# Rethinking Data Augmentation for Robust LiDAR Semantic Segmentation in Adverse Weather

Junsung Park<sup>1</sup>, Kyungmin Kim<sup>2</sup>, and Hyunjung Shim<sup>1\*</sup>

<sup>1</sup> Korea Advanced Institute of Science & Technology (KAIST) <sup>2</sup> Yonsei University

jshackist@kaist.ac.kr, kyungmin.kim@yonsei.ac.kr, kateshim@kaist.ac.kr

Abstract. Existing LiDAR semantic segmentation methods often struggle with performance declines in adverse weather conditions. Previous work has addressed this issue by simulating adverse weather or employing universal data augmentation during training. However, these methods lack a detailed analysis and understanding of how adverse weather negatively affects LiDAR semantic segmentation performance. Motivated by this issue, we identified key factors of adverse weather and conducted a toy experiment to pinpoint the main causes of performance degradation: (1) Geometric perturbation due to refraction caused by fog or droplets in the air and (2) Point drop due to energy absorption and occlusions. Based on these findings, we propose new strategic data augmentation techniques. First, we introduced a Selective Jittering (SJ) that jitters points in the random range of depth (or angle) to mimic geometric perturbation. Additionally, we developed a Learnable Point Drop (LPD) to learn vulnerable erase patterns with a Deep Q-Learning Network to approximate the point drop phenomenon from adverse weather conditions. Without precise weather simulation, these techniques strengthen the LiDAR semantic segmentation model by exposing it to vulnerable conditions identified by our data-centric analysis. Experimental results confirmed the suitability of the proposed data augmentation methods for enhancing robustness against adverse weather conditions. Our method achieves a notable 39.5 mIoU on the SemanticKITTIto-SemanticSTF benchmark, improving the baseline by 8.1%p and establishing a new state-of-the-art. Our code will be released at https: //github.com/engineerJPark/LiDARWeather.

Keywords: LiDAR Semantic Segmentation  $\cdot$  Robust Learning for Adverse Weather  $\cdot$  Data Augmentation

# 1 Introduction

LiDAR semantic segmentation is a fundamental task of 3D scene understanding, particularly in safety-critical applications like autonomous driving. However, existing LiDAR segmentation models [1,5,14,21,28,37] commonly lack robustness,

<sup>\*</sup> Corresponding author: Hyunjung Shim

showing significant performance degradation under adverse weather conditions such as snow, fog, rain, or wet surfaces.

To address this issue, recent studies [11, 33, 34] have introduced corruption benchmarks for adverse weather and proposed effective techniques for robust LiDAR segmentation against corruption. These robust methods are primarily divided into task-agnostic and simulation-based approaches. The task-agnostic approach [11, 25, 33] employs general machine-learning techniques for robustness without explicitly addressing LiDAR corruption caused by adverse weather. The simulation-based approach [3,7–9,23,26,35] artificially synthesizes weatherspecific data using physical equations for training. However, these efforts focus on detection tasks and consider only a single type of weather at a time.

While the simulation-based approach leverages the intrinsic properties of Li-DAR scans under target adverse weather conditions, simulating every weather type at all possible severities is impractical and often inaccurate. Instead, we adopt a data-centric perspective to analyze corrupted LiDAR data. For example, distortions caused by rain, snow, and fog often create similar patterns, as simulated in [9]. That is, all those adverse weather shows point drop patterns due to attenuation or occlusion by droplet [6,7,9]. While the simulation-based approach requires explicit modeling of the LiDAR for each weather condition, our data-centric approach potentially addresses complex distortions with a few unified patterns. Through existing research and our comprehensive analysis, we find that various adverse weather conditions create similar distortion patterns in Li-DAR data. Based on our analysis, we categorize these distortions into two types: (1) geometric perturbation and (2) point drop. Our toy experiments demonstrate that these distortions are closely related to segmentation performance degradation.

Rooted by the toy experiment, we introduce two novel and strategic data augmentation methods tailored to the LiDAR distortion caused by adverse weather. By incorporating these augmentations during training, we aim to bolster the model's robustness for each distortion type. For geometric perturbation, we introduce Selective Jittering (SJ). This method applies Additive Gaussian Noise (AGN) to alter the XYZ-coordinate and intensity within a selective local region. To handle a point drop pattern, we developed a Learnable Point Drop (LPD) that employs a Deep Q-Network (DQN) to strategically remove points. Our augmentation strategies are inspired by existing studies [9, 16, 27], which shows geometric perturbation typically involves small, random alterations to the original points. This insight led to our choice of jittering as a suitable technique for addressing the geometric perturbation. Also, our augmentation strategy for point drop is motivated by existing studies [7–9] which state that intensity and depth of LiDAR beams play a significant role. This prompts us to use a DQN to detect point drops and adapt to such point drop patterns.

Finally, we apply the proposed augmentations when training the LiDAR segmentation model with a clean source dataset. Experimental results demonstrate that our method achieves new state-of-the-art performance on SemanticSTF, a real adverse weather dataset. Specifically, we achieved an impressive +8.1%p mIoU gain over the baseline MinkowskiNet, nearly tripling the +2.5%p mIoU improvement from previous work [33]. Our method achieved remarkable performance gains, ranging from +0.4 to +10.3%p mIoU, across various model architectures (Minkowski, CENet, SPVCNN) and corruption dataset (SemanticSTF and SemanticKITTI-C). These results demonstrate the generalizability of the proposed augmentation methods.

In summary, this paper presents several key contributions.

- We identify two prevalent distortion types in LiDAR data caused by adverse weather, leading to performance degradation, through a data-centric analysis.
- We introduce two novel data augmentations tailored to each identified distortion type.
- Our method sets new state-of-the-art benchmarks on the SemanticKITTI to SemanticSTF benchmark. Notably, our improvement was +8.1%p mIoU without relying on precise simulations of adverse weather in LiDAR point inputs. This represents a tripling of the improvement over the baseline, compared to what previous methods have achieved [33].

# 2 Related Works

## 2.1 LiDAR Semantic Segmentation

Existing 3D LiDAR point cloud semantic segmentation methods can be categorized into three types based on the data representation: Point-based, Projectionbased, and Voxel-based.

Point-based methods [22, 28, 36] utilize the 3D points directly as input. KP-Conv [28] initially clustering local points, aggregating these local features, and then feeding them into kernel point convolutions. The Point Transformer [36] utilizes a transformer architecture to compute query points in each local region, obtained through k-nearest neighbors (kNN). The Point-Mixer [4] tried to adapt the MLP-Mixer [29] for point cloud applications. They achieve high performance but suffer from high computational costs due to the utilization of large-scale raw LiDAR data.

Projection-based methods [1, 10, 18] project LiDAR points into a 2D image and performs the semantic segmentation using architectures successful in a 2D image. RangeViT [1] directly adopts a ViT model pre-trained on 2D images, demonstrating that pre-trained power in 2D images can be effective as prior knowledge in range images. RangeFormer [10] proposes "RangeAug" to maximize the utility of range images created by projecting into 2D, producing multiple range image data to overcome the low performance in range image models. Projection-based methods achieve fast inference speed but present a sub-optimal performance due to missing information after the projection.

Voxel-based methods [5,14,37] perform efficient computations by dividing 3D space into a voxel grid and aggregating points within the same voxel. Minkowsk-iNet [5] voxelizes LiDAR points with a cubic grid and applies sparse convolution.

Cylinder3D [37] proposes cylindrical partitions, reflecting LiDAR's characteristic that the density of points depends on the distance. SphereFormer [14] utilizes a radial window and transformer architecture to aggregate long-range information and improve performance. Voxel-based methods achieve a balance between reasonable inference time and commendable segmentation performance.

# 2.2 LiDAR Data Augmentation

Inspired by 2D image augmentation, conventional LiDAR segmentation methods apply classic scaling, rotation, flipping, and translation to augment LiDAR data. Recently, several out-of-context augmentation techniques [12, 20, 30] that mix different LiDAR scans have been proposed. Mix3D [20] combines randomly selected two scans. Considering the sweeping mechanism of the LiDAR sensor, PolarMix [30] cuts LiDAR scans along the azimuth axis, then exchanges point cloud sectors and applies instance-level rotate-pasting. To reflect the spatial prior of the LiDAR point cloud, LaserMix [12] partitions LiDAR scans based on the laser beams and blends partitions from different LiDAR scans. Recent study [24] introduces a fast LiDAR domain augmentation module to address sensor-bias problems. To the best of our knowledge, our approach is the first augmentation method specifically designed to address data corruption under adverse weather conditions.

## 2.3 LiDAR under Adverse Weather Conditions

Robustness under harsh conditions is crucial in safety-critical applications. Adverse weather substantially degrade performance in real-world outdoor autonomous driving. Thus, there are several attempts to develop weather-robust models in fields such as 2D segmentation [15, 17], 3D detection [7, 8, 11], and 3D segmentation [11, 33]. Simulation-based approaches [7, 8] artificially synthesize data for single weather conditions through physical modeling and utilize it for training. We differ from these in that we do not model specific weather conditions explicitly. Recently proposed task-agnostic approaches [11, 25, 33] consider multiple weather conditions at once. However, they use general machine-learning techniques (such as teacher-student framework and feature prototype) to achieve robustness rather than specifically tackle LiDAR corruption caused by adverse weather. We differ from these in that we propose augmentation methods specifically tailored for adverse weather conditions based on the analysis of performance degradation in LiDAR data.

# 3 Finding Distortions to Augment

In this section, we aim to discuss the patterns of distortion that different adverse weather conditions impose on LiDAR data. Although adverse weather conditions are distinct in reality, studies have shown that their effects on LiDAR data often result in similar impacts. For instance, the distortions caused by "rain", "snow", and "fog" tend to produce similar point-missing patterns due to attenuation in the data, as demonstrated in [9]. Therefore, this section will focus on identifying the common distortion patterns caused by adverse weather through existing studies. Overall, existing studies describe the effect of adverse weather as four different types of distortion: (1) Point Drop due to energy absorption, (2) Occlusions caused by droplets of rain or snow and fog, (3) Geometric perturbation, and (4) Intensity distortion due to energy absorption.

#### 3.1 Distortion Factors from Adverse Weather

(D1) Point Drop. Several studies have explored how adverse weather conditions contribute to point drops in LiDAR data. Kilic *et al.* [9], Fersch *et al.* [6] and Shin *et al.* [26] describe point drops resulting from beam attenuation and beam missing due to droplets, fog, and frozen or wet ground. These studies collectively suggest that adverse weather conditions typically lead to point drops in LiDAR data.

(D2) Occlusions. Several studies have addressed occlusions caused by adverse weather conditions. Hahner *et al.* [7], Kilic *et al.* [9], Kong *et al.* [11] and Yan *et al.* [34] consider scenarios where beams colliding with snow collect signals at much shorter distances than the intended objects of collision. Upon reviewing these studies, we have concluded that adverse weather consistently leads to occlusions.

(D3) Geometric Perturbation. Some studies focus on geometric perturbation caused by adverse weather conditions. Kilic *et al.* [9], Li *et al.* [16] and Smith *et al.* [27] demonstrated geometric perturbation in adverse weather, such as fog, snow, and rain, by incorporating random noise into the coordinates. Through these studies, we have come to conclude that adverse weather universally causes geometric perturbation.

(D4) Intensity Distortion. Numerous studies have focused on intensity distortion caused by adverse weather conditions. Bijelic *et al.* [3], Shin *et al.* [26], Fersch *et al.* [6], Kong *et al.* [11], and Yan *et al.* [34] have collectively shown that adverse weather conditions like fog, wetness, and rain lead to a reduction in Li-DAR beam intensity, influencing the generation of synthetic data. Through these studies, we conclude that adverse weather commonly causes intensity distortion.

## **3.2** Toy Experiment

Based on previous research findings mentioned in Section 3.1, the types of distortions in LiDAR point clouds caused by adverse weather conditions converge into a set of common distortions. Therefore, from a data-centric perspective, the issues that we need to consider can be summarized as follows: (D1) Point Drop, (D2) Occlusions, (D3) Geometric Perturbation, and (D4) Intensity Distortion. Here, we aim to identify which distortion types negatively impact performance. To achieve this, we have generated four distortion types of toy synthetic data through the SemanticKITTI validation set.

Table 1: Results of toy experiments from *validation set* of SemanticKITTI [2]. Soft and Hard indicate the severity of the distortions.

Distortion	car	bi.cle	mt.cle	truck	oth-v.	pers.	bi.clst	mt.clst	road	parki.	sidew.	oth-g.	build.	fence	veget.	trunk	terra.	pole	traf.	mIoU
Clean	96.8	22.2	66.9	89.1	65.2	67.0	84.7	0.0	93.8	50.6	81.4	0.1	91.1	63.0	88.1	67.7	74.5	63.8	47.7	63.9
$\mathbf{D1}:\mathbf{Soft}$	96.3	17.9	63.4	87.1	63.7	64.3	81.4	0.0	92.7	47.5	79.7	0.1	90.1	62.3	87.3	64.2	74.1	61.1	41.6	61.8
D1 : Hard	86.1	0.2	12.7	21.0	40.0	10.1	24.4	0.0	4.2	10.2	11.1	0.0	71.7	47.8	77.0	34.6	27.6	30.4	17.6	27.7
D2: Soft	94.8	16.3	51.7	66.3	53.9	59.2	52.1	0.0	92.5	44.9	79.2	0.1	89.7	61.4	87.3	63.2	74.3	59.2	40.2	57.2
D2 : Hard	81.3	0.5	5.4	13.1	35.5	7.6	0.7	0.0	2.6	8.5	10.8	0.0	71.1	45.7	76.3	32.9	27.3	26.7	16.1	24.3
D3: Soft	96.2	15.4	56.7	58.6	51.3	51.8	78.1	0.0	66.3	33.9	46.1	0.0	86.1	61.6	84.4	62.0	56.8	61.5	44.0	53.2
D3 : Hard	93.8	9.7	38.3	19.4	31.4	35.2	55.2	0.0	9.9	12.3	12.2	0.0	48.9	40.8	71.4	55.5	31.2	58.7	40.8	35.0
D4: Soft	96.3	21.7	61.1	89.4	61.8	68.9	83.2	0.0	92.9	44.4	79.1	0.1	90.4	56.7	86.9	68.2	71.1	63.9	48.3	62.3
D4: Hard	95.0	17.2	50.4	81.4	57.0	64.9	79.7	0.0	90.7	38.8	64.7	0.9	88.4	45.3	83.5	62.8	50.5	60.1	47.9	56.8

- (D1) Point Drop: It considers the scenario where each LiDAR point disappears randomly and independently due to severe adverse weather. We removed individual points randomly to synthesize these data. We set the drop ratios at 0.5 and 0.9.
- (D2) Occlusions: The method assumes that occlusions occur predominantly in front of objects, mainly due to distortion from fog, snow, and rain. We synthesized this data by randomly selecting points and altering their depth to one-tenth of the original depth. The selecting ratios were also determined to be 0.5 and 0.9.
- (D3) Geometric Perturbation: This aspect assumes the distortion of point coordinates resulting from adverse weather conditions. We synthesized data by adding Gaussian noise to the coordinates of all points. The Gaussian noise levels were established at 0.05 and 0.25.
- (D4) Intensity Distortion: This distortion type assumes intensity attenuation caused by fog, raindrops, and snow particles. We synthesized this data by subtracting Gaussian noise from the intensity of all points. The Gaussian noise levels were established at 0.05 and 0.25.

We selected MinkowskiNet [5] for our toy experiment. This choice was based on existing research [11, 34], which indicates that MinkowskiNet is a standard and robust model.

In Tab. 1, it's evident that with increasing distortion in D1, D2, and D3, performance falls below half of the baseline. Contrarily, D4 maintains performance levels despite significant distortion. This performance degradation likely stems from changes in the local geometric structure where model computations occur. For this reason, D4 did not significantly impact performance.

Importantly, D2's notable performance drop is argued to result from a point drop akin to D1. It is due to occlusions altering the point's local geometric structure, an effect similar to point drop. This is evidenced in Fig. 1, where incorrect predictions from D2 show patterns resembling those in D1.

In summary, our observations indicate that geometric perturbation and point drop are the most impactful distortions from adverse weather. Data augmentation that replicates these phenomena could likely boost the model's robustness Rethinking Data Augmentation for Robust LiDAR Semantic Segmentation



**Fig. 1:** Visualization for the results of toy experiment. In all figures, green (or red) points indicate correct (or incorrect) predictions. Point drop, occlusion, and geometric perturbation are the main distortions that lead to performance degradation. Also, since the misprediction areas of point drop and occlusion largely overlap, it's reasonable to consider them as the same type of distortion.

against such conditions without explicit weather simulations. Thus, robust Li-DAR semantic segmentation models against adverse weather involve two key steps: (1) Training on specific adverse point drops detrimental to performance, and (2) Training with minor adjustments in point coordinates.

# 4 Methods

In this section, we propose solutions for the two challenges identified in Section 3.2. Our objective is to enhance the robustness of the model by addressing the common distortions exhibited by various adverse weather conditions, through data augmentation. This approach aims to achieve robustness without the necessity of simulating each adverse weather scenario explicitly.

As observed earlier, distortions resulting from adverse weather include point drop, occlusion, geometric perturbation, and intensity distortion. Of these, the main distortions identified as having the most detrimental impact on performance were (1) geometric perturbation caused by minor variations in point coordinates and (2) point drop caused by beam missing or occlusion. To address these challenges, we propose two techniques: Selective Jittering and Learnable Point Drop.

# 4.1 Selective Jittering

Selective Jittering is devised to address the first main distortion, geometric perturbation. Selective Jittering consists of two types: Depth-Selective Jittering (DSJ) and Angle-Selective Jittering (ASJ).

**Depth-Selective Jittering (DSJ).** DSJ employs a single scan to add Gaussian noise to the XYZ coordinates and the intensity value of points below a randomly selected depth range.

Angle-Selective Jittering (ASJ). It applies Gaussian noise to a randomly selected range of angles. By performing jittering across different ranges of depths or angles for each point, this method reflects the augmentation that some beams are affected by geometric perturbation in a non-uniform manner in adverse weather. This concept realistically represents the characteristics of transparent droplets like rain and snow [6].

A significant aspect of DSJ and ASJ is that they use only a single frame. DSJ and ASJ are efficient in data augmentation without additional LiDAR frames and can perform reasonable augmentation even when the given LiDAR data is not sequentially captured.

**Range Jittering (RJ).** Range Jittering has been proposed to simulate range distortion caused by droplets and fog, as mentioned in [9]. Unlike DSJ and ASJ, Range Jittering applies jittering only in the range direction of points. This method is used in place of using the original points in DSJ and ASJ.

### 4.2 Learnable Point Drop

Learnable Point Drop (LPD) was devised to address the second main distortion, point drop. LPD is designed to artificially create point drop due to occlusion, such as dense fog. LPD employs a Deep Q-Learning Network (DQN) [19] for identifying the drop ratio and drop region that lead to adverse effects on the model. The reward of the DQN is designed to identify point drops that increase the training loss and uncertainty of the LiDAR semantic segmentation model. Through LPD, the LiDAR semantic segmentation model is exposed to point drop scenarios caused by adverse weather conditions. Consequently, it learns to make accurate predictions even in the absence of points that are critical for segmentation in clean data environments. Since LPD is employed merely as a concept for data augmentation and exists as a separate module, it necessitates no alterations to the existing model's training scheme except for limiting the gradient norm to ensure the stability of DQN learning.

LPD module defines its current state by summing the loss  $L_{aug}$ , calculated from augmented data obtained through SJ, and the entropy  $H_{aug}$ , derived from logits. The loss  $L_{aug}$  is calculated using the original loss function employed by the model in use. The calculation of entropy is as follows:

$$H_{aug}(x) = -\frac{1}{N} \sum_{i=1}^{n} P(x_i) \log P(x_i)$$

In this process,  $x_i$  represents each point, and  $P(\cdot)$  is softmax to the logits. N denotes the number of points. LPD predicts the indices of points to drop, taking as input the sum of  $L_{aug}$  and  $H_{aug}$  added to the input point tensor. This process enables simultaneous learning of the drop ratio and drop region. Subsequently, the difference between the sum of the loss  $L_{LPD}$  and entropy  $H_{LPD}$ , obtained from the dropped points by LPD, and the sum of  $L_{aug}$  and  $H_{aug}$ , is defined as the reward.

In Comparison with random drop, random drop uniformly reduces points across all depths. Consequently, this approach is not suitable for simulating



**Fig. 2:** Description of Learnable Point Drop (LPD). (1) The LPD utilizes the loss  $L_{aug}$  and entropy  $H_{aug}$  derived from LiDAR segmentation model as its current state. (2) Upon receiving a point input, LPD performs a point drop, then recalculates the loss  $L_{LPD}$  and entropy  $H_{LPD}$ . (3) The difference between these new values and the previous ones is used to compute the reward.

point drops caused by fog, as it lacks the depth-specific characteristics of fogrelated distortions. Furthermore, randomly generated point drop scans may not achieve the severity of adverse weather conditions necessary to drop critical points that could deceive the model. A detailed comparison will be conducted in the supplementary material.

# 4.3 Overall Pipeline

The overall training procedure begins with data augmentation through SJ. The augmented points are input into the LiDAR semantic segmentation model to compute loss. The model then calculates logits for each point, applies softmax, and subsequently computes the entropy. The current state of the LPD module is set as the sum of the loss and entropy. The dropped points are then input into the LiDAR semantic segmentation model to calculate the corresponding loss and entropy. The difference between the sum of loss/entropy obtained from the augmented points and that obtained from the points determined by LPD is set as the reward. Through this process, the LPD module enables the LiDAR semantic segmentation model to effectively learn points distorted due to adverse weather. Overall process is detailed in Fig. 3

# 5 Experiments

### 5.1 Experimental Setup

**Dataset and Evaluation Metrics.** To assess robustness, we trained our model on 19 classes from SemanticKITTI and then validated its performance on *validation set* of SemanticSTF as same as [33]. The experiment process involves training on clean data followed by evaluating robustness in adverse conditions,



**Fig. 3:** The overall training process. (1) Initially, SJ is applied. (2) Following this, we calculate the loss, which leads to the update of the LPD state. (3) This process is followed by utilizing LPD to generate dropped points. (4) Subsequent to recalculating the loss, the reward is accordingly updated.

akin to existing research [11, 33, 34]. SemanticSTF is a dataset gathered under real-world adverse conditions such as rain, fog, and snow. Following the approach in [33], invalid data in SemanticSTF are mapped to an "ignore label". We utilized MinkowskiNet [5] for validation purposes due to its well-known robustness [11, 34]. Our evaluation metrics include the Intersection over Union (IoU) for each class and the mean IoU (mIoU) across all classes.

Implementation Details. MinkowskiNet-18/32width served as our baseline model. The learning rate was set at 0.24, with a weight decay of 0.0001. The mean and standard deviation of Gaussian noise used in SJ are 0 and 0.01 each. In cases of Learnable Point Drop, the norm of the gradient was limited to 100 for both the DQN and Segmentation model. All experiments were conducted using four A6000 GPUs for 15 epochs in batch size 2. The duration of all experiments ranged between 3 to 5 hours.

#### 5.2 Main Results

10

J.Park et al.

SemanticKITTI to SemanticSTF. According to Tab. 2 (a), our proposed method showed a significant improvement in SemanticSTF, with a +8.1 mIoU increase over the baseline and +5.6 mIoU over the recent competitor, PointDR. Our method demonstrated universally superior performance across all weather conditions, validating the validity of our choice in data distortion type and the effectiveness of our method. Our method substantially improved the performance of categories like *other vehicle*, *motorcyclist*, *sidewalk*, *pole*, and *traffic sign*, which had lower baseline performance, increasing their class IoU by about +5 to +10. Additionally, it significantly enhanced the mIoU for cars and persons by +19.0 and +9.6, respectively, compared to the baseline, which is crucial for safety in driving environments. More details are available in the supplementary material.

Also, in Tab. 2 (a), examining the performance improvement per weather condition, we note that the most significant increase in mIoU was in rain, with

**Table 2:** LiDAR segmentation results (mIoU) on the SemanticSTF validation set of models trained with (a) SemanticKITTI and (b) SynLiDAR [32]. D-fog and L-fog denote dense fog and light fog weather conditions in all experiments. \* denotes the reproduced result with the same segmentation backbone. The symbol ‡ indicates that the validation for model selection was performed on sequence 0 of SynLiDAR, rather than SemanticSTF. Best score is in **bold** and second best is underlined.

					(b) SynLiDAR $\rightarrow$ SemanticSTF								
(a) Semantick	ITTI-	$\rightarrow$ Sem	antics	STF		.0 .0	80	.e	wc				
	fog	fog	nin	MON	Methods	D-f	L-f	Ra	Sne	mIoU			
Methods	ģ	Ļ	Rá	5 mIoU	Oracle	51.9	54.6	57.9	53.7	54.7			
Oracle	51.9	54.6	57.9	53.7   54.7	Baseline	15.24	15.97	16.83	12.76	15.45			
Baseline LaserMix [13] PolarMix [31] PointDR <sup>*</sup> [33]	30.7 23.2 21.3 <b>37.3</b>	$30.1 \\ 15.5 \\ 14.9 \\ 33.5$	29.7 9.3 16.5 35.5	$\begin{array}{c ccccc} 25.3 & 31.4 \\ 7.8 & 14.7 \\ 9.3 & 15.3 \\ 26.9 & 33.9 \end{array}$	LaserMix [13] PolarMix [31] PointDR <sup>*</sup> [33] PointDR <sup>*</sup> ‡	$\begin{array}{c c} 15.32 \\ 16.47 \\ \underline{19.09} \\ 21.41 \end{array}$	17.95 18.69 20.28 <u>20.94</u>	18.55 19.63 <u>25.29</u> <b>25.48</b>	13.8 15.98 <u>18.98</u> <b>19.31</b>	$ \begin{array}{c c} 16.85 \\ 18.09 \\ 19.78 \\ \underline{20.47} \end{array} $			
Baseline+SJ+LPD Increments to baseline	$\frac{36.0}{+5.3}$	<b>37.5</b> +7.4	37.6 +7.9	<b>33.1 39.5</b> +7.8 +8.1	Baseline+SJ+LPD Increments to baseline Baseline+SJ+LPD‡ Increments to baseline	$  19.08 \\ +3.8 \\ 18.99 \\ +3.7 \\  $	20.65 + 4.7 <b>21.22</b> + 5.3	$21.97 + 5.1 \\ 23.14 + 6.3$	$17.27 + 4.5 \\ 17.28 + 4.5$	$\begin{vmatrix} 20.08 \\ +4.6 \\ 20.51 \\ +5.1 \end{vmatrix}$			

+7.9 mIoU. This result is indicative of our SJ augmentation effectively reflecting geometric perturbations caused by raindrops. This aspect will be further explored in the ablation study. Additionally, in snow, our method showed a remarkable increase of +7.8 mIoU, and in dense fog (+5.3 mIoU) and light fog (+7.4 mIoU), our methodology clearly outperformed the baseline. This indicates that our data augmentation effectively models the adverse impacts of weather on LiDAR semantic segmentation models. Detailed analysis of how our method enhances robustness against each weather condition, and the factors contributing to the improved performance, will be discussed in the ablation study.

Moreover, as previously mentioned, the training architectures designed to enhance robustness, such as those proposed in PointDR [11, 33, 34], can be orthogonally applied to our proposed data augmentation. Therefore, regardless of the choice of training architecture, the integration of our method is expected to yield superior performance compared to conventional approaches.

**SynLiDAR to SemanticSTF.** For SynLiDAR, unlike SemanticKITTI, where no specific validation set was available, sequence 0 was treated as the validation set. The best-performing models during these validation checkpoints were selected for reporting.

In contrast, PointDR demonstrated an unrealistic scenario in its baseline code, using SemanticSTF for validation during training and reporting only the highest-performing results. For a fair comparison using PointDR's criteria, our method shows a performance increase of +0.04 mIoU over PointDR. Using our more realistic protocol, our method demonstrates an even greater performance improvement of +0.32 mIoU against PointDR. Detailed results are shown in Tab. 2 (b).





**Fig. 4:** Qualitative results of our method from *validation set* of SemanticSTF. All models are trained in *train set* of SemanticKITTI. Green (or red) points indicate correct (or incorrect) predictions, showing our predictions are more accurate than PointDR in all weather, especially on *road* and *sidewalk*.

#### 5.3 Additional Experiments

Analysis of Proposed Component Methods. As shown in As Tab. 3, when ASJ was added to the baseline, we observed an improvement in robustness across all adverse weather conditions (increases of +2.6 mIoU in dense fog, +5.3 mIoU in light fog, +8.1 mIoU in rain, and +6.3 mIoU in snow), albeit with a -1.8 mIoU decrease in clean data performance. The introduction of DSJ further enhanced performance in dense fog (+4.1 mIoU), light fog (+6.3 mIoU), rain (+9.3 mIoU), snow (+4.6 mIoU), and an overall performance increase of +6.2 mIoU. Our Selective Jittering method's strong performance in rain conditions showcases its efficacy in rainy weather, where it considers geometric perturbations for selected beams instead of all.

The application of Range Jittering showed improvements in dense fog(+3.7 mIoU), light fog (+6.9 mIoU), rain(+6.0 mIoU), snow (+8.2 mIoU), and an overall performance increase of +7.3 mIoU, indicating the importance of simulating range distortion due to adverse weather as discussed in [9]. As seen earlier, in rainy weather, jittering only some beams aids in enhancing robustness. However, performing Range Jittering on the remaining original points leads to all points being jittered, reducing performance in rain conditions. Additionally, the improvement in snow conditions by using Range Jittering suggests that perturbing as many beams as possible is more effective for robustness enhancement in snowy weather.

The addition of LPD resulted in improvements in dense fog (+5.3 mIoU), light fog (+7.4 mIoU), rain (+7.9 mIoU), snow (+7.8 mIoU), and an overall increase of +8.1 mIoU against to the baseline. The substantial performance en-

Clean D-fog Snow L-fog Rain Methods mIoU Baseline 63.9 30.730.129.725.331.433.337.8 +ASJ62.1 (-1.8) 35.431.636.8 (+5.4) +DSJ63.0 (-0.9) 34.836.4**39.0** 29.9 37.6 (+6.2)+RJ33.437.035.7 **33.5** 61.2 (-2.7) 38.7 (+7.3) +LPD62.8 (-1.1) **36.0 37.5** 37.6 33.139.5 (+8.1)

**Table 3:** Experiments on the components of our methods with SemanticKITTI as source and SemanticSTF as target. The values in parentheses indicate the performance improvement or decrease over the baseline model.

hancements observed in conditions of dense and light fog indicate that the LPD effectively represents occlusion from foggy weather as intended. The performance reduction in snow conditions compared to the use of ASJ, DSJ, and Range Jittering (-0.4 mIoU) is likely due to LPD not exclusively targeting the frozen or wet ground for point dropout. LPD applies the point drop across all points, as shown in Fig. 5. This represents a limitation of LPD, necessitating the development of alternative methods to address this issue.

Overall, the use of all components led to a significant improvement in performance under adverse weather conditions, with an increase of +8.1 mIoU over the baseline at a reasonable cost of -1.1 mIoU in clean data. Ultimately, our method achieved state-of-the-art results on SemanticSTF, with an increase in mIoU across all adverse weather conditions compared to the baseline. The combined use of DSJ, ASJ, Range Jittering, and LPD demonstrated high performance, indicating that each component synergistically contributes to enhancing the model's robustness.



Fig. 5: Qualitative results of LPD from *validation set* of SemanticKITTI. It can be observed that the point drop facilitated by the LPD occurs not just locally but extends across all points. While LPD effectively mimics the occlusions and sparse point distributions seen in real adverse weather, it does not replicate the effect of wet ground surfaces.

Method	SemanticSTF	${\bf Semantic KITTI-C}$
CENet SPVCNN Minkowski	$14.2 \\ 28.1 \\ 31.4$	$49.3 \\ 52.5 \\ 53.0$
CENet+Ours SPVCNN+Ours Minkowski+Ours	$\begin{array}{ccc} 22.0 & (+7.8) \\ 38.4 & (+10.3) \\ 39.5 & (+8.1) \end{array}$	$\begin{array}{rrrr} 53.2 & (+3.9) \\ 52.9 & (+0.4) \\ 58.6 & (+5.6) \end{array}$

 Table 4: LiDAR segmentation results (mIoU) with the various architecture and datasets. All models were trained with SemanticKITTI.

**Experiments on Various Input Data Representation, Models and Datasets.** We evaluated our method using different data representations, architectures, and datasets, specifically: (1) projection-based CENet and (2) voxel-based SPVCNN for backbones, and (3) the synthetic corruption dataset SemanticKITTI-C [11]. Tab. 4 shows our results.

SemanticKITTI-C has a unique data distribution with more samples in classes, such as *sidewalks* and *vegetation*, than SemanticSTF. Our method consistently improves performance across various datasets and architectures. Specifically, for the SemanticSTF dataset, our method increases the mIoU of CENet by +7.8, SPVCNN by +10.3, and MinkowskiNet by +8.1 mIoU. Similarly, for the SemanticKITTI-C dataset, the improvements are +3.9 for CENet, +0.4 for SPVCNN, and +5.6 mIoU for MinkowskiNet. These consistent enhancements across different input data representations, models, and datasets highlight the robustness and generalizability of our approach.

# 6 Conclusion

Our study focused on analyzing LiDAR data distortions in adverse weather and enhancing LiDAR model robustness with a data-centric method. We identified that point drop and geometric distortion mainly affect model performance through toy experiments. We introduced novel data augmentation techniques, Selective Jittering, and Learnable Point Drop, leading to state-of-the-art performance on SemanticSTF. We showed that our methodology significantly contributes to robustness enhancement without complex data simulations. We expect that our research will contribute to improving the stability and reliability of LiDAR semantic segmentation.

Acknowledgements. This work was supported by IITP grant funded by MSIT (No.2021-0-02068, Artificial Intelligence Innovation Hub and RS-2019-II190075, Artificial Intelligence Graduate School Program(KAIST)) and KEIT grant funded by MOTIE (No.2022-0-00680, No.2022-0-01045), NRF funded by the MSIP (NRF-2022R1A2C3011154, RS-2023-00219019, RS-2023-00240135) and Samsung Electronics Co., Ltd (IO230508-06190-01).

# References

- Ando, A., Gidaris, S., Bursuc, A., Puy, G., Boulch, A., Marlet, R.: Rangevit: Towards vision transformers for 3d semantic segmentation in autonomous driving. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5240–5250 (2023)
- Behley, J., Garbade, M., Milioto, A., Quenzel, J., Behnke, S., Stachniss, C., Gall, J.: Semantickitti: A dataset for semantic scene understanding of lidar sequences. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 9297–9307 (2019)
- Bijelic, M., Gruber, T., Ritter, W.: A benchmark for lidar sensors in fog: Is detection breaking down? In: 2018 IEEE intelligent vehicles symposium (IV). pp. 760–767. IEEE (2018)
- Choe, J., Park, C., Rameau, F., Park, J., Kweon, I.S.: Pointmixer: Mlp-mixer for point cloud understanding. In: European Conference on Computer Vision. pp. 620–640. Springer (2022)
- Choy, C., Gwak, J., Savarese, S.: 4d spatio-temporal convnets: Minkowski convolutional neural networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 3075–3084 (2019)
- Fersch, T., Buhmann, A., Koelpin, A., Weigel, R.: The influence of rain on small aperture lidar sensors. In: 2016 German Microwave Conference (GeMiC). pp. 84– 87. IEEE (2016)
- Hahner, M., Sakaridis, C., Bijelic, M., Heide, F., Yu, F., Dai, D., Van Gool, L.: Lidar snowfall simulation for robust 3d object detection. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 16364– 16374 (2022)
- Hahner, M., Sakaridis, C., Dai, D., Van Gool, L.: Fog simulation on real lidar point clouds for 3d object detection in adverse weather. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 15283–15292 (2021)
- Kilic, V., Hegde, D., Sindagi, V., Cooper, A.B., Foster, M.A., Patel, V.M.: Lidar light scattering augmentation (lisa): Physics-based simulation of adverse weather conditions for 3d object detection. arXiv preprint arXiv:2107.07004 (2021)
- Kong, L., Liu, Y., Chen, R., Ma, Y., Zhu, X., Li, Y., Hou, Y., Qiao, Y., Liu, Z.: Rethinking range view representation for lidar segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 228–240 (2023)
- Kong, L., Liu, Y., Li, X., Chen, R., Zhang, W., Ren, J., Pan, L., Chen, K., Liu, Z.: Robo3d: Towards robust and reliable 3d perception against corruptions. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 19994–20006 (2023)
- Kong, L., Ren, J., Pan, L., Liu, Z.: Lasermix for semi-supervised lidar semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21705–21715 (2023)
- Kong, L., Ren, J., Pan, L., Liu, Z.: Lasermix for semi-supervised lidar semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21705–21715 (2023)
- Lai, X., Chen, Y., Lu, F., Liu, J., Jia, J.: Spherical transformer for lidar-based 3d recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17545–17555 (2023)
- Lee, S., Son, T., Kwak, S.: Fifo: Learning fog-invariant features for foggy scene segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18911–18921 (2022)

- 16 J.Park et al.
- Li, Y., Duthon, P., Colomb, M., Ibanez-Guzman, J.: What happens for a tof lidar in fog? IEEE Transactions on Intelligent Transportation Systems 22(11), 6670–6681 (2020)
- Li, Z., Wu, X., Wang, J., Guo, Y.: Weather-degraded image semantic segmentation with multi-task knowledge distillation. Image and Vision Computing 127, 104554 (2022)
- Milioto, A., Vizzo, I., Behley, J., Stachniss, C.: Rangenet++: Fast and accurate lidar semantic segmentation. In: 2019 IEEE/RSJ international conference on intelligent robots and systems (IROS). pp. 4213–4220. IEEE (2019)
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Riedmiller, M., Fidjeland, A.K., Ostrovski, G., et al.: Human-level control through deep reinforcement learning. nature 518(7540), 529–533 (2015)
- Nekrasov, A., Schult, J., Litany, O., Leibe, B., Engelmann, F.: Mix3d: Out-ofcontext data augmentation for 3d scenes. In: 2021 international conference on 3d vision (3dv). pp. 116–125. IEEE (2021)
- Puy, G., Boulch, A., Marlet, R.: Using a waffle iron for automotive point cloud semantic segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3379–3389 (2023)
- 22. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 652–660 (2017)
- Qian, K., Zhu, S., Zhang, X., Li, L.E.: Robust multimodal vehicle detection in foggy weather using complementary lidar and radar signals. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 444–453 (2021)
- Ryu, K., Hwang, S., Park, J.: Instant domain augmentation for lidar semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9350–9360 (2023)
- Saltori, C., Galasso, F., Fiameni, G., Sebe, N., Ricci, E., Poiesi, F.: Cosmix: Compositional semantic mix for domain adaptation in 3d lidar segmentation. In: European Conference on Computer Vision. pp. 586–602. Springer (2022)
- Shin, J., Park, H., Kim, T.: Characteristics of laser backscattering intensity to detect frozen and wet surfaces on roads. Journal of Sensors 2019(1), 8973248 (2019)
- 27. Smith, B.E., Gardner, A., Schneider, A., Flanner, M.: Modeling biases in laseraltimetry measurements caused by scattering of green light in snow. Remote sensing of environment **215**, 398–410 (2018)
- Thomas, H., Qi, C.R., Deschaud, J.E., Marcotegui, B., Goulette, F., Guibas, L.J.: Kpconv: Flexible and deformable convolution for point clouds. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 6411–6420 (2019)
- Tolstikhin, I.O., Houlsby, N., Kolesnikov, A., Beyer, L., Zhai, X., Unterthiner, T., Yung, J., Steiner, A., Keysers, D., Uszkoreit, J., et al.: Mlp-mixer: An allmlp architecture for vision. Advances in neural information processing systems 34, 24261–24272 (2021)
- Xiao, A., Huang, J., Guan, D., Cui, K., Lu, S., Shao, L.: Polarmix: A general data augmentation technique for lidar point clouds. Advances in Neural Information Processing Systems 35, 11035–11048 (2022)
- Xiao, A., Huang, J., Guan, D., Cui, K., Lu, S., Shao, L.: Polarmix: A general data augmentation technique for lidar point clouds. Advances in Neural Information Processing Systems 35, 11035–11048 (2022)

Rethinking Data Augmentation for Robust LiDAR Semantic Segmentation

- Xiao, A., Huang, J., Guan, D., Zhan, F., Lu, S.: Transfer learning from synthetic to real lidar point cloud for semantic segmentation. In: Proceedings of the AAAI conference on artificial intelligence. vol. 36, pp. 2795–2803 (2022)
- 33. Xiao, A., Huang, J., Xuan, W., Ren, R., Liu, K., Guan, D., El Saddik, A., Lu, S., Xing, E.P.: 3d semantic segmentation in the wild: Learning generalized models for adverse-condition point clouds. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9382–9392 (2023)
- Yan, X., Zheng, C., Xue, Y., Li, Z., Cui, S., Dai, D.: Benchmarking the robustness of lidar semantic segmentation models. International Journal of Computer Vision pp. 1–24 (2024)
- Yang, D., Liu, Z., Jiang, W., Yan, G., Gao, X., Shi, B., Liu, S., Cai, X.: Realistic rainy weather simulation for lidars in carla simulator. arXiv preprint arXiv:2312.12772 (2023)
- Zhao, H., Jiang, L., Jia, J., Torr, P.H., Koltun, V.: Point transformer. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 16259– 16268 (2021)
- Zhu, X., Zhou, H., Wang, T., Hong, F., Ma, Y., Li, W., Li, H., Lin, D.: Cylindrical and asymmetrical 3d convolution networks for lidar segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 9939–9948 (2021)