target domain labels or necessitating run-time adaptation.

We introduced PhyDAWS, a novel framework that incorporates physically-inspired modeling of LiDAR-weather interactions. PhyDAWS utilizes two complementary simulation approaches: a phenomenological technique capturing observed phenomena and a detailed Mie scattering-based method. These approaches generate physically grounded augmentations that move beyond heuristic or purely geometric strategies. Through extensive experiments on established SemanticKITTI-to-SemanticSTF and SynLiDAR-to-SemanticSTF domain generalization benchmarks, PhyDAWS demonstrated state-of-the-art performance. It achieved consistently higher mIoU scores across varied adverse weather conditions, improving upon prior state-of-the-art methods by significant margins (ranging from 4.1% to 10.3% for real-to-real and 0.3% to 5.2% for synthetic-to-real transfer). Furthermore, systematic ablation studies provided evidence for the individual contributions and synergistic effects of the phenomenological and Mie-scattering simulations, validating the effectiveness of the physically-inspired components.

In conclusion, this paper provides strong evidence for the efficacy and potential of physically-inspired data augmentation as a principled and effective strategy for achieving domain generalization in LiDAR point cloud semantic segmentation under diverse weather conditions. By systematically modeling the physical processes affecting sensor measurements, the PhyDAWS framework improves the reliability of 3D perception models when faced with unseen adverse weather. This contributes to enabling autonomous systems capable of navigating and operating more safely and reliably across the wide range of environmental conditions encountered in real-world operation.

CRediT authorship contribution statement

Jing Du: Writing – review & editing, Writing – original draft, Visualization, Validation, Methodology, Investigation, Funding acquisition, Formal analysis. **John Zelek:** Writing – review & editing, Validation, Supervision, Project administration, Methodology, Funding acquisition. **Chaolei Wang:** Writing – review & editing, Validation, Investigation. **Ting Han:** Writing – review & editing, Validation, Investigation. **Yip-ing Chen:** Writing – review & editing, Investigation. **Dedong Zhang:** Writing – review & editing. **Jonathan Li:** Writing – review & editing, Validation, Supervision, Project administration, Methodology, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was supported in part by the China Scholarship Council under PhD Scholarship 202208350003.

Data availability

Data will be made available on request.

References

- Awasthi, S., Varade, D., 2021. Recent advances in the remote sensing of alpine snow: A review. *GISci. Remote. Sens.* 58 (6), 852–888.
- Behley, J., Garbade, M., Milioto, A., Quenzel, J., Behnke, S., Stachniss, C., Gall, J., 2019. SemanticKITTI: A dataset for semantic scene understanding of LiDAR sequences. In: *Proc. ICCV. IEEE*, pp. 9296–9306. <http://dx.doi.org/10.1109/ICCV.2019.00939>.
- Bi, J., Song, Y., Jiang, Y., Sun, L., Wang, X., Liu, Z., Xu, J., Quan, S., Dai, Z., Yan, W., 2025. Lane detection for autonomous driving: Comprehensive reviews, current challenges, and future predictions. *IEEE Trans. Intell. Transp. Syst.* 26 (5), 5710–5746. <http://dx.doi.org/10.1109/TITS.2024.3524603>.
- Bijelic, M., Gruber, T., Mannan, F., Kraus, F., Ritter, W., Dietmayer, K., Heide, F., 2020. Seeing through fog without seeing fog: Deep multimodal sensor fusion in unseen adverse weather. In: *Proc. CVPR. IEEE*, pp. 11679–11689. <http://dx.doi.org/10.1109/CVPR42600.2020.01170>.
- Bijelic, M., Gruber, T., Ritter, W., 2018. Benchmarking image sensors under adverse weather conditions for autonomous driving. In: *Proc. IEEE Intell. Vehicles Symp., IV. IEEE*, pp. 1773–1779.
- Bohren, C.F., Huffman, D.R., 2008. *Absorption and Scattering of Light by Small Particles*. John Wiley & Sons.
- Chae, Y., Park, H., Kim, H., Yoon, K.J., 2025. Doppler-aware LiDAR-RADAR fusion for weather-robust 3D detection. In: *Proc. ICCV*. pp. 27197–27208.
- Chen, Y., Hu, V.T., Gavves, E., Mensink, T., Mettes, P., Yang, P., Snoek, C.G., 2020. Pointmixup: Augmentation for point clouds. In: *Proc. ECCV*. pp. 330–345.
- Chen, T., Ying, X., 2025. FPSMix: data augmentation strategy for point cloud classification. *Front. Comput. Sci.* 19 (2), 192701. <http://dx.doi.org/10.1007/S11704-023-3455-4>, URL: <https://doi.org/10.1007/s11704-023-3455-4>.
- Choy, C.B., Gwak, J., Savarese, S., 2019. 4D spatio-temporal ConvNets: Minkowski convolutional neural networks. In: *Proc. CVPR. IEEE*, pp. 3075–3084. <http://dx.doi.org/10.1109/CVPR.2019.00319>.
- Deng, J., Lu, J., Zhang, T., 2025. Quantity-quality enhanced self-training network for weakly supervised point cloud semantic segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* 47 (5), 3580–3596. <http://dx.doi.org/10.1109/TPAMI.2025.3532637>.
- Du, J., Cai, G., Wang, Z., Huang, S., Su, J., Marcato Junior, J., Smit, J., Li, J., 2021. ResDLPS-Net: Joint residual-dense optimization for large-scale point cloud semantic segmentation. *ISPRS J. Photogramm. Remote Sens.* 182, 37–51. <http://dx.doi.org/10.1016/j.isprsjprs.2021.09.024>.
- Du, J., Cai, G., Wang, Z., Su, J., Huang, M., Zelek, J.S., Jr., J.M., Li, J., 2025. MTCloud: Multi-type convolutional linkage network for point cloud instance segmentation. *Expert Syst. Appl.* 270, 126432. <http://dx.doi.org/10.1016/J.ESWA.2025.126432>.
- Du, J., Zelek, J., Li, J., 2024. Weather-aware autopilot: Domain generalization for point cloud semantic segmentation in diverse weather scenarios. *ISPRS J. Photogramm. Remote Sens.* 218, 204–219. <http://dx.doi.org/10.1016/j.isprsjprs.2024.09.006>.
- Grabner, M., Kvicera, V., 2013. Multiple scattering in rain and fog on free-space optical links. *J. Light. Technol.* 32 (3), 513–520.
- Grandhi, A., 2025. *Operational Performance of an Automated Shuttle: Comparative Insights and the Effect of Disengagements* (Master's thesis). The University of North Carolina at Charlotte, Charlotte, NC, USA.
- Hasirlioglu, S., Kamann, A., Doric, I., Brandmeier, T., 2016. Test methodology for rain influence on automotive surround sensors. In: *Proc. IEEE Int. Conf. Intelligent Transportation Systems. ITSC. IEEE*, pp. 2242–2247.
- Heo, J., Im, B., Shin, S., Ha, S.e., Kang, D., Song, T., Gu, Y., Han, S., 2025. Experimental analysis of sensor's performance degradation under adverse weather conditions. *Int. J. Precis. Eng. Manuf.* 1655–1672. <http://dx.doi.org/10.1007/s12541-025-01228-3>.
- Hu, S., Fang, Z., Deng, Y., Chen, X., Fang, Y., Kwong, S., 2025. Toward full-scene domain generalization in multi-agent collaborative bird's eye view segmentation for connected and autonomous driving. *IEEE Trans. Intell. Transp. Syst.* 26 (2), 1783–1796. <http://dx.doi.org/10.1109/TITS.2024.3506284>.
- Huang, X., Wang, J., Xia, Q., Chen, S., Yang, B., Li, X., Wang, C., Wen, C., 2025a. V2X-R: Cooperative LiDAR-4D radar fusion with denoising diffusion for 3D object detection. In: *Proc. CVPR*. pp. 27390–27400.
- Huang, X., Xu, Z., Wu, H., Wang, J., Xia, Q., Xia, Y., Li, J., Gao, K., Wen, C., Wang, C., 2025b. L4dr: Lidar-4dradar fusion for weather-robust 3d object detection. In: *Proc. AAAI*. 39, (4), pp. 3806–3814.
- Hulst, H.C., van de Hulst, H.C., 1981. *Light Scattering by Small Particles*. Courier Corporation.
- Ishimaru, A., et al., 1978. *Wave Propagation and Scattering in Random Media*. vol. 2, Academic Press New York.
- Iwana, B.K., Uchida, S., 2021. An empirical survey of data augmentation for time series classification with neural networks. *PLoS One* 16 (7), e0254841.
- Kaasalainen, S., Ahokas, E., Hyyppa, J., Suomalainen, J., 2005. Study of surface brightness from backscattered laser intensity: Calibration of laser data. *IEEE Geosci. Remote. Sens. Lett.* 2 (3), 255–259.
- Kettelgerdes, M., Sarmiento, N., Erdogan, H., Wunderle, B., Elger, G., 2024. Precise adverse weather characterization by deep-learning-based noise processing in automotive LiDAR sensors. *Remote. Sens.* 16 (13), 2407.
- Khoshbakhtnejad, E., Golezani, F.B., Mohammadian, B., Abou Yassine, A.H., Sojoudi, H., 2024. Trajectory and impact dynamics of snowflakes: Fundamentals and applications. *Powder Technol.* 120298.

- Kim, S., Lee, S., Hwang, D., Lee, J., Hwang, S.J., Kim, H.J., 2021. Point cloud augmentation with weighted local transformations. In: Proc. ICCV. pp. 548–557.
- Kim, J., Woo, J., Kim, J., Im, S., 2024. Rethinking LiDAR domain generalization: Single source as multiple density domains. In: Proc. ECCV. 15078, Springer, pp. 310–327. http://dx.doi.org/10.1007/978-3-031-72661-3_18.
- Kong, L., Xu, X., Ren, J., Zhang, W., Pan, L., Chen, K., Ooi, W.T., Liu, Z., 2025. Multi-modal data-efficient 3D scene understanding for autonomous driving. IEEE Trans. Pattern Anal. Mach. Intell. 47 (5), 3748–3765. <http://dx.doi.org/10.1109/TPAMI.2025.3535625>.
- Kuang, W., Zhao, X., Shen, Y., Wen, C., Lu, H., Zhou, Z., Chen, X., 2025. ResLPR: A LiDAR data restoration network and benchmark for robust place recognition against weather corruptions. arXiv preprint [arXiv:2503.12350](https://arxiv.org/abs/2503.12350).
- La Luna, A., Zhang, Z., Zheng, J., Song, Q., Yu, H., Ding, J., Yang, P., Saito, M., 2025. Scattering properties and lidar characteristics of Asian dust particles based on realistic shape models. Atmos. Chem. Phys. 25 (20), 13359–13377. <http://dx.doi.org/10.5194/acp-25-13359-2025>.
- Li, Y., Bai, F., Lyu, C., Qu, X., Liu, Y., 2025. A systematic review of generative adversarial networks for traffic state prediction: overview, taxonomy, and future prospects. Inf. Fusion 102915.
- Li, P., Liu, W., Li, D., Zhang, M., Wang, X., Sun, H., Ma, J., 2024. Eavesdropping risk evaluation for non-line-of-sight terahertz channels by a metallic wavy surface in rain. J. Opt. Soc. Am. B 41 (9), 1995–2002.
- Li, H., Pan, S.J., Wang, S., Kot, A.C., 2018. Domain generalization with adversarial feature learning. In: Proc. CVPR. IEEE, pp. 5400–5409. <http://dx.doi.org/10.1109/CVPR.2018.00566>.
- Li, J., Wang, J., Chen, J., Xu, T., 2025. Towards robust point cloud recognition with sample-adaptive auto-augmentation. IEEE Trans. Pattern Anal. Mach. Intell. 47 (4), 3003–3017. <http://dx.doi.org/10.1109/TPAMI.2025.3528392>.
- Li, B., Xu, Z., Li, J., Liu, X., Fang, J., Li, X., Yu, H., 2025. V2X-DG: Domain generalization for vehicle-to-everything cooperative perception. arXiv preprint [arXiv:2503.15435](https://arxiv.org/abs/2503.15435).
- Lin, F., Lin, T., Yao, Y., Ren, H., Wu, J., Cai, Q., 2025. VPA-Net: A visual perception assistance network for 3d lidar semantic segmentation. Pattern Recognit. 158, 111014. <http://dx.doi.org/10.1016/j.patcog.2024.111014>.
- Lin, H., Liu, Y., Wang, L., Qu, X., 2025. Big data-driven advancements and future directions in vehicle perception technologies: From autonomous driving to modular buses. IEEE Trans. Big Data 11 (3), 1568–1587. <http://dx.doi.org/10.1109/TBDATA.2025.3527208>.
- Liu, B., Zheng, C., Xu, X., Xu, C., Zhang, H., He, S., 2025. Rotation-adaptive point cloud domain generalization via intricate orientation learning. IEEE Trans. Pattern Anal. Mach. Intell. 47 (5), 4232–4239. <http://dx.doi.org/10.1109/TPAMI.2025.3535230>.
- Middleton, W.E.K., 1952. Vision Through the Atmosphere. University of Toronto Press.
- Nawaz, M., Khan, S., Daud, M., Asim, M., Anwar, G.A., Shahid, A.R., HO, H.P.A., Chan, T., Pak Kong, D., Yuan, W., 2025. Improving autonomous vehicle cognitive robustness in extreme weather with deep learning and thermal camera fusion. IEEE Open J. Veh. Technol. 6, 426–441.
- Park, J., Kim, K., Shim, H., 2024. Rethinking data augmentation for robust LiDAR semantic segmentation in adverse weather. In: Proc. ECCV. URL: <https://api.semanticscholar.org/CorpusID:270878732>.
- Peng, B., Wu, X., Jiang, L., Chen, Y., Zhao, H., Tian, Z., Jia, J., 2024. OA-CNNs: Omni-adaptive sparse CNNs for 3D semantic segmentation. In: Proc. CVPR. IEEE, pp. 21305–21315. <http://dx.doi.org/10.1109/CVPR52733.2024.02013>.
- Pruppacher, H.R., Klett, J.D., Wang, P.K., 1998. Microphysics of Clouds and Precipitation. Springer, Dordrecht.
- Qi, C.R., Su, H., Mo, K., Guibas, L.J., 2017a. PointNet: Deep learning on point sets for 3D classification and segmentation. In: Proc. CVPR. IEEE, pp. 77–85. <http://dx.doi.org/10.1109/CVPR.2017.16>.
- Qi, C.R., Yi, L., Su, H., Guibas, L.J., 2017b. PointNet++: Deep hierarchical feature learning on point sets in a metric space. In: Proc. NeurIPS. pp. 5099–5108.
- Rafi, T.H., Mahjabin, R., Ghosh, E., Ko, Y.W., Lee, J., 2024. Domain generalization for semantic segmentation: a survey. Artif. Intell. Rev. 57 (9), 247. <http://dx.doi.org/10.1007/S10462-024-10817-Z>.
- Raisuddin, A.M., Holmblad, J., Haghghi, H., Poledna, Y., Drechsler, M.F., Donzella, V., Aksoy, E.E., 2026. REHEARSE-3D: A multi-modal emulated rain dataset for 3D point cloud de-raining. Sensors 26 (2), <http://dx.doi.org/10.3390/s26020728>.
- Sanchez, J., Deschaud, J.-E., Goulette, F., 2025. COLA: Coarse-label multisource LiDAR semantic segmentation for autonomous driving. IEEE Trans. Robot. 41, 1742–1754. <http://dx.doi.org/10.1109/TRO.2025.3543302>.
- Sapkota, R., Raza, S., Shoman, M., Paudel, A., Karkee, M., 2025. Image, text, and speech data augmentation using multimodal LLMs for deep learning: A survey. arXiv preprint [arXiv:2501.18648](https://arxiv.org/abs/2501.18648).
- Seinfeld, J.H., Pandis, S.N., 2016. Atmospheric Chemistry and Physics: From Air Pollution to Climate Change. John Wiley & Sons.
- Srivastava, N., Hinton, G.E., Krizhevsky, A., Sutskever, I., Salakhutdinov, R., 2014. Dropout: a simple way to prevent neural networks from overfitting. J. Mach. Learn. Res. 15 (1), 1929–1958. <http://dx.doi.org/10.5555/2627435.2670313>.
- Su, H., Gao, H., Wang, X., Fang, X., Liu, Q., Huang, G., Li, X., Cao, Q., 2024. Object detection in adverse weather for autonomous vehicles based on sensor fusion and incremental learning. IEEE Trans. Instrum. Meas. 73, 1–10. <http://dx.doi.org/10.1109/TIM.2024.3472860>.
- Sun, C., Sun, P., Wang, J., Guo, Y., Zhao, X., 2024. Understanding LiDAR performance for autonomous vehicles under snowfall conditions. IEEE Trans. Intell. Transp. Syst. 25 (11), 16462–16472. <http://dx.doi.org/10.1109/TITS.2024.3409907>.
- Sun, T., Zhang, Z., Tan, X., Qu, Y., Xie, Y., 2024. Image understands point cloud: Weakly supervised 3D semantic segmentation via association learning. IEEE Trans. Image Process. 33, 1838–1852. <http://dx.doi.org/10.1109/TIP.2024.3372449>.
- Sun, X., Zhang, L., Xu, J., Yuan, J., Chen, T., 2025. Learnable bi-directional data augmentation for few-shot cross-domain point cloud classification. Neurocomputing 625, 129486. <http://dx.doi.org/10.1016/J.NEUCOM.2025.129486>, URL: <https://doi.org/10.1016/j.neucom.2025.129486>.
- Thurai, M., Bringi, V., 2005. Drop axis ratios from a 2D video disdrometer. J. Atmos. Ocean. Technol. 22 (7), 966–978.
- Tsakmakopoulou, D., Moustakas, K., 2024. Perception for connected autonomous vehicles under adverse weather conditions. In: Proc. IROS. IEEE, pp. 3161–3166. <http://dx.doi.org/10.1109/IROS58592.2024.10801295>.
- Ulbrich, C.W., 1983. Natural variations in the analytical form of the raindrop size distribution. J. Clim. Appl. Meteorol. 1764–1775.
- Wang, J., Lin, C., Nie, L., Huang, S., Zhao, Y., Pan, X., Ai, R., 2024a. Weatherdepth: Curriculum contrastive learning for self-supervised depth estimation under adverse weather conditions. In: Proc. ICRA. IEEE, pp. 4976–4982.
- Wang, Y., Wang, J., Li, J., Zhao, Z., Chen, G., Liu, A., Heng, P.A., 2024b. Pointpatchmix: Point cloud mixing with patch scoring. In: Proc. AAAI. pp. 5686–5694.
- Wang, P., Yao, W., Shao, J., He, Z., 2025. Test-time adaptation for geospatial point cloud semantic segmentation with distinct domain shifts. ISPRS J. Photogramm. Remote Sens. 229, 422–435. <http://dx.doi.org/10.1016/j.isprsjprs.2025.08.022>.
- Wang, Y., Zhao, W., Cao, C., Deng, T., Wang, J., Chen, W., 2024c. SFPNet: Sparse focal point network for semantic segmentation on general LiDAR point clouds. In: Proc. ECCV. 15063, Springer, pp. 403–421. http://dx.doi.org/10.1007/978-3-031-72652-1_24.
- Wu, X., Jiang, L., Wang, P., Liu, Z., Liu, X., Qiao, Y., Ouyang, W., He, T., Zhao, H., 2024. Point transformer V3: simpler, faster, stronger. In: Proc. CVPR. IEEE, pp. 4840–4851. <http://dx.doi.org/10.1109/CVPR52733.2024.00463>.
- Wu, X., Lao, Y., Jiang, L., Liu, X., Zhao, H., 2022. Point transformer V2: grouped vector attention and partition-based pooling. In: Proc. NeurIPS.
- Wu, Y., Zhu, Y., Zhang, K., Qian, J., Xie, J., Yang, J., 2025. WeatherGen: A unified diverse weather generator for LiDAR point clouds via spider mamba diffusion. In: Proc. CVPR. pp. 17019–17028.
- Xiao, A., Huang, J., Guan, D., Cui, K., Lu, S., Shao, L., 2022a. PolarMix: A general data augmentation technique for LiDAR point clouds. In: NeurIPS. URL: <https://arxiv.org/abs/2208.00223>.
- Xiao, A., Huang, J., Guan, D., Zhan, F., Lu, S., 2022b. Transfer learning from synthetic to real LiDAR point cloud for semantic segmentation. In: Proc. AAAI. AAAI Press, pp. 2795–2803, URL: <https://ojs.aaai.org/index.php/AAAI/article/view/20183>.
- Xiao, A., Huang, J., Xuan, W., Ren, R., Liu, K., Guan, D., El-Saddik, A., Lu, S., Xing, E.P., 2023. 3D semantic segmentation in the wild: Learning generalized models for adverse-condition point clouds. In: Proc. CVPR. <http://dx.doi.org/10.48550/arXiv.2304.00690>.
- Xiao, A., Zhang, X., Shao, L., Lu, S., 2024. A survey of label-efficient deep learning for 3D point clouds. IEEE Trans. Pattern Anal. Mach. Intell. 46 (12), 9139–9160.
- Xie, S., Kong, L., Zhang, W., Ren, J., Pan, L., Chen, K., Liu, Z., 2025. Benchmarking and improving bird's eye view perception robustness in autonomous driving. IEEE Trans. Pattern Anal. Mach. Intell. 47 (5), 3878–3894. <http://dx.doi.org/10.1109/TPAMI.2025.3535960>.
- Xu, J., Ma, X., Zhang, L., Zhang, B., Chen, T., 2024. Push-and-pull: A general training framework with differential augmentor for domain generalized point cloud classification. IEEE Trans. Circuits Syst. Video Technol. 34 (8), 7165–7175.
- Xu, C., Sankar, R., 2024. A comprehensive review of autonomous driving algorithms: Tackling adverse weather conditions, unpredictable traffic violations, blind spot monitoring, and emergency maneuvers. Algorithms 17 (11), 526.
- Yan, X., Zheng, C., Xue, Y., Li, Z., Cui, S., Dai, D., 2024. Benchmarking the robustness of LiDAR semantic segmentation models. Int. J. Comput. Vis. 132, 2674–2697. <http://dx.doi.org/10.1007/s11263-024-01991-2>.
- Yang, Z., Chen, L., Sun, Y., Li, H., 2024. Visual point cloud forecasting enables scalable autonomous driving. In: Proc. CVPR. IEEE, pp. 14673–14684. <http://dx.doi.org/10.1109/CVPR52733.2024.01390>.
- Yang, L., Hu, P., Yuan, S., Zhang, L., Liu, J., Shen, H., Zhu, X., 2025. Towards explicit geometry-reflectance collaboration for generalized LiDAR segmentation in adverse weather. In: Proc. CVPR. pp. 139–149.
- Yao, X., Bai, Y., Zhang, X., Zhang, Y., Sun, Q., Chen, R., Li, R., Yu, B., 2022. PCL: proxy-based contrastive learning for domain generalization. In: Proc. CVPR. IEEE, pp. 7087–7097. <http://dx.doi.org/10.1109/CVPR52688.2022.00696>.
- Yu, L., Zhong, H., Zou, L., Chen, K., Gao, P., 2025. Bridging domain gap of point cloud representations via self-supervised geometric augmentation. IEEE Trans. Circuits Syst. Video Technol. 35 (5), 4846–4856. <http://dx.doi.org/10.1109/TCSVT.2024.3525052>.
- Zhan, J., Duan, Y., Ding, J., Hu, X., Huang, X., Ma, J., 2024. Towards visibility estimation and noise-distribution-based defogging for LiDAR in autonomous driving. In: Proc. ICRA. pp. 18443–18449.

- Zhang, Y., Carballo, A., Yang, H., Takeda, K., 2023. Perception and sensing for autonomous vehicles under adverse weather conditions: A survey. *ISPRS J. Photogramm. Remote Sens.* 196, 146–177. <http://dx.doi.org/10.1016/j.isprsjprs.2022.12.021>.
- Zhao, H., Jiang, L., Jia, J., Torr, P.H.S., Koltun, V., 2021. Point transformer. In: *Proc. ICCV. IEEE*, pp. 16239–16248. <http://dx.doi.org/10.1109/ICCV48922.2021.01595>.
- Zhao, G., Ma, F., Qi, W., Liu, Y., Liu, M., Ma, J., 2025. CurbNet: Curb detection framework based on LiDAR point cloud segmentation. *IEEE Trans. Intell. Transp. Syst.* 26 (6), 8961–8974. <http://dx.doi.org/10.1109/TITS.2025.3540645>.
- Zhao, H., Zhang, J., Chen, Z., Zhao, S., Tao, D., 2024. UniMix: Towards domain adaptive and generalizable LiDAR semantic segmentation in adverse weather. In: *Proc. CVPR. IEEE*, pp. 14781–14791. <http://dx.doi.org/10.1109/CVPR52733.2024.01400>.
- Zhu, Q., Fan, L., Weng, N., 2024. Advancements in point cloud data augmentation for deep learning: A survey. *Pattern Recognit.* 153, 110532. <http://dx.doi.org/10.1016/J.PATCOG.2024.110532>.
- Zhu, X., Zhou, H., Wang, T., Hong, F., Ma, Y., Li, W., Li, H., Lin, D., 2021. Cylindrical and asymmetrical 3D convolution networks for LiDAR segmentation. In: *Proc. CVPR*. pp. 9939–9948.
- Zhuang, Y., Li, Q., Chen, Y., Huai, J., Li, M., Ma, T., Tang, Y., Liang, X., 2024. 3D-SeqMOS: A novel sequential 3D moving object segmentation in autonomous driving. *IEEE Trans. Intell. Transp. Syst.* 25 (8), 8782–8795. <http://dx.doi.org/10.1109/TITS.2024.3391286>.