Weather-Robust LiDAR Perception: Point Cloud Restoration from Adverse Weather

Chenghao Sun[1], **Pengpeng Sun**[1*], **Xiangmo Zhao**[1,2]

[1]School of Information Engineering, Chang'an University,China
[2]College of Urban Development and Modern Transportation Xi'an University of Architecture and Technology, China

Abstract

Adverse weather conditions—such as rain, fog, and snow—significantly degrade LiDAR point cloud quality, causing substantial performance deterioration in detection models trained on clean data. To address this, we propose LTDNet, a novel point cloud quality improvement network that restores degraded LiDAR scans by learning an end-to-end mapping from corrupted to clean geometry. LTDNet leverages position encoding, spatial–frequency joint feature extraction, weather-aware refinement, and probabilistic pruning to effectively recover structural integrity while suppressing weather-induced noise. To facilitate standardized evaluation, we introduce IQA3D, a new benchmark comprising both synthetic and real-world sequences under adverse weather. This dual-design benchmark serves two complementary purposes: synthetic sequences provide pixel-wise correspondences between degraded and clean point clouds for quantitatively assessing restoration fidelity, while real-world sequences enable evaluation of the practical impact of improvement methods on downstream 3D object detection under authentic weather conditions. This makes IQA3D particularly suitable for jointly measuring both perceptual quality and task-level robustness of point cloud improvement models. Extensive experiments on IQA3D demonstrate that LTDNet significantly improves detection performance across various state-of-the-art 3D detectors and three tested weather conditions, making it a practical and effective solution for robust LiDAR-based detection.

Introduction

LiDAR-based perception faces substantial challenges in adverse weather. Rain, fog, and snow cause signal attenuation, scattering, and unintended reflections, resulting in corrupted or missing point cloud data(Dreissig et al. 2023; Teufel et al. 2023; Piroli et al. 2023). This degradation reduces detection accuracy, effective range, and overall reliability of LiDAR-based perception, posing serious safety risks for autonomous vehicles(Vargas et al. 2021; Brzozowski and Parczewski 2023). Consequently, enhancing LiDAR robustness in diverse adverse weather remains a critical open challenge in autonomous perception(Sun et al. 2024).

*Corresponding author: pengpeng.sun@chd.edu.cn.

Existing solutions primarily fall into two categories. The first involves data augmentation by simulating adverse weather to improve model generalization(Kilic et al. 2021; Hahner et al. 2022; Bernuth, Volk, and Bringmann 2019), which often suffers from domain shift. The second focuses on enhancing the quality of degraded point clouds to provide cleaner inputs for downstream perception tasks, with most current efforts centered on denoising methods(Kurup and Bos 2021; Zhao et al. 2024; Seppänen, Ojala, and Tammi 2022; Heinzler et al. 2020; Yan et al. 2025; Sun et al. 2022). However, these methods are limited in addressing structural degradation caused by signal attenuation.

Moreover, there is currently no comprehensive benchmark to systematically evaluate how point cloud improvement methods affect downstream perception tasks, particularly in real-world degraded scenarios. The absence of standardized metrics and test datasets hinders objective comparisons and progress in this field.

To address these challenges, we propose LTDNet, a novel LiDAR degradation-to-clean network designed to enhance point cloud quality under adverse weather conditions. LTDNet serves as a plug-and-play preprocessing module that can be seamlessly integrated into existing LiDAR perception pipelines without modifying downstream model. Given a degraded point cloud represented as a 2D depth map, the network learns a mapping to a clean counterpart through a series of carefully designed components: spatial position encoding, a spatial-frequency joint feature extraction block, and a probabilistic pruning block. The resulting pseudo-clean point cloud preserves structural integrity and geometric fidelity, providing a more robust foundation for perception in challenging weather scenarios. In addition, we introduce a new evaluation benchmark to assess both the fidelity of point cloud restoration and its impact on downstream object detection performance, particularly under real-world adverse weather conditions. To address limitations of existing metrics, we introduce Degradation Degree Improvement (DDI), a new indicator that quantifies gains in geometric fidelity. In summary, our contributions can be summarized as follows:

• We propose LTDNet, a novel framework for improving LiDAR point cloud quality under adverse weather. By restoring structural integrity and geometric fidelity, LTDNet improves the robustness of downstream perception tasks like

object detection.

• We introduce IQA3D, a benchmark designed to (i) quantitatively assess point cloud improvement methods and (ii) evaluate their impact on object detection under real-world degradations.

• Extensive experiments on IQA3D show that LTDNet outperforms three existing advanced preprocessing methods, delivering superior quality restoration and consistent detection accuracy gains across various pretrained models.

Related Work

Datasets for Adverse Weather Perception

Recent research in autonomous driving has spurred the development of datasets focusing on challenging weather. The Canadian Adverse Driving Conditions(Pitropov et al. 2021) and Winter Adverse Driving datasets(Kurup and Bos 2021) provide LiDAR data under snow, with 3D or point-wise annotations, but they lack multi-weather diversity. The Boreas(Burnett et al. 2023) and Oxford RobotCar datasets(Maddern et al. 2017) cover a range of weather conditions, yet neither provides object annotations under adverse weather. In contrast, the DENSE dataset(Bijelic et al. 2020) offers over 10,000 km of driving data under diverse weather conditions with high-quality 3D annotations, and the RADIATE(Sheeny et al. 2021) dataset includes LiDAR data across multiple weather types, annotated with over 200,000 labeled instances. In this work, we select RADIATE and DENSE as base datasets for training and evaluating robustness under realistic weather degradations.

LiDAR-based 3D Object Detection

3D object detectors based on LiDAR can be categorized into point-based, voxel-based, and hybrid methods(Goswami et al. 2025). PointNet(Qi et al. 2017a) and its variants(Qi et al. 2017b) directly process raw point clouds by leveraging symmetric functions, preserving fine-grained spatial features but with limited scalability. VoxelNet(Zhou and Tuzel 2018) introduced voxelization to structure the data and applied mini-PointNet inside each voxel. SECOND(Yan, Mao, and Li 2018) further improved computational efficiency through sparse 3D convolutions. PointPillar(Lang et al. 2019) simplified voxel representation to vertical pillars, achieving real-time inference speeds. Voxel Transformer(Mao et al. 2021) introduces attention-based voxel feature aggregation to better capture long-range spatial relations in sparse point clouds. Hybrid models such as PV-RCNN(Shi et al. 2020) combine voxel context with precise point features to further enhance detection accuracy. However, most of these methods are trained and evaluated on clean datasets and exhibit performance degradation when exposed to noisy or incomplete point clouds caused by adverse weather. This shortcoming has motivated the development of more robust detection algorithms capable of addressing these environmental challenges.

Point Cloud Quality Improvement under Corruptions

To address weather-induced degradation in LiDAR point clouds, researchers have proposed various quality improvement methods, primarily focused on noise removal. Traditional rule-based filters(Charron, Phillips, and Waslander 2018; Kurup and Bos 2021; Huang et al. 2023; Roriz et al. 2021) leverage heuristics based on spatial density, intensity, and distance to remove outliers. These methods are computationally efficient and interpretable but may either over-filter useful points or fail to eliminate complex noise patterns. Learning-based methods aim to overcome these limitations. WeatherNet(Heinzler et al. 2020), a modified LiLaNet(Piewak et al. 2018), was the first CNN designed for point cloud denoising. SunnyNet(Chappell and Pereira 2022) enhances this with attention mechanisms. 4DenoiseNet(Seppänen, Ojala, and Tammi 2022) incorporates temporal information to improve filtering accuracy. 3DOutDet(Raisuddin et al. 2024) and its extension 3DUnOutDet(Raisuddin, Gouigah, and Aksoy 2024) utilize localized neighborhood convolutions and learn priors on sparsity and intensity to guide denoising.

While these approaches focus on noise removal, they typically fail to recover structural information lost due to severe degradation. To bridge this gap, we propose LTDNet, a learning-based framework that reconstructs high-quality point clouds from weather-corrupted.

Method

As illustrated in Figure 1, LTDNet adopts an encoder–decoder architecture. Given a degraded point cloud frame, we first project it into a 2D range image using an equirectangular projection with a single range channel. The resulting range image is then fed into two key modules: a Global Positional Encoding and a Local Positional Encoding, which explicitly compensate for the spatial distortion caused by the projection. The encoder comprises three Spatial-Frequency Blocks (SFBs), which progressively downsample the fused features. Each SFB captures local spatial and global frequency cues in parallel and integrates them using dual attention to ensure robustness to geometric degradation and noise. The decoder mirrors the encoder with three SFB-based blocks, enhanced with skip connections that preserve spatial details from the corresponding encoder layers. A Probabilistic Pruning Block (PPB) is applied at the end to remove outlier points and artifacts generated during reconstruction. To further enhance generalization under various adverse weather conditions, we introduce a Weather-Aware Block (WAB), which dynamically modulates intermediate features based on estimated degradation characteristics.

The key components of LTDNet—including positional encoding, spatial-frequency modeling, weather-aware adaptation, and point-wise pruning—are detailed below.

Positional Encoding

Global Positional Encoding When projecting 3D LiDAR points to a 2D range image, angular information (i.e., az-

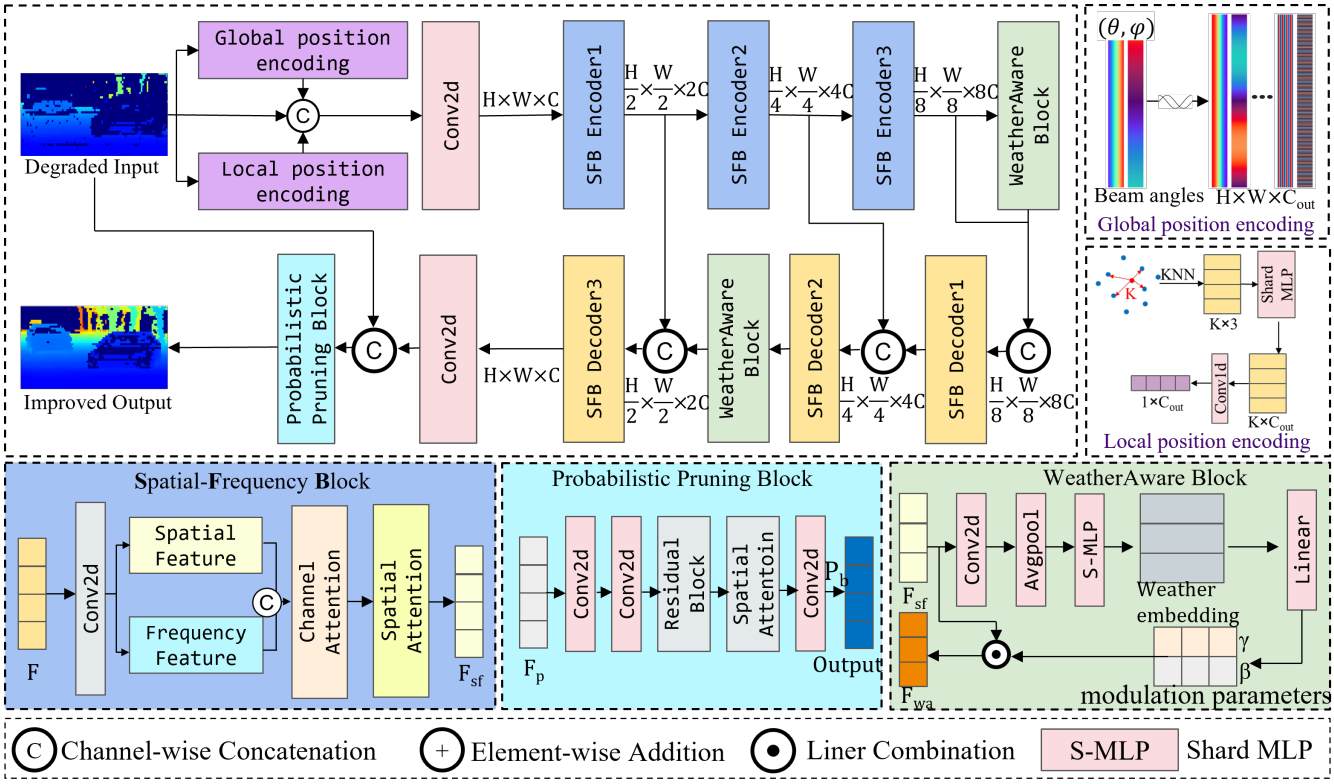

Figure 1: The architecture of the proposed network, LTDNet, which learns to map degraded LiDAR point clouds to clean point clouds. LTDNet adopts an encoder–decoder framework, augmented with global-local positional embeddings, a weather-aware block, and a probabilistic pruning-based post-processing block.

imuth θ and elevation φ) typically lost, weakening the spatial structure representation. While previous works(Liu et al. 2018; Zyrianov, Zhu, and Wang 2022) append raw angles (θ, φ) as additional channels, these low-frequency signals often lack the expressiveness needed to preserve fine-grained geometry. To address this, we encode angular positions using fourier positional encoding(Tancik et al. 2020), which decompose azimuth and elevation into multi-frequency components, introducing high-frequency biases that are linearly independent. We adopt a logarithmic spacing of frequency orders (e.g., $\log_2$ scale) to retain angular continuity. This creates a high-dimensional embedding that enhances the depth channel and improves global spatial encoding, as validated in R2DM(Nakashima and Kurazume 2024). A schematic of this process is shown in the top-right inset of Figure 1.

Local Positional Encoding To enhance local geometric understanding, we introduce a KNN-based relative position encoder. As shown in the right panel of Figure 1, for each point p, we identify its k-nearest neighbors and compute the range offsets, forming a k×3 coordinate matrix. This matrix is passed through a shared MLP to extract local geometric features, which are then compressed via a 1D convolution into a compact representation of the point's local context. This design explicitly models local structure and is implemented efficiently with the FAISS(Douze et al. 2025) for fast neighbor search.

Spatial-Frequency Block

The Spatial-Frequency Block (SFB) is designed to jointly exploit spatial locality and frequency-domain globality to enhance LiDAR feature representations under adverse weather conditions. Given an input feature $F \in \mathbb{R}^{C\times H\times W}$, we employ two parallel branches to extract spatial and spectral features, respectively.

In the spatial branch, the input tensor F is passed through two convolutional blocks $\mathcal{C}_S(\cdot)$, comprising standard convolution, batch normalization, and non-linear activation, to extract local structural features:

$$F_s = \mathcal{C}_S^2(F). \tag{1}$$

In the frequency branch, Discrete Fourier Transform (DFT) is applied to each channel independently, yielding amplitude $\mathbf{A}$ and phase $\mathbf{P}$ components. These are then fed into separate convolutional layers and recombined using Inverse DFT (IDFT):

$$F_f = \mathcal{F}^{-1}(\mathcal{C}_a^2(\mathbf{A}) + j\mathcal{C}_p^2(\mathbf{P})). \tag{2}$$

The outputs F_S and F_f are concatenated along the channel dimension and then fused sequentially through channel $\mathcal{A}_c$ and spatial attention $\mathcal{A}_s$ mechanisms to adaptively re-weight informative features, following the design of the CBAM module(Woo et al. 2018).

$$F_{sf} = \mathcal{A}_s(\mathcal{A}_c([F_s; F_f])). \tag{3}$$

The final output F_{sf} is a hybrid feature embedding that integrates structural continuity and frequency priors, facilitating more robust downstream perception.

Weather-Aware Block

Although LiDAR points degradation patterns under rain, fog, and snow share common traits, their subtle differences can undermine generalization. To address this, we propose a Weather-Aware Block (WAB) that dynamically adjusts feature representations based on learned degradation cues. As shown in the bottom panel of Figure 1, given an intermediate feature map F_{sf}, we first extract a compact weather embedding using global average pooling followed by a shared MLP:

$$F_{wb} = \mathrm{S} - \mathrm{MLP}(\mathrm{AvgPool}(F_{sf})). \tag{4}$$

This embedding is transformed into a pair of modulation parameters γ and β via a linear projection:

$$\gamma, \beta = Linear(F_{wb}), \tag{5}$$

where γ and β represent the modulation parameters used for scale scaling and bias translation, respectively. Subsequently, we apply Feature-wise Linear Modulation to modulate the original features:

$$F_{wa} = \gamma \cdot F_{sf} + \beta. \tag{6}$$

WAB enables the network to adaptively recalibrate intermediate activations under different weather-induced degradations.

Probabilistic Pruning Block

To further improve structural consistency in the reconstructed point cloud, we introduce a Probabilistic Pruning Block (PPB) to suppress physically implausible points such as floating artifacts or background noise. As illustrated in Figure 1(bottom), the block consists of two convolution layers with batch normalization and ReLU, followed by a lightweight ResBlock(He et al. 2016) to model contextual dependencies. A spatial attention(Woo et al. 2018) highlights regions likely to contain outliers.

A final convolution with sigmoid activation estimates a per-point retention probability $P_b \in [0, 1]$:

$$P_b = \sigma(\mathrm{Conv}(\mathrm{Attn}(\mathrm{ResBlock}(F_p)))). \tag{7}$$

Loss Function To train LTDNet and ensure its effectiveness in improving the quality of LiDAR point clouds, we define the loss function L as:

$$L = \frac{1}{N} \sum_{i=1}^{N} \left(|d_i - \hat{d}_i| \right), \tag{8}$$

where d_i and $\hat{d}_i$ represent the distance values in the clean range image and the predicted values of the model, respectively. N is the number of pixels in the range image.

Benchmark for Point Cloud Quality Improvement

Existing benchmarks for evaluating point cloud quality improvement methods under adverse weather primarily rely on synthetically degraded data generated from clean point clouds. For example, Robo3D(Kong et al. 2023) proposed datasets such as KITTI-C, nuScenes-C, and WOD-C by injecting artificial weather corruptions. However, these datasets lack real degraded point clouds as references, making it difficult to accurately assess the real benefits of improvement methods for downstream tasks.

To address this limitation, we construct IQA3D, a novel benchmark designed to evaluate point cloud quality improvement under adverse weather conditions. Built upon two public datasets—RADIATE and DENSE—it incorporates both real-world and simulated degradation scenarios. IQA3D consists of two complementary components:

• Simulated Degradation: We apply physically based weather degradation models(Kilic et al. 2021) to the clean point clouds to simulate three common types of adverse weather, thereby enabling fine-grained evaluation of point cloud quality improvement methods in a controlled setting.

• Real Degradation: We retain and integrate the original real-world adverse weather data from RADIATE and DENSE to evaluate the true gain of point cloud quality improvement methods for target detection models.

Weather Corruptions Generation To overcome the significant challenges and high costs associated with systematically collecting large-scale, strictly spatio-temporally aligned clear-adverse weather paired point cloud datasets under real-world, we employ LISA(Kilic et al. 2021), a widely used physics-based simulation method, for clear-to-adverse transformation. LISA leverages well-validated atmospheric radiative transfer and LiDAR signal propagation models to perform a physics-driven degradation of measured clear-weather point clouds, generating corresponding degraded point clouds under rainy, snowy, and foggy conditions.

Benchmark Configuration

Datasets We generate two simulated weather datasets—DENSE-W and RADIATE-W—each covering three weather types (rain, fog and snow), with three severity levels per type. In addition, we preserve the real degraded point clouds originally present in the RADIATE(Sheeny et al. 2021) and DENSE(Bijelic et al. 2020) datasets for real-world evaluation. As a result, each dataset is split into three subsets: (i) originally clean point clouds, (ii) synthetically degraded point clouds, and (iii) real-world degraded point clouds. Our quality improvement model is trained and validated on synthetic data, while downstream detection performance is evaluated on real-world data.

It is important to note that RADIATE's radar-coordinate 2D bounding boxes limit usability for LiDAR 3D tasks, while temporal misalignment between sensors causes spatial errors in coordinate conversions. We develop an automatic refinement pipeline that reconstructs 3D boxes via spatial association and height estimation, yielding annotations better

Method	R-RAIN		R-FOG		R-SNOW		D-RAIN		D-FOG		D-SNOW	
	DDI↑	NR-F1↑	DDI↑	NR-F1↑	DDI↑	NR-F1↑	DDI↑	NR-F1↑	DDI↑	NR-F1↑	DDI↑	NR-F1↑
LTDNet	**0.70**	91.6	**0.61**	**0.75**	91.3	**0.69**	**0.75**	91.3	**0.69**	89.2	**0.73**	92.7
DSOR	0.28	89.1	0.22	0.29	91.6	0.25	0.29	91.6	0.25	89.5	0.31	**94.2**
4DenoiseNet	0.31	90.2	0.32	0.34	90.3	0.30	0.34	90.3	0.30	87.9	0.36	92.5
3D-UnOutDet	0.35	**91.9**	0.31	0.37	**92.4**	0.31	0.37	**92.4**	0.31	**90.8**	0.36	93.6

Table 1: Quantitative results of point cloud quality improvement methods. R-RAIN and D-RAIN denote the rain-weather data from the RADIATE and DENSE datasets, respectively; the same naming convention applies to other weather conditions.

suited for LiDAR-based training and evaluation. Additional implementation details are available in the supplementary material.

Evaluated 3D Detection Methods We assess the effectiveness of our quality improvement model using three representative 3D object detectors: PV-RCNN(Shi et al. 2020), PointPillars(Lang et al. 2019), and CenterPoint(Yin, Zhou, and Krahenbuhl 2021). All detectors are trained using the OpenPCDet(Team 2020) with default configurations.

Baseline Methods

Since no existing methods have been specifically designed to improve LiDAR data degraded by adverse weather, this benchmark compares LTDNet with three representative denoising baselines: DSOR(Kurup and Bos 2021), 4DenoiseNet(Seppänen, Ojala, and Tammi 2022), and 3D-UnOutDet(Raisuddin, Gouigah, and Aksoy 2024).

Experiments

To evaluate the effectiveness of our proposed LTDNet for point cloud quality improvement under adverse weather conditions, we conduct two sets of experiments: (1) quality evaluation on degraded point clouds, and (2) downstream 3D object detection performance.

Point Cloud Quality Evaluation

We train and evaluate LTDNet using paired clean and synthetically degraded point clouds from the proposed DENSE-W and RADIATE-W datasets. For both DENSE-W and RADIATE-W, we divide the point cloud samples into training and testing sets with a 7:3 ratio. The projected range images have resolutions of 32 × 2250 (for RADIATE-W) and 64 × 900 (for DENSE-W), depending on the LiDAR sensor specifications. The model is trained for 200 epochs using the Adam optimizer with a learning rate of 1e-4. Experiments are conducted on a workstation equipped with an Intel i7-12700 CPU, 32GB RAM, and two NVIDIA 2080Ti GPUs.

To objectively assess point cloud quality improvement, we adopt two evaluation criteria: geometric fidelity and noise suppression. Specifically:

• Geometric Fidelity: We propose the Degradation De-gree Improvement (DDI) metric, which quantifies the relative improvement of the improved point cloud over the degraded input. This metric is defined as:

$$\mathrm{DDI} = 1 - \frac{\Delta_{imp}}{\Delta_{de}} = 1 - \frac{\|P_{imp} - P_{cl}\|_F}{\|P_{de} - P_{cl}\|_F}, \tag{9}$$

whereP_{cl}, P_{de}, and P_{imp} denote the clean, degraded, and improved point clouds, respectively, and $\|\cdot\|_F$ is the Frobenius norm. A DDI value of 1 indicates perfect improvement, while negative values penalize further distortion.

• Noise Suppression: We adopt the F1-score for noise removal (NR-F1) to measure the accuracy of noise filtering:

$$\mathrm{NR-F1} = 2 \times \frac{\mathrm{Precision} \times \mathrm{Recall}}{\mathrm{Precision} + \mathrm{Recall}}, \tag{10}$$

where Precision is the proportion of correctly retained clean points, and Recall is the proportion of successfully removed noise points.

Quantitative Results and Analysis Table 1 summarizes the performance of four quality improvement methods—LTDNet, DSOR, 4DenoiseNet, and 3D-UnOutDet—across three weather types (rain, fog, snow) and two datasets.

As shown in Table 1, all methods achieve high NR-F1 scores (∼90%), indicating effective and generalized noise suppression across weather conditions.

However, LTDNet significantly outperforms all baselines in terms of DDI, achieving values ranging from 0.61 to 0.81, which corresponds to a recovery of over 60% of the structural degradation. In contrast, baseline methods yield DDI scores around 0.30, demonstrating that denoising alone cannot resolve structural loss in degraded point clouds.

Additionally, LTDNet performs slightly better on the DENSE dataset, which may benefit from its denser point cloud distribution, offering richer structural cues for reconstruction.

Qualitative Results and Analysis Figure3 presents qualitative results of LTDNet on degraded point clouds collected under adverse weather. Specifically, we visualize a rain scene from the RADIATE-W, comparing the raw degraded point cloud (top) with the quality improved output from LTDNet (bottom).

The results demonstrate that LTDNet significantly improves point cloud quality across several dimensions. First, the continuity of road is notably improved, and occluded regions—such as areas behind vehicles—do not exhibit physically implausible artifacts like "see-through" points, indicating the network's ability to learn realistic occlusion-aware re-presentations. Furthermore, the improvement significantly improves object quality. As shown in the yellow

Detector	Pre -processing	RADIATE-W						DENSE-W					
		Rain		Fog		Snow		Rain		Fog		Snow	
		0–80m	0–40m	0–80m	0–40m	0–80m	0–40m	0–80m	0–40m	0–80m	0–40m	0–80m	0–40m
PV-RCNN	Non	38.43	64.38	33.30	59.16	36.53	62.12	44.52	70.28	39.74	63.40	42.89	68.32
	DSOR	39.17	65.82	33.67	59.21	37.98	63.01	44.67	71.03	39.65	63.57	43.91	70.56
	4DenoiseNet	39.41	65.76	33.72	59.54	37.52	62.87	44.55	71.62	39.83	63.48	43.31	71.22
	3D-UnOutDet	39.49	65.77	33.69	59.22	37.83	62.92	44.11	71.49	40.21	63.69	43.70	71.01
	LTDNet	**43.25**	**70.47**	**35.29**	**62.18**	**41.76**	**68.23**	**47.80**	**75.13**	**44.01**	**67.92**	**47.76**	**73.32**
PointPillars	Non	32.78	60.21	29.19	50.36	31.05	59.13	36.28	66.32	31.14	59.21	34.42	65.39
	DSOR	33.21	60.85	31.25	50.48	32.73	59.94	37.51	66.44	32.98	59.43	35.21	66.24
	4DenoiseNet	33.64	61.44	31.54	50.44	32.75	59.73	36.90	66.78	32.44	59.32	35.90	67.59
	3D-UnOutDet	33.69	61.52	31.80	50.52	32.81	59.82	37.37	66.75	32.61	59.65	35.95	67.82
	LTDNet	**38.54**	**66.54**	**33.12**	**52.94**	**35.87**	**65.82**	**40.21**	**72.76**	**37.19**	**64.82**	**40.15**	**71.24**
CenterPoint	Non	38.54	64.52	32.05	58.87	36.22	62.37	43.15	68.81	39.26	62.14	42.15	67.91
	DSOR	39.42	65.23	32.98	58.92	37.46	62.53	43.31	69.24	39.38	62.43	42.98	69.82
	4DenoiseNet	39.75	65.36	33.16	58.90	37.52	62.76	44.11	69.58	39.24	62.36	43.52	69.77
	3D-UnOutDet	39.65	65.33	33.12	58.92	37.56	62.68	43.97	69.64	39.82	62.81	43.97	69.21
	LTDNet	**42.87**	**70.36**	**35.21**	**61.47**	**42.20**	**68.92**	**46.93**	**74.32**	**45.39**	**67.43**	**47.92**	**73.26**

Table 2: Performance comparison of object detection under adverse weather with different data preprocessing methods. The table reports the Average Precision (AP) for the vehicle class at a 3D IoU threshold of 0.5.

zoom-in box, vehicle instances exhibit clearer contours, improved completeness, and more accurate geometric structures.

Experiments on Object Detection Methods

Quantitative Results and Analysis We evaluate the effectiveness of LTDNet in improving downstream 3D object detection performance under adverse weather conditions based on IQA3D benchmark. Three representative detectors are considered—PV-RCNN, PointPillars, and CenterPoint. All detectors are trained on clear-weather point clouds and tested on degraded inputs processed by different preprocessing algorithms, including DSOR, 4DenoiseNet, and 3D-UnOutDet. Raw degraded point clouds serve as the baseline. Detection performance is measured by AP@0.5 (vehicle class), under two distance ranges: 0–80m and 0–40m.

As summarized in Table 2, LTDNet consistently achieves the highest detection accuracy across all detector-backbone and weather combinations. While other baseline methods show minor improvements (typically less than 2 AP), LTDNet offers substantial gains of 4–6 AP points. For instance, on the RADIATE-W dataset in rainy scenes, PV-RCNN's AP increases from 38.43% (baseline) to 43.25% (+4.82%) using LTDNet. In snowy conditions, AP improves from 36.53% to 41.76%, outperforming other methods (e.g., DSOR: 37.98%, 4DenoiseNet: 37.52%).

Furthermore, LTDNet demonstrates robust cross-detector performance. In RADIATE-W, average AP gains are $\sim$4.1–4.4 points for PV-RCNN and CenterPoint, and $\sim$3.8 points for PointPillars. In DENSE-W, CenterPoint sees a gain of $\sim$5.5 points, while PV-RCNN and PointPillars improve by $\sim$4.8 and $\sim$4.6 points, respectively. This indicates weak coupling between LTDNet and the detection backbone, suggesting strong generalization and effective correction of weather-induced degradation regardless of detector architecture.

Qualitative Results and Analysis Figure 2 presents the qualitative result. We visualize the 3D object detection outcomes produced by the PV-RCNN model using three types of LiDAR point clouds as input: raw degraded point clouds captured under adverse weather conditions (top row), point clouds pre-processed by DSOR (middle row), and those refined by our LTDNet (bottom row). Ground truth bounding boxes are shown in blue, while predicted boxes are shown in red.

As highlighted by the green circles, under rain in both datasets, only the LTDNet-improved point clouds allow accurate detection of distant, partially degraded vehicles. In fog scenes from the RADIATE dataset, DSOR fails to detect a target vehicle (green circle) due to over-smoothing that mistakenly removes valid object points. In contrast, LTDNet preserves critical object geometry, leading to higher detection accuracy. Across both datasets under snowy conditions, the raw point clouds contain numerous false positives, likely caused by snow-induced noise. Both DSOR and LTDNet significantly reduce these spurious detections. Notably, LTDNet yields bounding boxes that more closely align with the ground truth, demonstrating its superior ability to enhance structural fidelity and detection performance in challenging weather scenarios.

Runtime

LTDNet demonstrates real-time performance, with an average inference time of 38ms on RADIATE-W and 41ms on DENSE-W, measured on an Intel i7-12700 CPU and an NVIDIA RTX 2080 Ti GPU.

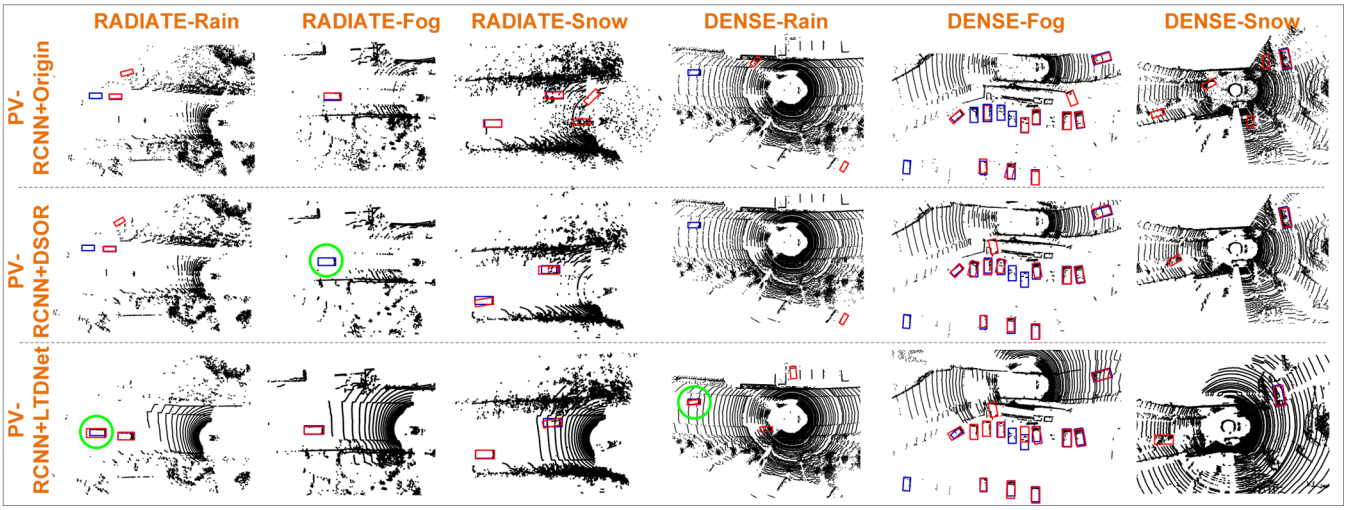

Figure 2: Target detection results of PV-RCNN under different preprocessing methods. Each row corresponds to a preprocessing method, from top to bottom: no preprocessing, DSOR, and our proposed LTDNet. Blue bounding box represents the ground truth, while red boxes represent model predictions.

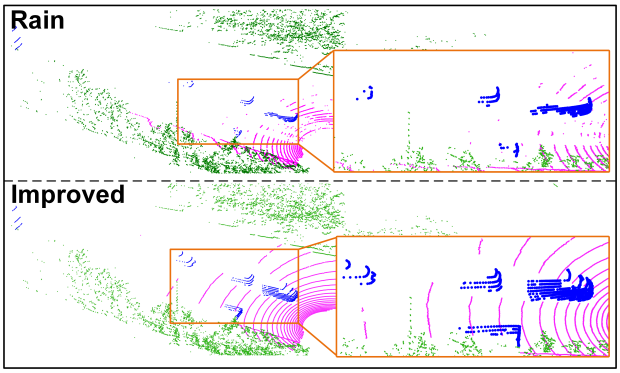

Figure 3: The quality improvement effect of LTDNet on degraded point clouds. Purple points are the ground, blue for vehicles, and green for other points such as trees and buildings. The objects in the yellow enlarged box are all vehicles.

Ablation Study

We conduct an ablation study to evaluate the contribution of each key module in LTDNet, as shown in Table 3. The 3D object detection performance (AP) using PV-RCNN is assessed under three adverse weather conditions. Model A, with only the PE module, achieves a baseline AP of 40.2. Introducing the SFB (Model B) improves this to 42.4, demonstrating the value of spatial–frequency feature fusion. Adding the PPB (Model C) further boosts AP to 43.1 by effectively removing artifacts and outliers. Finally, Model D, incorporating the WAB, reaches the highest AP of 43.6, highlighting the advantage of weather-aware feature modeling. These results show that each component contributes positively to LTDNet's overall performance.

Model	PE	SFB	PPB	WAB	AP
A	✓				40.2
B	✓	✓			42.4
C	✓	✓	✓		43.1
D	✓	✓	✓	✓	43.6

Table 3: Ablation of various components of LTDNet.

Conclusion

This paper presents LTDNet, a novel LiDAR point cloud quality improvement framework that learns an end-to-end mapping from degraded to clean point clouds, surpassing traditional denoising methods. LTDNet integrates position encoding, a spatial–frequency joint feature extraction block, probabilistic pruning, and a weather-aware refinement module to restore structural integrity and suppress noise. To evaluate restoration quality and its impact on downstream perception under real-world adverse weather, we introduce the IQA3D benchmark. Experiments show that LTDNet effectively corrects degradation from rain, fog, and snow, outperforming existing baselines and enhancing the weather robustness of 3D object detection models.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (Grant Nos. 52302401 and 52441205); the 111 Project on Information of Vehicle-Infrastructure Sensing and ITS (Grant No. B14043); and the Fundamental Research Funds for the Central Universities, CHD (Grant No. 300102245716). The authors extend their sincere thanks to all the supporters and contributors to this research.

References

Bernuth, A.; Volk, G.; and Bringmann, O. 2019. Simulating Photo-realistic Snow and Fog on Existing Images for Enhanced CNN Training and Evaluation. In *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, 41–46. [Online; accessed 2025-02-10].

Bijelic, M.; Gruber, T.; Mannan, F.; Kraus, F.; Ritter, W.; Dietmayer, K.; and Heide, F. 2020. Seeing Through Fog Without Seeing Fog: Deep Multimodal Sensor Fusion in Unseen Adverse Weather. ArXiv:1902.08913 [cs].

Brzozowski, M.; and Parczewski, K. 2023. Problems related to the operation of autonomous vehicles in adverse weather conditions. *Combustion Engines*, 62. ISBN: 2300-9896.

Burnett, K.; Yoon, D. J.; Wu, Y.; Li, A. Z.; Zhang, H.; Lu, S.; Qian, J.; Tseng, W.-K.; Lambert, A.; Leung, K. Y.; Schoellig, A. P.; and Barfoot, T. D. 2023. Boreas: A multi-season autonomous driving dataset. *The International Journal of Robotics Research*, 42(1-2): 33–42.

Chappell, B. A.; and Pereira, T. M. D. 2022. SunnyNet: A neural network approach to 3D non-LTE radiative transfer. *Astronomy & Astrophysics*, 658: A182. Publisher: EDP Sciences.

Charron, N.; Phillips, S.; and Waslander, S. L. 2018. Denoising of lidar point clouds corrupted by snowfall. In *2018 15th Conference on Computer and Robot Vision (CRV)*, 254–261. IEEE. [Online; accessed 2024-07-13].

Douze, M.; Guzhva, A.; Deng, C.; Johnson, J.; Szilvasy, G.; Mazaré, P.-E.; Lomeli, M.; Hosseini, L.; and Jégou, H. 2025. The Faiss library. ArXiv:2401.08281 [cs].

Dreissig, M.; Scheuble, D.; Piewak, F.; and Boedecker, J. 2023. Survey on LiDAR Perception in Adverse Weather Conditions. In *2023 IEEE Intelligent Vehicles Symposium (IV)*, 1–8. ISSN: 2642-7214.

Goswami, P.; Vaishnav, R.; Anand, T.; and Dayal, P. 2025. A Comprehensive Review on LiDAR Based 3D Deep Learning Object Detection Algorithms. In *2025 International Conference on Computer, Electrical & Communication Engineering (ICCECE)*, 1–6. IEEE. ISBN 979-8-3315-1389-4.

Hahner, M.; Sakaridis, C.; Bijelic, M.; Heide, F.; Yu, F.; Dai, D.; and Van Gool, L. 2022. Lidar snowfall simulation for robust 3d object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 16364–16374. [Online; accessed 2024-11-16].

He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778. Las Vegas, NV, USA: IEEE. ISBN 978-1-4673-8851-1. [Online; accessed 2025-08-01].

Heinzler, R.; Piewak, F.; Schindler, P.; and Stork, W. 2020. Cnn-based lidar point cloud de-noising in adverse weather. *IEEE Robotics and Automation Letters*, 5(2): 2514–2521. Publisher: IEEE titleTranslation: CNN.

Huang, H.; Yan, X.; Yang, J.; Cao, Y.; and Zhang, X. 2023. LIDSOR: A filter for removing rain and snow noise points from LiDAR point clouds in rainy and snowy weather. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 48: 733–740. Publisher: Copernicus GmbH.

Kilic, V.; Hegde, D.; Sindagi, V.; Cooper, A. B.; Foster, M. A.; and Patel, V. M. 2021. Lidar Light Scattering Augmentation (LISA): Physics-based Simulation of Adverse Weather Conditions for 3D Object Detection. ArXiv:2107.07004 [cs].

Kong, L.; Liu, Y.; Li, X.; Chen, R.; Zhang, W.; Ren, J.; Pan, L.; Chen, K.; and Liu, Z. 2023. Robo3D: Towards Robust and Reliable 3D Perception against Corruptions. 19994–20006. [Online; accessed 2025-07-03].

Kurup, A.; and Bos, J. 2021. DSOR: A Scalable Statistical Filter for Removing Falling Snow from LiDAR Point Clouds in Severe Winter Weather. Issue: arXiv:2109.07078 arXiv:2109.07078 [cs].

Lang, A. H.; Vora, S.; Caesar, H.; Zhou, L.; Yang, J.; and Beijbom, O. 2019. Pointpillars: Fast encoders for object detection from point clouds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 12697–12705. [Online; accessed 2024-09-12].

Liu, R.; Lehman, J.; Molino, P.; Petroski Such, F.; Frank, E.; Sergeev, A.; and Yosinski, J. 2018. An intriguing failing of convolutional neural networks and the CoordConv solution. In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc. [Online; accessed 2025-08-01].

Maddern, W.; Pascoe, G.; Linegar, C.; and Newman, P. 2017. 1 year, 1000 km: The Oxford RobotCar dataset. *The International Journal of Robotics Research*, 36(1): 3–15.

Mao, J.; Xue, Y.; Niu, M.; Bai, H.; Feng, J.; Liang, X.; Xu, H.; and Xu, C. 2021. Voxel Transformer for 3D Object Detection. 3164–3173. [Online; accessed 2025-11-12].

Nakashima, K.; and Kurazume, R. 2024. LiDAR Data Synthesis with Denoising Diffusion Probabilistic Models. ArXiv:2309.09256 [cs].

Piewak, F.; Pinggera, P.; Schafer, M.; Peter, D.; Schwarz, B.; Schneider, N.; Enzweiler, M.; Pfeiffer, D.; and Zollner, M. 2018. Boosting lidar-based semantic labeling by cross-modal training data generation. In *Proceedings of the European conference on computer vision (ECCV) workshops*, 0–0.

Piroli, A.; Dallabetta, V.; Kopp, J.; Walessa, M.; Meissner, D.; and Dietmayer, K. 2023. Energy-based detection of adverse weather effects in lidar data. *IEEE Robotics and Automation Letters*, 8(7): 4322–4329. ISBN: 2377-3766 publisher: IEEE.

Pitropov, M.; Garcia, D. E.; Rebello, J.; Smart, M.; Wang, C.; Czarnecki, K.; and Waslander, S. 2021. Canadian Adverse Driving Conditions dataset. *The International Journal of Robotics Research*, 40(4-5): 681–690.

Qi, C. R.; Su, H.; Mo, K.; and Guibas, L. J. 2017a. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 652–660.

Qi, C. R.; Yi, L.; Su, H.; and Guibas, L. J. 2017b. Pointnet++: Deep hierarchical feature learning on point sets in a

metric space. *Advances in neural information processing systems*, 30.

Raisuddin, A. M.; Cortinhal, T.; Holmblad, J.; and Aksoy, E. E. 2024. 3D-OutDet: A Fast and Memory Efficient Outlier Detector for 3D LiDAR Point Clouds in Adverse Weather. In *2024 IEEE Intelligent Vehicles Symposium (IV)*, 2862–2868. ISSN: 2642-7214.

Raisuddin, A. M.; Gouigah, I.; and Aksoy, E. E. 2024. 3d-unoutdet: A fast and efficient unsupervised snow removal algorithm for 3d lidar point clouds. *Authorea Preprints*.

Roriz, R.; Campos, A.; Pinto, S.; and Gomes, T. 2021. DIOR: A hardware-assisted weather denoising solution for LiDAR point clouds. *IEEE Sensors Journal*, 22(2): 1621–1628. Publisher: IEEE.

Seppänen, A.; Ojala, R.; and Tammi, K. 2022. 4denoisenet: Adverse weather denoising from adjacent point clouds. *IEEE Robotics and Automation Letters*, 8(1): 456–463. Publisher: IEEE.

Sheeny, M.; De Pellegrin, E.; Mukherjee, S.; Ahrabian, A.; Wang, S.; and Wallace, A. 2021. Radiate: A radar dataset for automotive perception in bad weather. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 1–7. IEEE. [Online; accessed 2024-07-28].

Shi, S.; Guo, C.; Jiang, L.; Wang, Z.; Shi, J.; Wang, X.; and Li, H. 2020. Pv-rcnn: Point-voxel feature set abstraction for 3d object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10529–10538. [Online; accessed 2024-09-12].

Sun, C.; Sun, P.; Wang, J.; Guo, Y.; and Zhao, X. 2024. Understanding lidar performance for autonomous vehicles under snowfall conditions. *IEEE Transactions on Intelligent Transportation Systems*, 25(11): 16462–16472.

Sun, P.; Sun, C.; Wan, L.; Wang, R.; and Zhao, X. 2022. Objects detection with 3-D roadside LiDAR under snowy weather. *IEEE Sensors Journal*, 22(23): 23051–23063.

Tancik, M.; Srinivasan, P.; Mildenhall, B.; Fridovich-Keil, S.; Raghavan, N.; Singhal, U.; Ramamoorthi, R.; Barron, J.; and Ng, R. 2020. Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in neural information processing systems*, 33: 7537–7547.

Team, O. D. 2020. *Openpcdet: An open-source toolbox for 3d object detection from point clouds*.

Teufel, S.; Gamerdinger, J.; Volk, G.; Gerum, C.; and Bringmann, O. 2023. Enhancing robustness of LiDAR-Based perception in adverse weather using point cloud augmentations. In *2023 IEEE Intelligent Vehicles Symposium (IV)*, 1–6. IEEE. ISBN 979-8-3503-4691-6.

Vargas, J.; Alsweiss, S.; Toker, O.; Razdan, R.; and Santos, J. 2021. An overview of autonomous vehicles sensors and their vulnerability to weather conditions. *Sensors*, 21(16): 5397. ISBN: 1424-8220 publisher: MDPI.

Woo, S.; Park, J.; Lee, J.-Y.; and Kweon, I. S. 2018. CBAM: Convolutional Block Attention Module. ArXiv:1807.06521 [cs].

Yan, X.; Yang, J.; Huang, H.; Liang, Y.; Ma, Y.; Li, Y.; and Zhang, Y. 2025. AdverseNet: A Unified LiDAR Point Cloud Denoising Network for Autonomous Driving in Adverse Weather. *IEEE Sensors Journal*. ISBN: 1530-437X publisher: IEEE.

Yan, Y.; Mao, Y.; and Li, B. 2018. Second: Sparsely embedded convolutional detection. *Sensors*, 18(10): 3337. ISBN: 1424-8220 publisher: Multidisciplinary Digital Publishing Institute.

Yin, T.; Zhou, X.; and Krahenbuhl, P. 2021. Center-Based 3D Object Detection and Tracking. 11784–11793. [Online; accessed 2025-08-01].

Zhao, X.; Wen, C.; Wang, Y.; Bai, H.; and Dou, W. 2024. TripleMixer: A 3D Point Cloud Denoising Model for Adverse Weather. ArXiv:2408.13802 [cs].

Zhou, Y.; and Tuzel, O. 2018. Voxelnet: End-to-end learning for point cloud based 3d object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4490–4499.

Zyrianov, V.; Zhu, X.; and Wang, S. 2022. Learning to Generate Realistic LiDAR Point Clouds. ArXiv:2209.03954 [cs].