

# 3D Semantic Segmentation in the Wild: Learning Generalized Models for Adverse-Condition Point Clouds


Aoran Xiao<sup>1</sup>, Jiaxing Huang<sup>1</sup>, Weihao Xuan<sup>2</sup>, Ruijie Ren<sup>3</sup>, Kangcheng Liu<sup>1</sup>

Dayan Guan<sup>4</sup>, Abdulmotaleb El Saddik<sup>4,6</sup>, Shijian Lu<sup>1,†</sup>, Eric Xing<sup>4,5</sup>

<sup>1</sup>Nanyang Technological University <sup>2</sup>Waseda University <sup>3</sup>Technical University of Denmark

<sup>4</sup>Mohamed bin Zayed University of Artificial Intelligence

<sup>5</sup>Carnegie Mellon University <sup>6</sup>University of Ottawa



(a) A LiDAR scan captured in a snowy day

(b) Point-level annotations

Figure 1. We introduce SemanticSTF, an adverse-weather LiDAR point cloud dataset with dense point-level annotations that can be exploited for the study of point cloud semantic segmentation under all-weather conditions (including fog, snow, and rain). The graph on the left shows one scan sample captured on a snowy day, and the one on the right shows the corresponding point-level annotations.

## Abstract

*Robust point cloud parsing under all-weather conditions is crucial to level-5 autonomy in autonomous driving. However, how to learn a universal 3D semantic segmentation (3DSS) model is largely neglected as most existing benchmarks are dominated by point clouds captured under normal weather. We introduce SemanticSTF, an adverse-weather point cloud dataset that provides dense point-level annotations and allows to study 3DSS under various adverse weather conditions. We study all-weather 3DSS modeling under two setups: 1) domain adaptive 3DSS that adapts from normal-weather data to adverse-weather data; 2) domain generalizable 3DSS that learns all-weather 3DSS models from normal-weather data. Our studies reveal the challenge while existing 3DSS methods encounter adverse-weather data, showing the great value of SemanticSTF in steering the future endeavor along this very meaningful research direction. In addition, we design a domain randomization technique that alternatively randomizes the geometry styles of point clouds and aggregates their embeddings, ultimately leading to a generalizable model that can improve 3DSS under various adverse weather effectively. The SemanticSTF and related codes are available at <https://github.com/xiaoaoran/SemanticSTF>.*

## 1. Introduction

3D LiDAR point clouds play an essential role in semantic scene understanding in various applications such as self-driving vehicles and autonomous drones. With the recent advance of LiDAR sensors, several LiDAR point cloud datasets [2, 11, 49] such as SemanticKITTI [2] have been proposed which greatly advanced the research in 3D semantic segmentation (3DSS) [19, 41, 62] for the task of point cloud parsing. As of today, most existing point cloud datasets for outdoor scenes are dominated by point clouds captured under normal weather. However, 3D vision applications such as autonomous driving require reliable 3D perception under all-weather conditions including various adverse weather such as fog, snow, and rain. How to learn a weather-tolerant 3DSS model is largely neglected due to the absence of related benchmark datasets.

Although several studies [3, 33] attempt to include adverse weather conditions in point cloud datasets, such as the STF dataset [3] that consists of LiDAR point clouds captured under various adverse weather, these efforts focus on object detection benchmarks and do not provide any point-wise annotations which are critical in various tasks such as 3D semantic and instance segmentation. To address this gap, we introduce *SemanticSTF*, an adverse-weather point cloud dataset that extends the STF Detection Benchmark by providing point-wise annotations of 21 semantic categories, as illustrated in Fig. 1. Similar to STF, SemanticSTF cap-

<sup>†</sup> Corresponding author

tures four typical adverse weather conditions that are frequently encountered in autonomous driving including dense fog, light fog, snow, and rain.

SemanticSTF provides a great benchmark for the study of 3DSS and robust point cloud parsing under adverse weather conditions. Beyond serving as a well-suited test bed for examining existing fully-supervised 3DSS methods that handle adverse-weather point cloud data, SemanticSTF can be further exploited to study two valuable weather-tolerant 3DSS scenarios: 1) domain adaptive 3DSS that adapts from normal-weather data to adverse-weather data, and 2) domain generalizable 3DSS that learns all-weather 3DSS models from normal-weather data. Our studies reveal the challenges faced by existing 3DSS methods while processing adverse-weather point cloud data, highlighting the significant value of SemanticSTF in guiding future research efforts along this meaningful research direction.

In addition, we design PointDR, a new baseline framework for the future study and benchmarking of all-weather 3DSS. Our objective is to learn robust 3D representations that can reliably represent points of the same category across different weather conditions while remaining discriminative across categories. However, robust all-weather 3DSS poses two major challenges: 1) LiDAR point clouds are typically sparse, incomplete, and subject to substantial geometric variations and semantic ambiguity. These challenges are further exacerbated under adverse weather conditions, with many missing points and geometric distortions due to fog, snow cover, etc. 2) More noises are introduced under adverse weather due to snow flicks, rain droplets, etc. PointDR addresses the challenges with two iterative operations: 1) *Geometry style randomization* that expands the geometry distribution of point clouds under various spatial augmentations; 2) *Embedding aggregation* that introduces contrastive learning to aggregate the encoded embeddings of the randomly augmented point clouds. Despite its simplicity, extensive experiments over point clouds of different adverse weather conditions show that PointDR achieves superior 3DSS generalization performance.

The contribution of this work can be summarized in three major aspects. *First*, we introduce SemanticSTF, a large-scale adverse-weather point cloud benchmark that provides high-quality point-wise annotations of 21 semantic categories. *Second*, we design PointDR, a point cloud domain randomization baseline that can be exploited for future study and benchmarking of 3DSS under all-weather conditions. *Third*, leveraging SemanticSTF, we benchmark existing 3DSS methods over two challenging tasks on domain adaptive 3DSS and domain generalized 3DSS. The benchmarking efforts lay a solid foundation for future research on this highly meaningful problem.

## 2. Related Works

**3D semantic segmentation** aims to assign point-wise semantic labels for point clouds. It has been developed rapidly over the past few years, largely through the development of various deep neural networks (DNNs) such as standard convolutional network for projection-based methods [9, 30, 46, 50, 59], multi-layer perceptron (MLP)-based networks [19, 34, 34], 3D voxel convolution-based networks [7, 62], or hybrid networks [6, 27, 41, 51, 57]. While existing 3DSS networks are mainly evaluated over normal weather point clouds, their performance for adverse weather point clouds is far under-investigated. The proposed SemanticSTF closes the gap and provides a solid ground for the study and evaluation of all-weather 3DSS. By enabling investigations into various new research directions, SemanticSTF represents a valuable tool for advancing the field.

**Vision recognition under adverse conditions.** Scene understanding under adverse conditions has recently attracted increasing attention due to the strict safety demand in various outdoor navigation and perception tasks. In 2D vision, several large-scale datasets have been proposed to investigate perceptions tasks in adverse visual conditions including localization [29], detection [56], and segmentation [36]. On the other hand, learning 3D point clouds of adverse conditions is far under-explored due to the absence of comprehensive dataset benchmarks. The recently proposed datasets such as STF [3] and CAD3 [33] contain LiDAR point clouds captured under adverse weather conditions. However, these studies focus on the object detection task [15, 16] with bounding-box annotations, without providing any point-wise annotations. Our introduced SemanticSTF is the first large-scale dataset that consists of LiDAR point clouds in adverse weather conditions with high-quality dense annotations, to the best of our knowledge.

**Domain generalization** [4, 31] aims to learn a generalizable model from single or multiple related but distinct source domains where target data is inaccessible during model learning. It has been widely studied in 2D computer vision tasks [1, 21, 26, 61] while few studies explore it in point cloud learning. Recently, [25] studies domain generalization for 3D object detection by deforming point clouds via vector fields. Differently, this work is the first attempt that explores domain generalization for 3DSS.

**Unsupervised domain adaptation** is a method of transferring knowledge learned from a labeled source domain to a target domain by leveraging the unlabeled target data. It has been widely studied in 2D image learning [12, 14, 20, 22–24] and 3D point clouds [15, 16, 28, 39, 52, 53, 58]. Recently, domain adaptive 3D LiDAR segmentation has drawn increasing attention due to the challenge in point-wise annotation. Different UDA approaches have been designed to mitigate discrepancies across LiDAR point clouds of different domains. For example, [46, 60] project point

clouds into depth images and leverage 2D UDA techniques while [37, 48, 49, 55] directly work in the 3D space. However, these methods either work for *synthetic-to-real* UDA scenarios [46, 49] or *normal-to-normal* point cloud adaptation [55], ignoring *normal-to-adverse* adaptation which is highly practical in real applications. Our SemanticSTF dataset fills up this blank and will inspire more development of new algorithms for normal-to-adverse adaptation.

### 3. The SemanticSTF Dataset

#### 3.1. Background

LiDAR sensors send out laser pulses and measure their flight time based on the echoes it receives from targets. The travel distance as derived from the time-of-flight and the registered angular information (between the LiDAR sensors and the targets) can be combined to compute the 3D coordinates of target surface which form point clouds that capture the 3D shape of the targets. However, the active LiDAR pulse system can be easily affected by the scattering media such as particles of rain droplets and snow [10, 18, 32, 35], leading to shifts of measured distances, variation of echo intensity, point missing, etc. Hence, point clouds captured under adverse weather usually have clear distribution discrepancy as compared with those collected under normal weather as illustrated in Fig. 1. However, existing 3DSS benchmarks are dominated by normal-weather point clouds which are insufficient for the study of universal 3DSS under all-weather conditions. To this end, we propose SemanticSTF, a point-wise annotated large-scale adverse-weather dataset that can be explored for the study of 3DSS and point cloud parsing under various adverse weather conditions.

#### 3.2. Data Selection and Split

We collect SemanticSTF by leveraging the STF benchmark [3], a multi-modal adverse-weather dataset that was jointly collected in Germany, Sweden, Denmark, and Finland. The data in STF have multiple modalities including LiDAR point clouds and they are collected under various adverse weather conditions such as snow and fog. However, STF provides bounding-box annotations only for the study of 3D detection tasks. In SemanticSTF, we manually selected 2,076 scans captured by a Velodyne HDL64 S3D LiDAR sensor from STF that cover various adverse weather conditions including 694 snowy, 637 dense-foggy, 631 light-foggy, and 114 rainy (all rainy LiDAR scans in STF). During the selection, we pay special attention to the geographical diversity of the point clouds aiming for minimizing data redundancy. We ignore the factor of daytime/nighttime since LiDAR sensors are robust to lighting conditions. We split SemanticSTF into three parts including 1,326 full 3D scans for training, 250 for validating, and 500 for testing. All three splits have approximately the same

proportion of LiDAR scans of different adverse weathers.

#### 3.3. Data Annotation

Point-wise annotation of LiDAR point clouds is an extremely laborious task due to several factors, such as 3D view changes, inconsistency between point cloud display and human visual perception, sweeping occlusion, point sparsity, etc. However, point-wise annotating of adverse-weather point clouds is even more challenging due to two new factors. *First*, the perceived distance shifts under adverse weather often lead to various geometry distortions in the collected points which make them different from those collected under normal weather. This presents significant challenges for annotators who must recognize various objects and assign a semantic label to each point. *Second*, LiDAR point clouds collected under adverse weather often contain a significant portion of invalid regions that consist of indiscernible semantic contents (e.g., thick snow cover) that make it difficult to identify the ground type. The existence of such invalid regions makes point-wise annotation even more challenging.

We designed a customized labeling pipeline to handle the annotation challenges while performing point-wise annotation of point clouds in SemanticSTF. Specifically, we first provide labeling instructions and demo annotations and train a team of professional annotators to provide point-wise annotations of a set of selected STF LiDAR scans. To achieve reliable high-quality annotations, the annotators leverage the corresponding 2D camera images and Google Street views as extra references while identifying the category of each point in this initial annotation process. After that, the annotators cross-check their initial annotations for identifying and correcting labeling errors. At the final stage, we engaged professional third parties who provide another round of annotation inspection and correction.

Annotation of SemanticSTF is a highly laborious and time-consuming task. For instance, while labeling downtown areas with the most complex scenery, it took an annotator an average of 4.3 hours to label a single LiDAR scan. Labeling a scan captured in a relatively simpler scenery, such as a highway, also takes an average of 1.6 hours. In addition, an additional 30-60 minutes are required per scan for verification and correction by professional third parties. In total, annotating the entire SemanticSTF dataset takes over 6,600 man-hours.

While annotating SemanticSTF, we adopted the same set of semantic classes as in the widely-studied semantic segmentation benchmark, SemanticKITTI [2]. Specifically, we annotate the 19 evaluation classes of SemanticKITTI, which encompass most traffic-related objects in autonomous driving scenes. Additionally, following [36], we label points with indiscernible semantic contents caused by adverse weather (e.g. ground covered by snowdrifts) as *invalid*. Fur-




Figure 2. Number of annotated points per class in SemanticSTF.

thermore, we label points that do not belong to the 20 categories or are indistinguishable as *ignored*, which are not utilized in either training or evaluations. Detailed descriptions of each class can be found in the appendix.

### 3.4. Data Statistics

SemanticSTF consists of point-wise annotations of 21 semantic categories, and Fig. 2 shows the detailed statistics of the point-wise annotations. It can be seen that classes *road*, *sidewalk*, *building*, *vegetation*, and *terrain* appear most frequently whereas classes *motor*, *motorcyclist*, and *bicyclist* have clearly lower occurrence frequency. Such class imbalance is largely attributed to the various object sizes and unbalanced distribution of object categories in transportation scenes, and it is also very common in many existing benchmarks. Overall, the statistics and distribution of different object categories are similar to that of other 2D and 3D semantic segmentation benchmarks such as Cityscapes [8], ACDC [36], and SemanticKITTI [2].

To the best of our knowledge, SemanticSTF is the first large-scale adverse-weather 3DSS benchmark that provides high-quality point-wise annotations. Table 1 compares it with several existing point cloud datasets that have been widely adopted for the study of 3D detection and semantic segmentation. We can observe that existing datasets are either collected under normal weather conditions or collected for object detection studies with bounding-box annotations only. 3DSS benchmark under adverse weather is largely blank, mainly due to the great challenge in point-wise annotations of adverse-weather point clouds as described in previous subsections. From this sense, SemanticSTF fills up this blank by providing a large-scale benchmark and test bed which will be very useful to future research in universal 3DSS under all weather conditions.

### 3.5. Data illustration

Fig. 3 provides examples of point cloud scans captured under adverse weather conditions in SemanticSTF (in row 1) as well as the corresponding annotations (in row 2). Compared with normal-weather point clouds, point clouds captured under adverse weather exhibit four distinct properties: 1) Snow coverage and snowflakes under *snowy* weather introduce many white points (labeled as “invalid”) as illustrated in Fig. 3(a). The thick snow coverage may lead to object deformation as well; *Rainy* conditions may cause specular reflection of laser signals from water on the ground

Dataset	#Cls	Type	Annotation	Fog	Rain	Snow
KITTI [13]	8	real	bounding box	✗	✗	✗
nuScenes [5]	23	real	bounding box	✗	✗	✗
Waymo [40]	4	real	bounding box	✗	✗	✗
STF [3]	5	real	bounding box	✓	✓	✓
SemanticKITTI [2]	25	real	point-wise	✗	✗	✗
nuScenes-LiDARSeg [11]	32	real	point-wise	✗	✗	✗
Waymo-LiDARSeg [40]	21	real	point-wise	✗	✗	✗
SynLiDAR [49]	32	synth.	point-wise	✗	✗	✗
SemanticSTF (ours)	21	real	point-wise	✓	✓	✓

Table 1. Comparison of SemanticSTF against existing outdoor LiDAR benchmarks. #Cls means the class number.

and produce many noise points as shown in Fig.3(b); 3) *Dense fog* may greatly reduce the working range of LiDAR sensors, leading to small spatial distribution of the collected LiDAR points as illustrated in Fig. 3(c); 4) Point clouds under *light fog* have similar characteristics as normal-weather point clouds as illustrated in Fig. 3(d). The distinct properties of point clouds under different adverse weather introduce different types of domain shift from normal-weather point clouds which complicate 3DSS greatly as discussed in Section 5. They also verify the importance of developing universal 3DSS models that can perform well under all weather conditions.

## 4. Point Cloud Domain Randomization

Leveraging SemanticSTF, we explore domain generalization (DG) for semantic segmentation of LiDAR point clouds under all weather conditions. Specifically, we design PointDR, a domain randomization technique that helps to train a generalizable segmentation model from normal-weather point clouds that can work well for adverse-weather point clouds in SemanticSTF.

### 4.1. Problem Definition

Given labeled point clouds of a source domain  $\mathcal{S} = \{S_k = \{x_k, y_k\}\}_{k=1}^K$  where  $x$  represents a LiDAR point cloud scan and  $y$  denotes its point-wise semantic annotations, the goal of domain generalization is to learn a segmentation model  $F$  by using the source-domain data only that can perform well on point clouds from an unseen target domain  $\mathcal{T}$ . We consider a 3D point cloud segmentation model  $F$  that consists of a feature extractor  $E$  and a classifier  $G$ . Note under the setup of domain generalization, target data will not be accessed in training as they could be hard and even impossible to acquire at the training stage.

### 4.2. Point Cloud Domain Randomization

Inspired by domain randomization studies in 2D computer vision research [42, 43], we explore how to employ domain randomization for learning domain generalizable models for point clouds. Specifically, we design PointDR,






Figure 3. Examples of LiDAR point cloud scans captured under different adverse weather including snow, rain, dense fog, and light fog (the first row) and corresponding dense annotations in SemanticSTF (the second row).

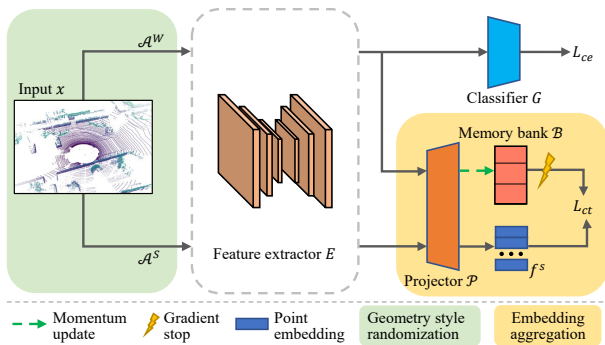


Figure 4. The framework of our point cloud randomization method (PointDR): *Geometry style randomization* creates different point cloud views with various spatial perturbations while *embedding aggregation* encourages the feature extractor to aggregate randomized point embeddings to learn perturbation-invariant representations, ultimately leading to a generalizable segmentation model.

a point cloud randomization technique that consists of two complementary designs including *geometry style randomization* and *embedding aggregation* as illustrated in Fig. 4. **Geometry style randomization** aims to enrich the geometry styles and expand the distribution of training point cloud data. Given a point-cloud scan  $x$  as input, we apply weak and strong spatial augmentation to obtain two copies of  $x$  including a weak-view  $x^w = \mathcal{A}^W(x)$  and a strong-view  $x^s = \mathcal{A}^S(x)$ . For the augmentation schemes of  $\mathcal{A}^W$ , we follow existing supervised learning methods [41] and adopt the simple random rotation and random scaling. While for the augmentation schemes of  $\mathcal{A}^S$ , we further adopt random dropout, random flipping, random noise perturbation, and random jittering on top of  $\mathcal{A}^W$  to obtain a more diverse and complex copy of the input point cloud scan  $x$ .

**Embedding aggregation** aims to aggregate encoded embeddings of randomized point clouds for learning domain-

invariant representations. We adopt contrastive learning [17] as illustrated in Fig. 4. Given the randomized point clouds  $x^w$  and  $x^s$ , we first feed them into the *feature extractor*  $E$  and a *projector*  $\mathcal{P}$  (a two-layer MLP) which outputs normalized point feature embeddings  $f^w$  and  $f^s$ , respectively ( $f = \mathcal{P}(E(x))$ ).  $\bar{f}_C^w \in \mathbb{R}^{D \times C}$  ( $D$ : feature dimension;  $C$ : number of semantic classes) is then derived by class-wise averaging the feature embeddings  $f^w$  in a batch, which is stored in a memory bank  $\mathcal{B} \in \mathbb{R}^{D \times C}$  that has no backpropagation and is momentum updated by iterations (i.e.,  $\mathcal{B} \leftarrow m \times \mathcal{B} + (1 - m) \times \bar{f}_C^w$  with a momentum coefficient  $m$ ). Finally, we employ each point feature embedding  $f_i^s$  of the strong-view  $f^s$  as query and feature embeddings in  $\mathcal{B}$  as keys for contrastive learning, where the key sharing the same semantic class as the query is positive key  $\mathcal{B}_+$  and the rest are negative keys. The contrastive loss is defined as

$$\mathcal{L}_{ct} = \frac{1}{N} \sum_{i=1}^N -\log \frac{\exp(f_i^s \mathcal{B}_+ / \tau)}{\sum_{j=1}^C \exp(f_i^s \mathcal{B}_j / \tau)} \quad (1)$$

where  $\tau$  is a temperature hyper-parameter [47]. Note there is no back-propagation for the “ignore” class in optimizing the contrastive loss.

Contrastive learning pulls point feature embeddings of the same classes closer while pushing away point feature embeddings of different classes. Therefore, optimizing the proposed contrastive loss will aggregate randomized point cloud features and learn perturbation-invariant representations, ultimately leading to a robust and generalizable segmentation model. The momentum-updated memory bank provides feature prototypes of each semantic class for more robust and stable contrastive learning.

Combining the supervised cross-entropy loss  $\mathcal{L}_{ce}$  for weakly-augmented point clouds in Eq. 1, the overall train-

Methods	car	bicle	mtcle	truck	oth-v.	pers.	bi.clst	mt.clst	road	parki.	sidew.	oth-g.	build.	fence	veget.	trunk	terra.	pole	traf.	D-fog	L-fog	Rain	Snow	mIoU
Oracle	89.4	42.1	0.0	59.9	61.2	69.6	39.0	0.0	82.2	21.5	58.2	45.6	86.1	63.6	80.2	52.0	77.6	50.1	61.7	51.9	54.6	57.9	53.7	54.7
SemanticKITTI→SemanticSTF																								
Baseline	55.9	0.0	0.2	1.9	10.9	10.3	<b>6.0</b>	0.0	61.2	10.9	32.0	0.0	67.9	41.6	49.8	27.9	40.8	<b>29.6</b>	17.5	29.5	26.0	28.4	21.4	24.4
Dropout [38]	62.1	0.0	<b>15.5</b>	3.0	<b>11.5</b>	5.4	2.0	0.0	58.4	12.8	26.7	1.1	72.1	43.6	52.9	<b>34.2</b>	43.5	28.4	15.5	29.3	25.6	29.4	24.8	25.7
Perturbation	<b>74.4</b>	0.0	0.0	<b>23.3</b>	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	25.9
PolarMix [48]	57.8	<b>1.8</b>	3.8	16.7	3.7	26.5	0.0	<b>2.0</b>	65.7	2.9	32.5	0.3	71.0	<b>48.7</b>	53.8	20.5	45.4	25.9	15.8	29.7	25.0	28.6	25.6	26.0
MMD [26]	63.6	0.0	2.6	0.1	11.4	<b>28.1</b>	0.0	0.0	<b>67.0</b>	<b>14.1</b>	37.9	0.3	67.3	41.2	57.1	27.4	47.9	28.2	16.2	30.4	28.1	<b>32.8</b>	25.2	26.9
PCL [54]	65.9	0.0	0.0	17.7	0.4	8.4	0.0	0.0	59.6	12.0	35.0	<b>1.6</b>	<b>74.0</b>	47.5	<b>60.7</b>	15.8	<b>48.9</b>	26.1	<b>27.5</b>	28.9	27.6	30.1	24.6	26.4
PointDR (Ours)	67.3	0.0	4.5	19.6	9.0	18.8	2.7	0.0	62.6	12.9	<b>38.1</b>	0.6	73.3	43.8	56.4	32.2	45.7	28.7	27.4	<b>31.3</b>	<b>29.7</b>	31.9	<b>26.2</b>	<b>28.6</b>
SynLiDAR→SemanticSTF																								
Baseline	27.1	<b>3.0</b>	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 [38]	28.0	<b>3.0</b>	1.4	9.6	0.0	17.1	0.8	0.7	<b>34.2</b>	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	<b>17.9</b>	5.1	16.3	16.7	19.3	13.4	15.2
PolarMix [48]	<b>39.2</b>	1.1	1.2	8.3	<b>1.5</b>	17.8	0.8	0.7	23.3	1.3	17.5	0.4	45.2	24.8	<b>46.2</b>	20.1	38.7	7.6	1.9	16.1	15.5	19.2	15.6	15.7
MMD [26]	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	<b>43.5</b>	9.9	2.6	17.3	16.3	20.0	12.7	15.1
PCL [54]	30.9	0.8	1.4	10.0	0.4	23.3	<b>4.0</b>	<b>7.9</b>	28.5	1.3	<b>17.7</b>	<b>1.2</b>	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 (Ours)	37.8	2.5	<b>2.4</b>	<b>23.6</b>	0.1	<b>26.3</b>	2.2	3.3	27.9	<b>7.7</b>	17.5	0.5	<b>47.6</b>	<b>25.3</b>	45.7	<b>21.0</b>	37.5	<b>17.9</b>	<b>5.5</b>	<b>19.5</b>	<b>19.9</b>	<b>21.1</b>	<b>16.9</b>	<b>18.5</b>

Table 2. Experiments on domain generalization with SemanticKITTI [2] or SynLiDAR [49] as source and SemanticSTF as target.

ing objective of PointDR can be formulated by:

$$\mathcal{L}_{\text{PointDR}} = \mathcal{L}_{ce} + \lambda_{ct} \mathcal{L}_{ct} \quad (2)$$

## 5. Evaluation of Semantic Segmentation

SemanticSTF can be adopted for benchmarking different learning setups and network architectures on point cloud segmentation. We perform experiments over two typical learning setups including domain generalization and unsupervised domain adaptation. In addition, we evaluate several state-of-the-art point-cloud segmentation networks to examine their generalization capabilities.

### 5.1. Domain Generalization

We first study domain generalizable point cloud segmentation. For DG, we can only access an annotated source domain during training and the trained model is expected to generalize well to *unseen* target domains. Leveraging SemanticSTF, we build two DG benchmarks and examine how PointDR helps learn a universal 3DSS model that can work under different weather conditions.

The first benchmark is *SemanticKITTI* [2]  $\rightarrow$  *SemanticSTF* where SemanticKITTI is a large-scale real-world 3DSS dataset collected under normal weather conditions. This benchmark serves as a solid testing ground for evaluating domain generalization performance from normal to adverse weather conditions. The second benchmark is *SynLiDAR* [49]  $\rightarrow$  *SemanticSTF* where SynLiDAR is a large-scale synthetic 3DSS dataset. The motivation of this benchmark is that learning a universal 3DSS model from synthetic point clouds that can work well across adverse weather is of high research and application value considering the

challenges in point cloud collection and annotation. Note this benchmark is more challenging as the domain discrepancy comes from both normal-to-adverse weather distribution shift and synthetic-to-real distribution shift.

**Setup.** We use all 19 evaluating classes of SemanticKITTI in both domain generalization benchmarks. The category of *invalid* in SemanticSTF is mapped to the *ignored* since SemanticKITTI and SynLiDAR do not cover this category. We adopt MinkowskiNet [7] (with TorchSparse library [41]) as the backbone model, which is a sparse convolutional network that provides state-of-the-art performance with decent efficiency. We adopt the evaluation metrics of Intersection over the Union (IoU) for each segmentation class and the mean IoU (mIoU) over all classes. All experiments are run over a single NVIDIA 2080Ti (11GB). More implementation details are provided in the appendix.

**Baseline Methods.** Since domain generalizable 3DSS is far under-explored, there is little existing baseline that can be directly adopted for benchmarking. We thus select two closely related approaches as baseline to evaluate the proposed PointDR. The first approach is data augmentation and we select three related augmentation methods including *Dropout* [38] that randomly drops out points to simulate LiDAR points missing in adverse weather, *Noise perturbation* that adds random points in the 3D space to simulate noise points as introduced by particles like falling snow, and *PolarMix* [48] that mixes point clouds of different sources for augmentation. The second approach is to adapt 2D domain generalization methods for 3DSS. We select two 2D domain generalization methods including the widely studied *MMD* [26] and the recently proposed *PCL* [54].

**Results.** Table 2 shows experimental results over the valida-

Method	$\mathcal{L}_{ce}$	$\mathcal{L}_{ct}$	$\mathcal{B}$	mIoU
Baseline	✓			24.4
PointDR-CT	✓	✓		27.4
PointDR	✓	✓	✓	28.6

Table 3. Ablation study of PointDR over domain generalized segmentation task SemanticKITTI→SemanticSTF.

tion set of SemanticSTF. For both benchmarks, the *Baseline* is a source-only model that is trained by using the training data of SemanticKITTI or SynLiDAR. We can see that the *Baseline* achieves very low mIoU while evaluated over the validation set of SemanticSTF, indicating the large domain discrepancy between point clouds of normal and adverse weather conditions. In addition, all three data augmentation methods improve the model generalization consistently but the performance gains are limited especially for the challenging benchmark SynLiDAR→SemanticSTF. The two 2D generalization methods both help SemanticKITTI → SemanticSTF clearly but show very limited improvement over SynLiDAR → SemanticSTF. The proposed PointDR achieves the best generalization consistently across both benchmarks, demonstrating its superior capability to learn perturbation-invariant point cloud representations and effectiveness while handling all-weather 3DSS tasks.

We also evaluate the compared domain generalization methods over each individual adverse weather condition as shown in Table 2. It can be observed that the three data augmentation methods work for data captured in rainy and snowy weather only. The 2D generalization method MMD shows clear effectiveness for point clouds under dense fog and rain while PCL works for point clouds under rainy and snowy weather instead. We conjecture that the performance variations are largely attributed to the different properties of point clouds captured under different weather conditions. For example, more points are missing in rain while object points often deform due to the covered snow (more illustrations are provided in the appendix). Such data variations lead to different domain discrepancies across weather which further leads to different performances of the compared methods. As PointDR learns perturbation-tolerant representations, it works effectively across different adverse weather conditions. We also provide qualitative results, please refer to the appendix for details.

**Ablation study.** We study different PointDR designs to examine how they contribute to the overall generalization performance. As Table 3 shows, we report three models over the benchmark “SemanticKITTI → SemanticSTF”: 1) *Baseline* that is trained with  $\mathcal{L}_{ce}$ . 2) *PointDR-CT* that is jointly trained with  $\mathcal{L}_{ce}$  and  $\mathcal{L}_{ct}$  without using the memory bank  $\mathcal{B}$ . 3) The complete *PointDR* that is trained with  $\mathcal{L}_{ce}$ ,  $\mathcal{L}_{ct}$  and the memory bank  $\mathcal{B}$ . We evaluate the three models over the validation set of SemanticSTF and Table 3

shows experimental results. We can see that the *Baseline* performs poorly at 24.4% due to clear domain discrepancy between point clouds of normal weather and adverse weather. Leveraging the proposed contrastive loss,  $\mathcal{L}_{ct}$  achieves clearly better performance at 27.4%, indicating that learning perturbation-invariance is helpful for universal LiDAR segmentation of all-weather conditions. On top of that, introducing the momentum-updated memory bank  $\mathcal{B}$  further improves the segmentation performance at 28.6%. This is because the feature embeddings in  $\mathcal{B}$  serve as the class prototypes which help the optimization of the segmentation network, finally leading to more robust representations of 3DSS that perform better over adverse weather point clouds.

## 5.2. Domain Adaptation

We also study SemanticSTF over a domain adaptive point cloud segmentation benchmark SemanticKITTI → SemanticSTF. Specifically, we select four representative UDA methods including ADDA [44], entropy minimization (Ent-Min) [45], self-training [63], and CoSMix [37] for adaptation from the source SemanticKITTI [2] toward the target SemanticSTF. Following the state-of-the-art [37, 48, 49] on synthetic-to-real adaptation, we adopt MinkowskiNet [7] as the segmentation backbone for all compared methods. Table 4 shows experimental results over the validation set of SemanticSTF. We can see that all UDA methods outperform the *Source-only* consistently under the normal-to-adverse adaptation setup. At the other end, the performance gains are still quite limited, showing the great improvement space along domain adaptive 3DSS from normal to adverse weather conditions.

In addition, we examined the adaptability of the four UDA methods in relation to each individual adverse weather condition. Specifically, we trained each of the four methods for adaptation from SemanticKITTI to SemanticSTF data for each adverse weather condition. Table 5 shows the experimental results over the validation set of SemanticSTF. We can see all four methods outperform the *Source-only* method under *Dense-fog* and *Light-fog*, demonstrating their effectiveness in mitigating domain discrepancies. However, for *rain* and *Snow*, only CoSMix achieved marginal performance gains while the other three UDA methods achieved limited performance improvements. We conjecture that snow and rain introduce large deformations on object surfaces or much noise, making adaptation from normal to adverse weather more challenging. CoSMix works in the input space by directly mixing source and target points, allowing it to perform better under heavy snow and rain which have larger domain gaps. However, all methods achieved relatively low segmentation performance, indicating the significance of our research and the large room for improvement in our constructed benchmarks.

Methods	car	bicle	mt.cle	truck	oth-v.	pers.	bi.clst	mt.clst	road	parki.	sidew.	oth-g.	build.	fence	veget.	trunk	terra.	pole	traf.	mIoU
Oracle	89.4	42.1	0.0	59.9	61.2	69.6	39.0	0.0	82.2	21.5	58.2	45.6	86.1	63.6	80.2	52.0	77.6	50.1	61.7	54.7
Source-only	64.8	0.0	0.0	13.8	1.8	5.0	2.1	0.0	62.7	7.5	34.0	0.0	66.7	36.2	53.9	31.3	44.3	24.0	14.2	24.3
ADDA [44]	65.6	0.0	0.0	21.0	1.3	2.8	1.3	<b>16.7</b>	64.7	1.2	35.4	0.0	66.5	41.8	<b>57.2</b>	32.6	42.2	23.3	26.4	26.3
Ent-Min [45]	69.2	0.0	10.1	31.0	5.3	2.8	<b>2.6</b>	0.0	<b>65.9</b>	2.6	35.7	0.0	<b>72.5</b>	<b>42.8</b>	52.4	32.5	<b>44.7</b>	24.7	21.1	27.2
Self-training [63]	<b>71.5</b>	0.0	10.3	<b>33.1</b>	7.4	5.9	1.3	0.0	65.1	6.5	<b>36.6</b>	0.0	67.8	41.3	51.7	<b>32.9</b>	42.9	25.1	25.0	27.6
CoSMix [37]	65.0	<b>1.7</b>	<b>22.1</b>	25.2	<b>7.7</b>	<b>33.2</b>	0.0	0.0	64.7	<b>11.5</b>	31.1	<b>0.9</b>	62.5	37.8	44.6	30.5	41.1	<b>30.9</b>	<b>28.6</b>	<b>28.4</b>

Table 4. Comparison of state-of-the-art domain adaptation methods on SemanticKITTI→SemanticSTF adaptation. SemanticKITTI serves as the source domain and the entire SemanticSTF including all four weather conditions serves as the target domain.

Method	Dense-fog	Light-fog	Rain	Snow
Source-Only	26.9	25.2	27.7	23.5
ADDA [44]	31.5	27.9	27.4	23.4
Ent-Min [45]	31.4	28.6	30.3	24.9
Self-training [63]	<b>31.8</b>	29.3	27.9	25.1
CoSMix [37]	31.6	<b>30.3</b>	<b>33.1</b>	<b>32.9</b>

Table 5. Comparison of state-of-the-art domain adaptation methods on SemanticKITTI→SemanticSTF adaptation for individual adverse weather conditions. We train a separate model for each weather-specific subset of SemanticSTF and evaluate the trained model on the weather condition it has been trained for.

3DSS Model	D-fog	L-fog	Rain	Snow	All
RandLA-Net [19]	26.5	26.0	25.1	22.7	25.3
SalsaNext [9]	16.0	9.6	7.8	3.5	9.1
SPVCNN [41]	30.4	22.8	21.7	18.3	22.4
SPVNAS [41]	25.5	18.3	17.0	13.0	18.0
Cylinder3D [62]	14.8	7.4	5.7	4.0	7.3

Table 6. Performance of state-of-the-art 3DSS models that are pre-trained over SemanticKITTI and tested on validation set of SemanticSTF for individual weather conditions and jointly for *all* weather conditions.

### 5.3. Network Models vs All-Weather 3DSS

We also study how different 3DSS network architectures generalize when they are trained with normal-weather point clouds and evaluated over SemanticSTF. Specifically, we select five representative 3DSS networks [9, 19, 41, 62] that have been widely adopted in 3D LiDAR segmentation studies. In the experiments, each selected network is first pre-trained with SemanticKITTI [2] and then evaluated over the validation set of SemanticSTF. We directly use the officially released code and the pre-trained weights for evaluation. Table 6 shows experimental results. We can observe that the five pre-trained models perform very differently though they all achieve superior segmentation over SemanticKITTI. Specifically, RandLA-Net [19], SPVCNN [41], and SPVNAS [41] perform clearly better than SalsaNext [9] and Cylinder3D [62]. In addition, none of the five pre-trained models perform well, verifying the clear domain discrepancy between point clouds of normal and adverse weather conditions. The experiments further indicate the great value of SemanticSTF in the future exploration of robust point cloud parsing under all weather conditions. In addition, the supervised performance of these 3DSS networks over SemanticSTF is provided in the appendix.

## 6. Conclusion and Outlook

This paper presents SemanticSTF, a large-scale dataset and benchmark suite for semantic segmentation of LiDAR

point clouds under adverse weather conditions. SemanticSTF provides high-quality point-level annotations for point clouds captured under adverse weather including dense fog, light fog, snow and rain. Extensive studies have been conducted to examine how state-of-the-art 3DSS methods perform over SemanticSTF, demonstrating its significance in directing future research on domain adaptive and domain generalizable 3DSS under all-weather conditions.

We also design PointDR, a domain randomization technique that aims to use normal-weather point clouds to train a domain generalizable 3DSS model that can work well over adverse-weather point clouds. PointDR consists of two novel designs including geometry style randomization and embedding aggregation which jointly learn perturbation-invariant representations that generalize well to various new point-cloud domains. Extensive experiments show that PointDR achieves superior point cloud segmentation performance as compared with the state-of-the-art.

## Acknowledgement

This study is funded BY the Ministry of Education Singapore, under the Tier-1 scheme with project number RG18/22. It is also supported under the RIE2020 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from Singapore Telecommunications Limited (Singtel), through Singtel Cognitive and Artificial Intelligence Lab for Enterprises (SCALE@NTU).



## References

- [1] Yogesh Balaji, Swami Sankaranarayanan, and Rama Chellappa. Metareg: Towards domain generalization using meta-regularization. *Advances in neural information processing systems*, 31, 2018. [2](#)
- [2] Jens Behley, Martin Garbade, Andres Milioto, Jan Quenzel, Sven Behnke, Cyrill Stachniss, and Jurgen Gall. Semantickitti: A dataset for semantic scene understanding of lidar sequences. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9297–9307, 2019. [1](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [3] Mario Bijelic, Tobias Gruber, Fahim Mannan, Florian Kraus, Werner Ritter, Klaus Dietmayer, and Felix Heide. Seeing through fog without seeing fog: Deep multimodal sensor fusion in unseen adverse weather. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11682–11692, 2020. [1](#), [2](#), [3](#), [4](#)
- [4] Gilles Blanchard, Gyemin Lee, and Clayton Scott. Generalizing from several related classification tasks to a new unlabeled sample. *Advances in neural information processing systems*, 24, 2011. [2](#)
- [5] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020. [4](#)
- [6] Ran Cheng, Ryan Razani, Ehsan Taghavi, Enxu Li, and Bingbing Liu. 2-s3net: Attentive feature fusion with adaptive feature selection for sparse semantic segmentation network. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12547–12556, 2021. [2](#)
- [7] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4d spatio-temporal convnets: Minkowski convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3075–3084, 2019. [2](#), [6](#), [7](#)
- [8] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016. [4](#)
- [9] Tiago Cortinhal, George Tzelepis, and Eren Erdal Aksoy. Salsanext: Fast, uncertainty-aware semantic segmentation of lidar point clouds. In *International Symposium on Visual Computing*, pages 207–222. Springer, 2020. [2](#), [8](#)
- [10] A Filgueira, H González-Jorge, Susana Lagüela, L Díaz-Vilariño, and Pedro Arias. Quantifying the influence of rain in lidar performance. *Measurement*, 95:143–148, 2017. [3](#)
- [11] Whye Kit Fong, Rohit Mohan, Juana Valeria Hurtado, Lubing Zhou, Holger Caesar, Oscar Beijbom, and Abhinav Valada. Panoptic nuscenes: A large-scale benchmark for lidar panoptic segmentation and tracking. *IEEE Robotics and Automation Letters*, 7(2):3795–3802, 2022. [1](#), [4](#)
- [12] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, pages 1180–1189. PMLR, 2015. [2](#)
- [13] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013. [4](#)
- [14] Dayan Guan, Jiaxing Huang, Aoran Xiao, and Shijian Lu. Domain adaptive video segmentation via temporal consistency regularization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8053–8064, 2021. [2](#)
- [15] Martin Hahner, Christos Sakaridis, Mario Bijelic, Felix Heide, Fisher Yu, Dengxin Dai, and Luc Van Gool. Lidar snowfall simulation for robust 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16364–16374, 2022. [2](#)
- [16] Martin Hahner, Christos Sakaridis, Dengxin Dai, and Luc Van Gool. 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*, pages 15283–15292, 2021. [2](#)
- [17] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020. [5](#)
- [18] Robin Heinzler, Philipp Schindler, Jürgen Seekircher, Werner Ritter, and Wilhelm Stork. Weather influence and classification with automotive lidar sensors. In *2019 IEEE intelligent vehicles symposium (IV)*, pages 1527–1534. IEEE, 2019. [3](#)
- [19] Qingyong Hu, Bo Yang, Linhai Xie, Stefano Rosa, Yulan Guo, Zhihua Wang, Niki Trigoni, and Andrew Markham. Randla-net: Efficient semantic segmentation of large-scale point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11108–11117, 2020. [1](#), [2](#), [8](#)
- [20] Jiaxing Huang, Dayan Guan, Aoran Xiao, and Shijian Lu. Cross-view regularization for domain adaptive panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10133–10144, 2021. [2](#)
- [21] Jiaxing Huang, Dayan Guan, Aoran Xiao, and Shijian Lu. Fsd: Frequency space domain randomization for domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6891–6902, 2021. [2](#)
- [22] Jiaxing Huang, Dayan Guan, Aoran Xiao, and Shijian Lu. Model adaptation: Historical contrastive learning for unsupervised domain adaptation without source data. *Advances in Neural Information Processing Systems*, 34:3635–3649, 2021. [2](#)
- [23] Jiaxing Huang, Dayan Guan, Aoran Xiao, Shijian Lu, and Ling Shao. Category contrast for unsupervised domain adaptation in visual tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1203–1214, 2022. [2](#)

- [24] Guoliang Kang, Lu Jiang, Yi Yang, and Alexander G Hauptmann. Contrastive adaptation network for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4893–4902, 2019. 2
- [25] Alexander Lehner, Stefano Gasperini, Alvaro Marcos-Ramiro, Michael Schmidt, Mohammad-Ali Nikouei Mahani, Nassir Navab, Benjamin Busam, and Federico Tombari. 3d-vfield: Adversarial augmentation of point clouds for domain generalization in 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17295–17304, 2022. 2
- [26] Hao Liang Li, Sinno Jialin Pan, Shiqi Wang, and Alex C Kot. Domain generalization with adversarial feature learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5400–5409, 2018. 2, 6
- [27] Zhijian Liu, Haotian Tang, Yujun Lin, and Song Han. Point-voxel cnn for efficient 3d deep learning. *Advances in Neural Information Processing Systems*, 32, 2019. 2
- [28] Zhipeng Luo, Zhongang Cai, Changqing Zhou, Gongjie Zhang, Haiyu Zhao, Shuai Yi, Shijian Lu, Hongsheng Li, Shanghang Zhang, and Ziwei Liu. Unsupervised domain adaptive 3d detection with multi-level consistency. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8866–8875, 2021. 2
- [29] Will Maddern, Geoffrey Pascoe, Chris Linegar, and Paul Newman. 1 year, 1000 km: The oxford robotcar dataset. *The International Journal of Robotics Research*, 36(1):3–15, 2017. 2
- [30] Andres Milioto, Ignacio Vizzo, Jens Behley, and Cyrill Stachniss. Rangenet++: Fast and accurate lidar semantic segmentation. In *2019 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pages 4213–4220. IEEE, 2019. 2
- [31] Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. Domain generalization via invariant feature representation. In *International Conference on Machine Learning*, pages 10–18. PMLR, 2013. 2
- [32] Thierry Peynot, James Underwood, and Steven Scheduling. Towards reliable perception for unmanned ground vehicles in challenging conditions. In *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1170–1176. IEEE, 2009. 3
- [33] Matthew Pitropov, Danson Evan Garcia, Jason Rebello, Michael Smart, Carlos Wang, Krzysztof Czarnecki, and Steven Waslander. Canadian adverse driving conditions dataset. *The International Journal of Robotics Research*, 40(4-5):681–690, 2021. 1, 2
- [34] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017. 2
- [35] Julian Ryde and Nick Hillier. Performance of laser and radar ranging devices in adverse environmental conditions. *Journal of Field Robotics*, 26(9):712–727, 2009. 3
- [36] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Acde: The adverse conditions dataset with correspondences for semantic driving scene understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10765–10775, 2021. 2, 3, 4
- [37] Cristiano Saltori, Fabio Galasso, Giuseppe Fiameni, Nicu Sebe, Elisa Ricci, and Fabio Poiesi. Cosmix: Compositional semantic mix for domain adaptation in 3d lidar segmentation. *ECCV*, 2022. 3, 7, 8
- [38] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014. 6
- [39] Peng Su, Kun Wang, Xingyu Zeng, Shixiang Tang, Dapeng Chen, Di Qiu, and Xiaogang Wang. Adapting object detectors with conditional domain normalization. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XI 16*, pages 403–419. Springer, 2020. 2
- [40] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2446–2454, 2020. 4
- [41] Haotian Tang, Zhijian Liu, Shengyu Zhao, Yujun Lin, Ji Lin, Hanrui Wang, and Song Han. Searching efficient 3d architectures with sparse point-voxel convolution. In *European conference on computer vision*, pages 685–702. Springer, 2020. 1, 2, 5, 6, 8
- [42] Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. In *2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pages 23–30. IEEE, 2017. 4
- [43] Jonathan Tremblay, Aayush Prakash, David Acuna, Mark Brophy, Varun Jampani, Cem Anil, Thang To, Eric Cameracci, Shaad Boochoon, and Stan Birchfield. Training deep networks with synthetic data: Bridging the reality gap by domain randomization. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 969–977, 2018. 4
- [44] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7167–7176, 2017. 7, 8
- [45] Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, and Patrick Pérez. Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2517–2526, 2019. 7, 8
- [46] Bichen Wu, Xuanyu Zhou, Sicheng Zhao, Xiangyu Yue, and Kurt Keutzer. Squeezesegv2: Improved model structure and unsupervised domain adaptation for road-object segmentation from a lidar point cloud. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 4376–4382. IEEE, 2019. 2, 3

- [47] Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3733–3742, 2018. [5](#)
- [48] Aoran Xiao, Jiaying Huang, Dayan Guan, Kaiwen Cui, Shijian Lu, and Ling Shao. Polarmix: A general data augmentation technique for lidar point clouds. *NeurIPS*, 2022. [3](#), [6](#), [7](#)
- [49] Aoran Xiao, Jiaying Huang, Dayan Guan, Fangneng Zhan, and Shijian Lu. Transfer learning from synthetic to real lidar point cloud for semantic segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 2795–2803, 2022. [1](#), [3](#), [4](#), [6](#), [7](#)
- [50] Aoran Xiao, Xiaofei Yang, Shijian Lu, Dayan Guan, and Jiaying Huang. Fps-net: A convolutional fusion network for large-scale lidar point cloud segmentation. *ISPRS Journal of Photogrammetry and Remote Sensing*, 176:237–249, 2021. [2](#)
- [51] Jianyun Xu, Ruixiang Zhang, Jian Dou, Yushi Zhu, Jie Sun, and Shiliang Pu. Rpvnet: A deep and efficient range-point-voxel fusion network for lidar point cloud segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16024–16033, 2021. [2](#)
- [52] Qiangeng Xu, Yin Zhou, Weiyue Wang, Charles R Qi, and Dragomir Anguelov. Spg: Unsupervised domain adaptation for 3d object detection via semantic point generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15446–15456, 2021. [2](#)
- [53] Jihan Yang, Shaoshuai Shi, Zhe Wang, Hongsheng Li, and Xiaojuan Qi. St3d: Self-training for unsupervised domain adaptation on 3d object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10368–10378, 2021. [2](#)
- [54] Xufeng Yao, Yang Bai, Xinyun Zhang, Yuechen Zhang, Qi Sun, Ran Chen, Ruiyu Li, and Bei Yu. Pcl: Proxy-based contrastive learning for domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7097–7107, 2022. [6](#)
- [55] Li Yi, Boqing Gong, and Thomas Funkhouser. Complete & label: A domain adaptation approach to semantic segmentation of lidar point clouds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15363–15373, 2021. [3](#)
- [56] Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, and Trevor Darrell. Bdd100k: A diverse driving dataset for heterogeneous multitask learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2636–2645, 2020. [2](#)
- [57] Feihu Zhang, Jin Fang, Benjamin Wah, and Philip Torr. Deep fusionnet for point cloud semantic segmentation. In *European Conference on Computer Vision*, pages 644–663. Springer, 2020. [2](#)
- [58] Weichen Zhang, Wen Li, and Dong Xu. Srdan: Scale-aware and range-aware domain adaptation network for cross-dataset 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6769–6779, 2021. [2](#)
- [59] Yang Zhang, Zixiang Zhou, Philip David, Xiangyu Yue, Zerong Xi, Boqing Gong, and Hassan Foroosh. Polarnet: An improved grid representation for online lidar point clouds semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9601–9610, 2020. [2](#)
- [60] Sicheng Zhao, Yezhen Wang, Bo Li, Bichen Wu, Yang Gao, Pengfei Xu, Trevor Darrell, and Kurt Keutzer. epointda: An end-to-end simulation-to-real domain adaptation framework for lidar point cloud segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 3500–3509, 2021. [2](#)
- [61] Kaiyang Zhou, Ziwei Liu, Yu Qiao, Tao Xiang, and Chen Change Loy. Domain generalization: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. [2](#)
- [62] Xinge Zhu, Hui Zhou, Tai Wang, Fangzhou Hong, Yuexin Ma, Wei Li, Hongsheng Li, and Dahua Lin. Cylindrical and asymmetrical 3d convolution networks for lidar segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9939–9948, 2021. [1](#), [2](#), [8](#)
- [63] Yang Zou, Zhiding Yu, Xiaofeng Liu, BVK Kumar, and Jinsong Wang. Confidence regularized self-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5982–5991, 2019. [7](#), [8](#)