

# Walking Your LiDOG: A Journey Through Multiple Domains for LiDAR Semantic Segmentation

Cristiano Saltori<sup>1</sup>, Aljoša Ošep<sup>2</sup>, Elisa Ricci<sup>1,3</sup>, Laura Leal-Taixé<sup>4</sup>

<sup>1</sup> University of Trento, <sup>2</sup> TU Munich, <sup>3</sup> Federazione Bruno Kessler, <sup>4</sup> NVIDIA

cristiano.saltori@unitn.it

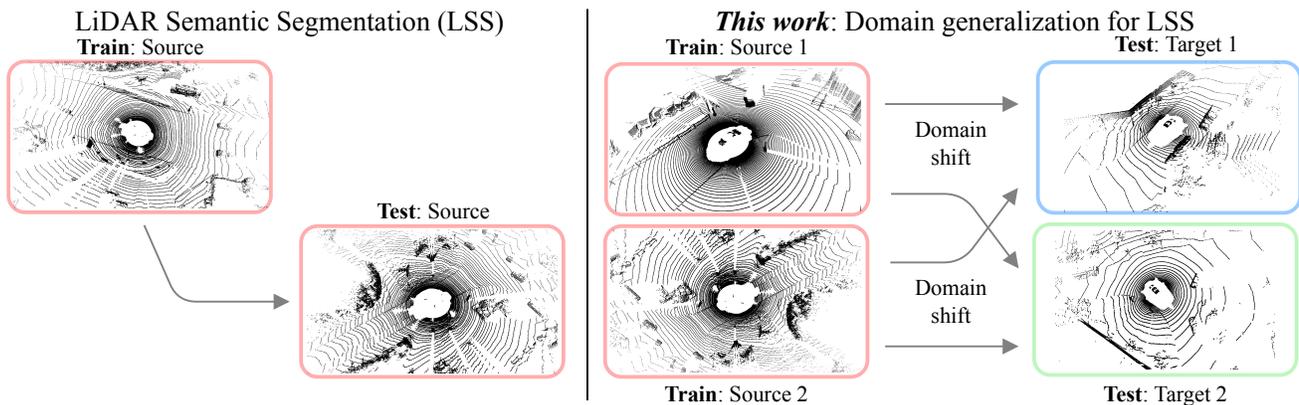


Figure 1: **Domain Generalization for LiDAR Semantic Segmentation (DG-LSS)**. *Left*: Existing LSS methods are trained and evaluated on point clouds drawn from the same domain. *Right*: We focus on studying LSS under domain shifts, where the test samples are drawn from a different data distribution. Our paper aims to address the generalization aspect of this task.

## Abstract

The ability to deploy robots that can operate safely in diverse environments is crucial for developing embodied intelligent agents. As a community, we have made tremendous progress in within-domain LiDAR semantic segmentation. However, do these methods generalize across domains? To answer this question, we design the first experimental setup for studying domain generalization (DG) for LiDAR semantic segmentation (DG-LSS). Our results confirm a significant gap between methods, evaluated in a cross-domain setting: for example, a model trained on the source dataset (SemanticKITTI) obtains 26.53 mIoU on the target data, compared to 48.49 mIoU obtained by the model trained on the target domain (nuScenes). To tackle this gap, we propose the first method specifically designed for DG-LSS, which obtains 34.88 mIoU on the target domain, outperforming all baselines. Our method augments a sparse-convolutional encoder-decoder 3D segmentation network with an additional, dense 2D convolutional decoder that learns to classify a birds-eye view of the point cloud. This simple auxiliary task encourages the 3D network to learn features that are robust to sensor placement shifts and resolution, and are transferable across domains. With this work, we aim to in-

spire the community to develop and evaluate future models in such cross-domain conditions.

## 1. Introduction

We address the challenge of achieving accurate and robust semantic segmentation of LiDAR point clouds. LiDAR semantic segmentation (LSS) is one of the most fundamental perception problems in mobile robot navigation, with applications ranging from mapping [28], localization [32], and online dynamic situational awareness [45].<sup>1</sup>

**Status Quo.** State-of-the-art LSS methods [8, 16, 44, 68] perform well when trained and evaluated using the same sensory setup and environment (*i.e.*, source domain, Fig. 1, left). However, their performance degrades significantly in the presence of domain shifts (Fig. 1, right), commonly caused by differences in sensory settings, *e.g.*, a new type of sensor, or recording environments, *e.g.*, geographic regions with different road layouts or types of vehicles. One way to mitigate this is to collect multi-domain datasets for pre-training,

<sup>1</sup>Semantic point classifier was part of the perception stack in early autonomous vehicles, such as Stanley, that won the DARPA challenge in 2005.

similar to what has been done in the image domain using inexpensive cameras that are widely available [9, 30]. However, building a crowd-sourced collection of multi-domain LiDAR datasets is currently not feasible.

**Stirring the pot.** As a first step towards LiDAR segmentation models that are robust to domain shifts, we present the first-of-its-kind experimental test-bed for studying *Domain Generalization* (DG) in the context of *LiDAR Semantic Segmentation* (LSS). In our DG-LSS setup we train and evaluate models on different domains, including two synthetic [39] and two real-world densely labeled datasets [3, 13], recorded in different geographic regions with different sensors. This evaluation setup reveals a significant gap in terms of mean intersection-over-union between models trained on the source and target domains: for example, a model transferred from SemanticKITTI to nuScenes dataset obtains 26.53 mIoU compared 48.49 mIoU obtained by the model fully trained on the target domain. *Could this gap be alleviated with DG techniques?*

**Insights.** To address this challenge, we propose *LiDOG* (**LiDAR D**omain **G**eneralization) as a simple yet effective method specifically designed for DG-LSS. In addition to reasoning about the scene semantics in 3D space, LiDOG projects features from the sparse 3D decoder onto the 2D bird’s-eye-view (BEV) plane along the vertical axis and learns to estimate a dense 2D semantic layout of the scene. In this way, LiDOG encourages the 3D network to learn features that are robust to variations in, *e.g.*, type of sensor or geo locations, and thereby can be transferred across different domains. This directly leads to increased robustness toward domain shifts and yields +8.35 mIoU improvement on the target domain, confirming the efficacy of our approach. Our experimental evaluation confirms this approach is consistently more effective compared to prior efforts in data augmentations [61, 34, 38], domain adaptation techniques [49, 22], and image-based DG techniques [7, 35], applied to the LiDAR semantic segmentation.

**Contributions.** We make the following key contributions. We (i) present the *first* study on *domain generalization* in the context of *LiDAR semantic segmentation*. To this end, we (ii) carefully construct a test-bed for studying DG-LSS using two synthetic and two densely-labeled real-world datasets recorded in different cities with different sensors. This allows us to (iii) rigorously study how prior efforts proposed in related domains can be used to tackle domain shift and (iv) propose LiDOG, a simple yet strong baseline that learns robust, generalizable, and domain-invariant features by learning semantic priors in the 2D birds-eye view. Despite its simplicity, we achieve state-of-the-art performance in all the generalization directions.<sup>2</sup>

<sup>2</sup>Our code is available at <https://saltoricristiano.github.io/lidog/>.

## 2. Related Work

**LiDAR Semantic Segmentation (LSS).** LiDAR-based semantic scene understanding has been used in autonomous vehicle perception since the dawn of autonomous driving [45]. State-of-the-art methods are data-driven and rely on different LiDAR data representations and network architectures. Point-based methods [36, 16, 44, 63, 52] learn representations directly from unordered point sets, whereas range-view-based methods [33, 50, 51, 62, 1] learn a representation from a range pseudo-image that represents the LiDAR point cloud. State-of-the-art methods [8, 15, 68, 27, 43] map 3D points into a 3D quantized grid and learn representations using sparse 3D convolutions [15, 8]. The works mentioned above focus on *within-domain* performance, *i.e.*, they train and test on data from the same dataset and do not evaluate robustness to domain shifts. By contrast, we focus on *cross-domain* robustness, *i.e.*, we train on a source domain and test on a target domain.

**Domain Adaptation for LSS.** Domain shift in LiDAR perception [49, 40, 23, 37] is a well-documented issue as sensors are expensive, and crowd-sourcing data collection is difficult. *Unsupervised* domain adaptation (UDA) techniques allow us to reduce the domain shift by looking at data from the target domain, but without needing any labels in that domain. Several methods reduce domain shift by mimicking sensory characteristics of the target domain data, often in a data-driven manner [22, 64, 54, 55]. Alternatively, domain shift can be reduced by feature alignment or prediction consistency [51, 20, 19], or by training a point cloud completion network to map domains into a canonical domain [58]. More recently, mixup strategies [60, 59] have been introduced in the context of UDA for LSS [38, 53]. In LiDAR-based 3D object detection, scaling [49] and self-training [40, 57] are often used strategies. All aforementioned methods must be exposed to (unlabeled) data from the target domain to adapt the learned representation. In this work, we study DG-LSS with the goal of learning as-general-and-robust-as-possible representations in the first place, before exposure to *any data from a particular target domain* during training.

**Domain Generalization.** Alternatively, robustness to domain shifts can be tackled via *Domain Generalization* (DG). In DG, we aim to learn domain-invariant features by training solely on the source domain data [66]. Due to the importance of learning robust representations, DG has been widely studied in the image domain [7, 66, 65, 35] in both, multi-[14, 26, 67] and single-source [5, 4, 65, 48, 24] settings. Image-based DG approaches reduce domain shift via domain alignment [29, 14], meta-learning [26, 2], data and style augmentations [42, 67, 65], ensembling techniques [11, 31], disentangled learning [21, 47], self-supervised learning [5, 4], regularization strategies [48, 17] and, reinforcement learning [17, 24]. Data augmentations have also been employed

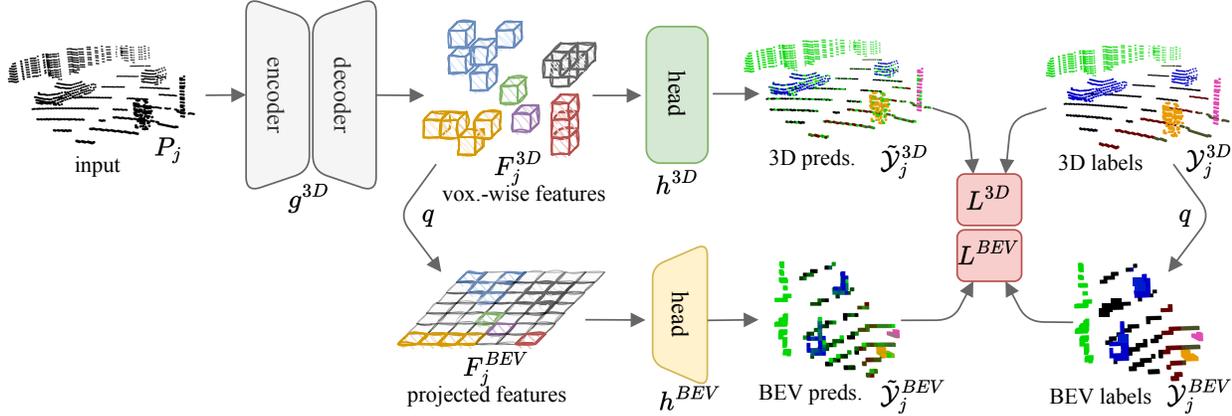


Figure 2: **LiDOG overview.** We encode our input LiDAR scan  $P_j$  using the 3D backbone  $g^{3D}$  to learn the occupied voxels’ feature representations  $F_j^{3D}$ . (*Upper branch - main task*) We apply a sparse segmentation head on  $F_j^{3D}$  and supervise with 3D semantic labels,  $\mathcal{Y}_j^{3D}$ . (*Lower branch - auxiliary task*) We project those features along the height-axis to obtain a dense 2D bird’s-eye (BEV) view features  $F_j^{BEV}$ , and apply several 2D convolutional layers to learn the 2D BEV representation. We supervise the BEV auxiliary task by using BEV-view of semantic labels,  $\mathcal{Y}_j^{BEV}$ . We train jointly on both  $L^{3D}$  and  $L^{BEV}$ .

in the context of DG in LiDAR-based 3D object detection in [25] to learn to detect vehicles that have been damaged in car accidents. In our work, we propose the first method specifically designed for DG-LSS and introduce a dense 2D convolutional decoder that learns features robust to domain shift. In our experimental evaluation (Sec. 5), we adapt and compare to several state-of-the-art data augmentation techniques, such as Mix3D [34], PointCutMix [61], and PolarMix [53] that apply mixup strategies [60, 59] to 3D point clouds. Moreover, we adapt state-of-the-art methods from the image domain RobustNet [7] and IBN [35]. We show that our method has substantially better generalization across all experiments. Finally, *parallel work* by [41] accumulates consecutive point clouds over time to introduce robustness towards point clouds of different domains and, consequentially, improve generalization capabilities.

### 3. A LiDAR Journey Across Domains

In this section, we first recap the standard LiDAR Semantic Segmentation (LSS) setting for completeness. Next, we formally introduce the problem of Domain Generalization for LiDAR Semantic Segmentation (DG-LSS), which we study in this paper.

**LiDAR Semantic Segmentation.** In LSS, we are given a set of labeled training instances in the form of point clouds  $\mathcal{T}_{train} = \{P_j = \{\mathbf{p}_i \in \mathbb{R}^3\}_{i=1}^{N_j}\}$ . Each point cloud  $P_j \sim \mathcal{D}$  consists of  $N_j$  points, and is drawn from a certain point cloud distribution  $\mathcal{D}$ . Moreover, point clouds in the training set are labeled according to a predefined semantic class vocabulary  $\mathcal{K} = \{1, \dots, K\}$  of  $K$  categorical labels. Given such a training set  $\mathcal{T}_{train}$  along with class vocabulary  $\mathcal{K}$ , the goal of LSS is to learn a function  $f(\mathbf{p}_i) \rightarrow \mathcal{K}$ , that maps each point  $\mathbf{p}_i \in P_j, \forall j$ , to one of  $K$  semantic classes, such

that the prediction error on the *same domain*  $\mathcal{D}$  is minimized. In other words, LSS models are trained *and* evaluated on (disjunct) sets of point clouds that are drawn from the same data distribution  $\mathcal{D}$ . We note that it is standard practice for LSS methods to be evaluated on multiple datasets<sup>3</sup>. However, standard multi-dataset evaluation follows the setting described above: models are trained *and* evaluated for *each* domain *separately*.

#### Domain Generalization for LiDAR Semantic Segmentation.

In this work, we release the assumption that train and test splits were sampled from the same domain and explicitly study *cross-domain generalization* in LSS. We are given  $M$  labeled datasets  $\mathcal{T}_m \sim \mathcal{D}_m, m \in \{1, \dots, M\}$  sampled from different data distributions  $\mathcal{D}_m$ . Importantly, all datasets must be labeled according to a common semantic vocabulary  $\mathcal{K}$ , *i.e.*, we are studying DG in closed-set conditions [66]. Then, the task can be defined as follows: we use  $\mathcal{T}_{\{m\}}$  datasets for training, and the held-out datasets  $\mathcal{T}_{\{m'\}}, \{m'\} \cap \{m\} = \emptyset$ , solely for testing.

The goal of DG-LSS is, therefore, to train a mapping  $f(\mathbf{p}_i) \rightarrow \mathcal{K}$ , such that the prediction performance on the target datasets  $\mathcal{T}_{\{m'\}} \sim \mathcal{D}_{\{m'\}}$  is maximized. We only have access to data and labels from the target datasets during evaluation. The model has never *seen any target data* before evaluation, neither labeled nor unlabeled.

### 4. Walking Your LiDOG

**Overview.** We provide a schematic overview of LiDOG: *Lidar D*omain *G*eneralization network in Fig. 2. LiDOG leverages a 3D sparse-convolutional encoder-decoder neural network (Fig. 2,  $g^{3D}$ ) that encodes the input point cloud  $P_j$

<sup>3</sup>At the moment, SemanticKITTI and nuScenes-lidarseg are commonly used for the evaluation.

into sparse 3D features  $F^{3D}$  (Fig. 2). In order to tackle the domain shift problem, we propose to augment this network with a dense top-down prediction auxiliary task during model training. In particular, we project sparse 3D features  $F^{3D}$  from the decoder network along the height axis to obtain a 2D bird’s-eye view (BEV) representation  $F^{BEV}$  of the features learned by the 3D network. Then, we jointly optimize the losses for the 3D and BEV network heads. Importantly, the auxiliary BEV head is not required during model inference. Its sole purpose is to effectively learn domain-agnostic features during training.

### 4.1. 3D Segmentation Branch

The upper branch of LiDOG (Fig. 2) aims to learn the main segmentation task, *i.e.*, it learns the mapping function  $f(\mathbf{p}_i) \rightarrow \mathcal{K}$ . We first process the input LiDAR point cloud  $P_j$  into a 3D occupancy (voxel) grid  $V_j$ . During this phase, the continuous 3D input space is evenly sampled into discrete 3D cells. Points falling into the same cell are merged. We then learn a feature representation of the occupancy grid  $V_j$  via a 3D encoder-decoder network  $g^{3D}$  based on the well-established sparse-convolutional backbone [8]. The encoder consists of a series of sparse 3D convolutional downsampling layers that learn a condensed 3D representation, while the decoder upsamples the features back to the original resolution with sparse convolutional upsampling layers. We use sparse batch-normalization layers. The output of  $g^{3D}$  is a set of voxel-wise features  $F_j^{3D} = g^{3D}(V_j)$ , where each feature vector is associated with a voxel cell in  $V_j$ . We employ a sparse segmentation head to obtain semantic posteriors that predict the voxel-wise categorical distribution from the voxel features  $F_j^{3D}$ :  $\tilde{\mathcal{Y}}_j^{3D} = p(\mathcal{K}|V_j) = \sigma(h^{3D}(F_j^{3D}))$ , where  $\sigma$  denotes the softmax activation function.

### 4.2. Dense Auxiliary Task

In the lower branch of LiDOG, we add the auxiliary task that encourages the network to learn a representation that is robust to domain shifts. We transform 3D features to bird’s-eye view (BEV) space, and train a network to estimate BEV semantic scene layout.

**Sparse-to-dense feature projection.** We start our dense-to-sparse segmentation task by projecting voxel-wise features  $F_j^{3D}$  into dense BEV-features feature space  $F_j^{BEV}$ . Given a voxel  $v_i \in V_j$  center coordinates  $(x_i, y_i, z_i)$ , we compute its corresponding BEV coordinates  $(x_i^{BEV}, z_i^{BEV})$  as:

$$q(x_i, z_i) = (x_i^{BEV}, z_i^{BEV}) = \left( \left\lfloor \frac{x_i}{x_q} \right\rfloor, \left\lfloor \frac{z_i}{z_q} \right\rfloor \right), \quad (1)$$

where  $q$  is the projection function, and  $x_q$  and  $z_q$  are the quantization parameters. We project only voxels within projection bounds  $B^{3D} = [(-b_x, b_x), (-b_z, b_z)]$ , where  $b_x$  and  $b_z$  are the bound values along the  $x$  and  $z$  axis, and

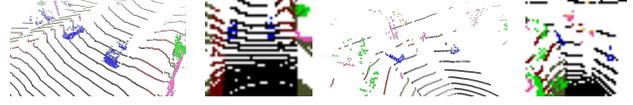


Figure 3: LiDAR point clouds and their corresponding BEV views: SemanticKITTI (left) and nuScenes (right). After projection, BEV images are geometrically more similar.

shift projected coordinates to fit in the BEV-features height and width. Consequently, we define BEV features  $F_j^{BEV}$  by initializing them to zero and by filling non-empty cells as  $F_j^{BEV}(x_i^{BEV}, z_i^{BEV}) = F_i^{3D}$ . When multiple voxels project to the same cell, we randomly keep one.

**Sparse-to-dense label projection.** We obtain source labels by following the same procedure used for  $F_j^{BEV}$ . After voxelization, each voxel  $v_i$  is associated to a label category  $y_i^{3D} \in \mathcal{Y}^{3D}$ . We use Eq. 1 and the same bounds  $B^{3D}$  to compute the label coordinates  $(x_i^{BEV}, z_i^{BEV})$  of BEV image pixels. Then, we define BEV labels as:  $\mathcal{Y}_j^{BEV}(x_i^{BEV}, z_i^{BEV}) = y_i^{3D}$ .

**BEV predictions.** We predict BEV semantic posteriors by employing a standard 2D segmentation head  $h^{BEV}$  which predicts semantic classes in the dense 2D BEV plane and takes as input BEV features  $F_j^{BEV}$ . First, we reduce the dimensionality of  $F_j^{BEV}$  through a pooling operation. Then, we feed  $h^{BEV}$  with  $F_j^{BEV}$  and obtain BEV semantic predictions:  $\tilde{\mathcal{Y}}_j^{BEV} = \sigma(h^{BEV}(F_j^{BEV}))$ , where  $\sigma$  denotes the softmax activation function.

### 4.3. Losses

Given 3D semantic predictions  $\tilde{\mathcal{Y}}_j^{3D}$  and labels  $\mathcal{Y}_j^{3D}$ , as well as BEV semantic predictions  $\tilde{\mathcal{Y}}_j^{BEV}$  and labels  $\mathcal{Y}_j^{BEV}$ , we train both LiDOG network heads jointly using soft DICE loss [18], *i.e.*,  $L^{3D} = dice(\tilde{\mathcal{Y}}_j^{3D}, \mathcal{Y}_j^{3D})$ , and  $L^{BEV} = dice(\tilde{\mathcal{Y}}_j^{BEV}, \mathcal{Y}_j^{BEV})$ . Finally, we average contributions of both losses:  $L_{tot} = \frac{1}{2}(L^{BEV} + L^{3D})$ . DICE loss has been shown in literature [18, 38] to perform favorably for rarer classes compared to the cross-entropy loss.

### 4.4. Why Does This Work?

We demonstrate empirically in Sec. 5 that the dense BEV prediction proxy task significantly improves model performance in terms of cross-domain generalization. Why does this simple objective yield, in practice, more robust feature representations? There are several possible sources of domain shift, *e.g.*, different sensors, different environments, and scans are visibly different (Fig. 3). After the projection, we obtain a denser label space that is less sensitive to individual sensor characteristics: in the BEV space scans from different domains look *more alike*, Fig. 3. We effectively regularize the model training by promoting the network robustness to changes in the sensor acquisition patterns. We

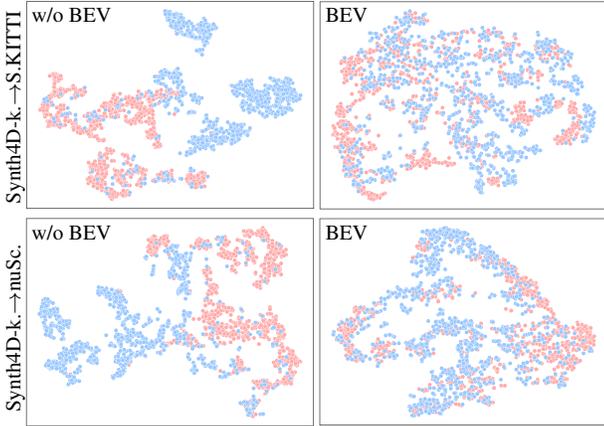


Figure 4: **Feature visualization:** t-SNE visualization [46] of the point embeddings for the *road* class, obtained by training our network without (w/o BEV, *left*) and with BEV task (BEV, *right*). *Top:* Synth4D-KITTI→SemanticKITTI. *Bottom:* Synth4D-KITTI→nuScenes. As shown, projected features are well-aligned when training the network with the auxiliary BEV task.

validate this intuition in Fig. 4 by visualizing learned point embeddings [46] for the *road* class with and without BEV auxiliary task, across multiple domains. As can be seen, projected embeddings are consistently closer when employing BEV auxiliary task.

## 5. Experimental Evaluation

In this section, we report our experimental evaluation. In Sec. 5.1, we describe our evaluation protocol, datasets and baselines. In Sec. 5.2-5.3, we report our *synth-to-real* and *real-to-real* results, and discuss qualitative results. In Sec. 5.4, we ablate our proposed approach. We provide implementation details in the appendix.

### 5.1. Experimental Protocol

We base our evaluation protocol on two synthetic (Synth4D [39]) and two real-world datasets (SemanticKITTI [3] and nuScenes [13]). We study (i) generalization from *synthetic*→*real data* in *single-* and *multi-source* settings, *i.e.*, trained on one or multiple source datasets, as well as (ii) *real*→*real* generalization.

**Synthetic datasets.** Synth4D [39] is a synthetic dataset generated using CARLA [10] simulator. It provides two synthetic source domains, each one with 20k labeled scans, emulating Velodyne64HDLE and Velodyne32HDLE sensors, similar to those used to record SemanticKITTI (*Synth4D-KITTI*) and nuScenes-lidarseg datasets (*Synth4D-nuScenes*). These datasets were generated in the same synthetic environment, mimicking different sensor types. We use the official training splits in all experiments.

**Real datasets.** SemanticKITTI [3] was recorded in vari-

ous regions of Karlsruhe, Germany, with Velodyne64HDLE sensor, and provides 20k labeled scans. nuScenes [13] was recorded in Boston, USA, and Singapore, with a sparser Velodyne32HDLE sensor and provides 40k labeled LiDAR scans. We use the official training and validation splits for both datasets in all experiments.

**Class vocabulary.** To ensure this task is well-defined, we formalize cross-dataset consistent and compatible semantic class vocabulary, which ensures there is a one-to-one mapping between all semantic classes. We follow [39] which provides a common label space among Synth4D, SemanticKITTI and nuScenes-lidarseg consisting of seven semantic classes: *vehicle*, *person*, *road*, *sidewalk*, *terrain*, *manmade* and *vegetation*. We perform all evaluation with respect to these classes.

**Baselines.** We are the first in studying DG-LSS, therefore, we construct our set of baselines by considering the previous efforts in bridging the domain shift in other LiDAR tasks as well DG methods for images.

*Data augmentation baselines (Aug).* Data augmentations are used in supervised learning to improve model generalization and, in DA, to reduce domain shift. We re-implement and evaluate in the DG context three data augmentation approaches: Mix3D [34], PointCutMix [61] and CoSMix [38]. Mix3D concatenates both points and labels of different scenes, obtaining a single mixed-up scene. PointCutMix and CoSMix follow the same strategy but mix either patches or semantic regions, respectively. We follow the official implementations and apply these data augmentation strategies during source training. For single-source training, we mix between samples of the same domain, while for multi-source training, we mix between samples of the two source domains.

*Image-based baselines (2D DG).* Image-based DG is a widely explored field. However, not all the methods can be extended to DG-LSS due to their image-related assumptions, *e.g.*, style consistency. We identify and re-implement two image-based approaches that can be extended to DG-LSS: IBN [35] and RobustNet [7]. IBN makes use of both batch and instance normalization layers in the network blocks to boost generalization performance. We follow the IBN-C block scheme and use it to build the sparse model. RobustNet [7] makes use of IBN blocks and introduces an instance whitening loss. Similarly, we start from our IBN implementation and use their instance-relaxed whitening loss [7] during training.

*UDA for LiDAR baselines (3D UDA).* Most of the UDA for LSS approaches assume (unlabeled) target data available for training/adaptation [57], hence, we cannot adapt and evaluate them in the DG setting fairly. We identify two UDA baselines that assume only weak supervision from the target

domain, *e.g.*, vehicle dimensions or sensor specs. We reimplement and extend SN [49] and RayCast [23]. SN [49] uses the knowledge of average vehicle dimensions from source and target domains to re-scale source instances. We employ DBSCAN to isolate vehicle instances in both source and target domains and estimate the ratio between their dimensions. At training time, we follow their pipeline and train our segmentation model on the re-scaled point clouds. RayCast [22] employs ray casting to re-sample source data by mimicking target sensor sampling. We use the official code to re-sample source data and use mapped data during training. Notice that both these baselines use a-priori knowledge from the target domain, *i.e.*, vehicle dimensions or target sensor specifications, and can thus be considered as weakly-supervised baselines.

## 5.2. Synth → Real Evaluation

In this section, we train our model on one or more synthetic source domains and evaluate on both real target domains, *i.e.*, SemanticKITTI and nuScenes-lidarseg. We report two models as lower and upper bounds in all the tables for reference: *source* (*i.e.*, a model trained on the source domain without the help of DG techniques) and *target* (*i.e.*, model directly trained on the target data).

**Single-source.** In Tab. 1, we report the results for single-source *Synth4D-KITTI→Real*. We report analogous *Synth4D-nuScenes→Real* results with similar findings in the appendix. We observe a 36.84 mIoU gap (SemanticKITTI) between the *source* (24.69 mIoU) and *target* (61.53 mIoU) models. Among the baselines, augmentation-based methods are the most effective: PointCutMix and Mix3D achieve 42.14 mIoU and 30.36 mIoU on *Synth4D-KITTI→Real*, respectively. Surprisingly, these approaches outperform 3D DA baselines, that leverage prior information about the target domain, *e.g.*, vehicle dimensions (SN) or target sensor specs (RayCast), confirming the efficacy of data augmentations for DG. LiDOG significantly reduces the domain gap between *source* and *target* models and outperforms all the compared baselines in all the scenarios. For example, on *Synth4D-KITTI→Real* (Tab. 1), we obtain 44.18 mIoU, a +19.49 improvement over the *source* model.

In Fig. 5, we discuss qualitative results, comparing the *source* model, the top-performing baseline, Mix3D, our proposed LiDOG, and the ground-truth labels. As we can see, the *source* model predictions are often incorrect and mingled in many small and large regions. Mix3D brings (limited) improvements. For example, notice how *road* (Fig. 5 top) is incorrectly segmented. LiDOG achieves overall top performance, improving on both small (*vehicle*) and large (*building* and *road*) areas. We report additional qualitative results in the appendix.

**Multi-source.** In Tab. 2, we report the results for the multi-

source training, where we train our model on both synthetic datasets, and evaluate the model on both real datasets. Compared to the single source (Tab. 1), in this setting we train models on synthetic data that mimics two different sensor types. The *source* model improves from 24.69 → 31.82 mIoU on SemanticKITTI and from 20.62 → 25.60 mIoU on nuScenes as a result of training on multiple synthetic datasets. As can be seen, LiDOG consistently improves over top-scoring baselines, RobustNet (+3.04 mIoU) and Mix3D (+4.56 mIoU). On average, LiDOG significantly improves over the *source* model with +10.62 mIoU and +14.63 mIoU on SemanticKITTI and nuScenes, respectively. We conclude that LiDOG is consistently a top-performer for classes and scenarios over the *source* model, except for the *terrain* class. This may be a limitation of our BEV projection. We assume it occurs when multiple classes are spatially overlapping due to the top-down projection, *e.g.*, mixing *terrain* with *vegetation*.

## 5.3. Real → Real Evaluation

In this setting, we only train and evaluate our models on real-world recordings. We report *SemanticKITTI→nuScenes* results in Tab. 3 and *nuScenes→SemanticKITTI* results in Tab. 4. Overall, we observe that the performance gap of *Real→Real* (Tab. 3-4) is lower compared to the gap in *Synth→Real* setting (Tab. 1-2 and Tab. 1 in the appendix).

**SemanticKITTI→nuScenes.** We observe a performance gap of 21.96 mIoU between *source* (26.53 mIoU) and *target* (48.49 mIoU) models. All the baselines consistently improve over the *source* model, however, by a smaller margin as compared to *Synth→Real*. Again, LiDOG improves on all the classes compared to the *source* performance and achieves top generalization performance on all the classes except for *vehicle* and *manmade*. On average, LiDOG obtains 34.88 mIoU and obtains +8.35 mIoU improvement over the *source* model.

**nuScenes→SemanticKITTI.** We observe a performance gap of 31.98 mIoU with *source* and *target* models achieving 29.55 mIoU and 61.53 mIoU, respectively. Augmentation-based methods are once again the most effective baselines, with PointCutMix achieving 39.78 mIoU. LiDOG outperforms all the compared baselines, achieves 41.22 mIoU and obtains +11.67 over the *source* model.

## 5.4. Ablations

**BEV auxiliary task.** We study whether LiDOG generalization capabilities are due to the additional segmentation head, or specifically due to the dense BEV segmentation head. The use of multiple prediction heads for ensembling is a well-known practice for improving prediction robustness [6, 12]. To this purpose, we remove the lower branch of LiDOG and attach an additional 3D branch. The resulting

Table 1: **Synth4D-KITTI→Real, single-source**. Our approach (LiDOG) improves upon *Source* model on both real datasets: +19.49 mIoU for SemanticKITTI and +16.52 mIoU for nuScenes, outperforming all baselines. Lower bound (*red*): a model trained on the source domain without the help of DG techniques. Upper bound (*blue*): model directly trained on the target data.

S	T	Info	Method	Vehicle	Person	Road	Sidewalk	Terrain	Manmade	Vegetation	mIoU
		Lower bound	Source	34.33	4.47	36.38	13.27	19.48	23.1	41.81	24.69
Synth4D-KITTI	SemanticKITTI	Aug	Mix3D [34]	59.99	14.02	55.75	17.71	25.67	39.26	53.00	37.92
			PointCutMix [61]	58.26	14.35	67.74	13.91	<b>27.26</b>	<b>49.21</b>	<b>64.22</b>	42.14
			CoSMix [38]	43.01	14.16	61.84	12.38	18.84	18.00	57.56	32.26
	2D DG	IBN [35]	24.95	9.18	56.82	17.85	7.21	21.59	44.65	26.04	
		RobustNet [7]	50.85	14.97	58.71	7.83	19.96	42.58	44.52	34.20	
	3D UDA	SN [49]	49.38	14.83	68.53	18.45	25.62	37.49	59.48	39.11	
		RayCast [23]	51.73	4.33	56.43	18.07	23.91	36.23	40.52	33.03	
	3D DG	Ours (LiDOG)	<b>72.86</b>	<b>17.1</b>	<b>71.85</b>	<b>28.48</b>	11.53	46.02	61.42	<b>44.18</b>	
	Upper bound	Target	81.57	23.23	82.09	57.64	46.84	64.14	75.21	61.53	
			Lower bound	Source	19.99	6.80	32.85	6.81	11.95	45.07	20.86
Synth4D-KITTI	nuScenes	Aug	Mix3D [34]	26.8	<b>22.68</b>	41.90	7.31	13.54	48.17	52.09	30.36
			PointCutMix [61]	23.97	19.49	44.27	6.29	12.14	48.85	55.44	30.06
			CoSMix [38]	16.94	15.78	52.97	2.09	5.55	15.6	43.85	21.83
	2D DG	IBN [35]	21.32	11.96	36.83	5.39	11.25	35.91	37.46	22.87	
		RobustNet [7]	26.19	7.47	43.85	2.29	13.93	43.63	46.32	26.24	
	3D UDA	SN [49]	23.14	14.08	51.60	11.38	14.02	46.91	50.83	30.28	
		RayCast [23]	19.54	11.78	56.53	6.66	8.45	45.66	36.99	26.52	
	3D DG	Ours (LiDOG)	<b>31.29</b>	19.62	<b>64.66</b>	<b>14.21</b>	<b>15.63</b>	<b>57.3</b>	<b>57.27</b>	<b>37.14</b>	
	Upper bound	Target	47.42	20.18	80.89	37.02	34.99	64.27	54.67	48.49	

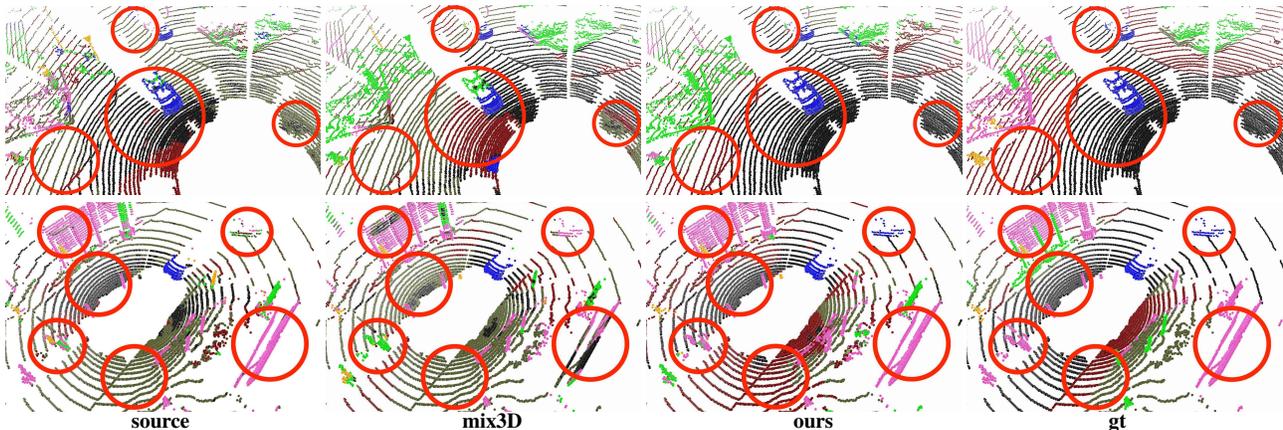


Figure 5: **Qualitative results**. *Top*: Synth4D-KITTI→SemanticKITTI, *bottom*: Synth4D-KITTI→nuScenes. LiDOG consistently improves results over the *source* model and outperforms Mix3D [34] with more homogeneous predictions.

architecture has two sparse 3D heads  $h^{3D}$  taking as input  $F^{3D}$ . During inference, predictions are obtained by averaging the predictions from both heads. Fig. 6 reports the generalization performance of this alternate version (*Double*) compared to LiDOG (*Ours*). As can be seen, by simply

learning two decoders and ensembling predictions we obtain lower performance as compared to utilizing the proposed BEV network heads to improve the generalization.

**BEV prediction area.** First, we study the impact on the generalization ability of LiDOG w.r.t. BEV area size in

Table 2: **(Synth4D-nuScenes + Synth4D-KITTI)→Real, multi-source**. Baselines significantly improve performance relative to the *source* model. Specifically, with LiDOG we observe +10.62 mIoU improvement on SemanticKITTI and +14.63 mIoU on nuScenes. Our approach (LiDOG) outperforms all the compared approaches. Lower bound (*red*): a model trained on the source domain without the help of DG techniques. Upper bound (*blue*): model directly trained on the target data.

S	T	Info	Method	Vehicle	Person	Road	Sidewalk	Terrain	Manmade	Vegetation	mIoU
Synth4D-nuScenes+Synth4D-KITTI	SemanticKITTI	Lower bound	Source	45.14	8.39	41.74	17.28	18.44	33.81	57.93	31.82
		Aug	Mix3D [34]	45.64	7.72	59.11	22.12	16.86	<b>50.03</b>	51.59	36.15
			PointCutMix [61]	48.38	6.28	60.8	20.53	7.56	33.62	54.32	33.07
			CoSMix [38]	45.81	0.93	59.28	22.81	18.13	43.14	40.91	33.00
		2D DG	IBN [35]	58.9	12.02	67.09	27.89	7.53	33.91	57.62	37.85
			RobustNet [7]	59.28	13.11	66.55	7.75	<b>30.87</b>	35.46	<b>62.75</b>	39.40
		3D UDA	SN [49]	35.07	5.95	60.36	<b>27.96</b>	13.3	26.88	54.88	32.06
			RayCast [23]	58.06	3.3	63.41	25.33	22.03	35.39	41.64	35.59
		3D DG	Ours (LiDOG)	<b>66.23</b>	<b>18.87</b>	<b>67.39</b>	24.13	15.22	46.21	59.03	<b>42.44</b>
		Upper bound	Target	81.57	23.23	82.09	57.64	46.84	64.14	75.21	61.53
Synth4D-nuScenes+Synth4D-KITTI	nuScenes	Lower bound	Source	20.64	13.4	28.44	7.61	10.23	52.88	46.01	25.60
		Aug	Mix3D [34]	28.35	24.6	57.16	13.96	9.48	60.19	55.98	35.67
			PointCutMix [61]	25.6	14.8	54.29	11.58	6.66	58.09	51.53	31.79
			CoSMix [38]	23.13	8.26	52.24	11.04	10.88	58.29	53.43	31.04
		2D DG	IBN [35]	26.11	16.95	55.16	13.59	12.67	56.64	52.92	33.43
			RobustNet [7]	27.03	16.48	50.61	6.66	15.73	57.68	<b>56.47</b>	32.95
		3D UDA	SN [49]	19.08	13.89	53.18	14.73	9.69	53.54	45.39	29.93
			RayCast [23]	25.37	9.03	60.27	10.03	11.15	54.72	45.56	30.87
		3D DG	Ours (LiDOG)	<b>30.78</b>	<b>25.3</b>	<b>73.2</b>	<b>19.83</b>	<b>16.03</b>	<b>62.65</b>	53.80	<b>40.23</b>
		Upper bound	Target	47.42	20.18	80.89	37.02	34.99	64.27	54.67	48.49

Table 3: **SemanticKITTI→nuScenes, single-source**. We train our model on SemanticKITTI and evaluate it on the nuScenes dataset. LiDOG improves over the *source* model by +8.35 mIoU. Lower bound (*red*): a model trained on the source domain with-out the help of DG techniques. Upper bound (*blue*): model directly trained on the target data.

S	T	Info	Method	Vehicle	Person	Road	Sidewalk	Terrain	Manmade	Vegetation	mIoU
SemanticKITTI	nuScenes	Lower bound	Source	22.92	0.03	63.33	16.09	7.42	35.40	40.53	26.53
		Aug	Mix3D [34]	33.74	11.22	58.54	12.91	5.28	<b>50.36</b>	48.59	31.52
			PointCutMix [61]	22.75	2.68	59.37	10.47	7.04	27.9	42.74	24.71
			CoSMix [38]	<b>35.91</b>	0.00	58.13	11.57	8.95	45.17	49.11	29.83
		2D DG	IBN [35]	29.93	0.01	56.77	18.70	12.09	37.67	33.83	27.00
			RobustNet [7]	25.50	0.03	62.40	15.12	9.40	29.58	43.62	26.52
		3D UDA	SN [49]	21.35	0.01	60.48	15.06	6.16	31.85	45.69	25.80
			RayCast [23]	28.82	0.00	59.25	16.08	12.51	49.72	49.82	30.89
		3D DG	Ours (LiDOG)	23.97	<b>14.86</b>	<b>70.63</b>	<b>24.59</b>	<b>13.97</b>	45.27	<b>50.85</b>	<b>34.88</b>
		Upper bound	Target	47.42	20.18	80.89	37.02	34.99	64.27	54.67	48.49

Table 4: **nuScenes**→**SemanticKITTI**, **single-source**. We train our model on nuScenes and evaluate it on the SemanticKITTI dataset. LiDOG improves over the *source* model by +11.67 mIoU. Lower bound (*red*): a model trained on the source domain with-out the help of DG techniques. Upper bound (*blue*): model directly trained on the target data.

S	T	Info	Method	Vehicle	Person	Road	Sidewalk	Terrain	Manmade	Vegetation	mIoU	
		Lower bound	Source	27.16	6.12	29.19	8.75	22.87	51.27	61.47	29.55	
nuScenes	SemanticKITTI	Aug	Mix3D [34]	37.86	6.74	41.95	5.73	27.59	41.21	65.41	32.36	
			PointCutMix [61]	55.22	<b>13.92</b>	45.96	5.43	30.47	<b>56.50</b>	<b>70.98</b>	39.78	
			CoSMix [38]	44.58	13.88	36.10	10.19	29.32	54.43	69.08	36.80	
	2D DG	IBN [35]	22.00	11.32	37.24	0.21	13.11	21.76	50.33	22.28		
		RobustNet [7]	32.94	10.98	39.85	14.70	28.27	50.42	58.47	33.66		
	3D UDA	SN [49]	25.69	5.46	19.59	2.17	23.47	27.65	61.07	23.58		
		RayCast [23]	28.30	16.09	45.80	9.44	20.56	38.56	61.83	31.51		
	3D DG	Ours (LiDOG)	<b>60.07</b>	9.03	<b>47.44</b>	<b>16.40</b>	<b>32.58</b>	54.21	68.82	<b>41.22</b>		
			Upper bound	Target	81.57	23.23	82.09	57.64	46.84	64.14	75.21	61.53

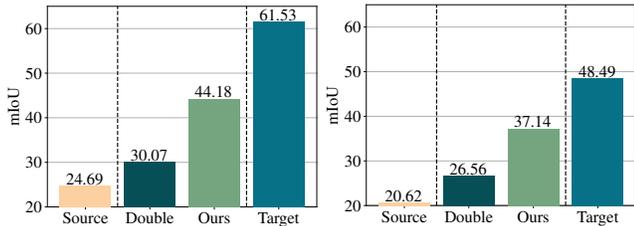


Figure 6: **Effectiveness of the BEV head:** We compare 2D BEV decoder (*Ours*) to simply adding an additional 3D segmentation head (*Double*) on SemanticKITTI (*left*) and nuScenes (*right*). Source: *Synth4D – KITTI*.

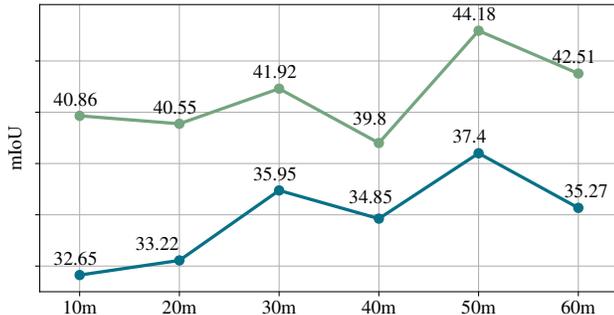


Figure 7: **BEV prediction area:** We study the impact of BEV area size on SemanticKITTI (*top*) and nuScenes (*bottom*). 50x50m area is consistently the best-performing option on both datasets. Source: *Synth4D – KITTI*.

Synth4D-KITTI→Real setting. We experiment with projection bounds  $B^{3D}$  ranging from 10x10 to 60x60m. As can be seen in Fig. 7, there is a near-linear relation between the BEV area and generalization performance.

## 6. Conclusions

**Impact.** With the advent of autonomous agents, there will be an increasing need for robust perception models that can

generalize across different environments and sensory configurations. Robustness to domain shift in semantic LiDAR scene understanding is therefore an important challenge that needs to be addressed. Fully supervised fine-tuning can be used to bridge domain gaps, however, it comes at the high cost of collecting and annotating data for specific domains and is therefore not a scalable solution.

**Domain generalization.** To the best of our knowledge, this is the first work that investigates domain generalization techniques to learn robust and generalizable representations for LiDAR semantic segmentation. We evaluate and discuss a comprehensive set of baselines from related fields and study their generalization across synthetic and real datasets, proving the need for more effective solutions. To this end, we propose LiDOG that encourages a 3D semantic backbone to learn features robust to domain shifts and is consistently more accurate than the compared baselines.

**Limitations.** The main limitation of LiDOG is that its core contribution, the auxiliary bird’s-eye prediction task, may lead to overlap after the projection to the BEV plane when classes overlap vertically, *e.g.*, the terrain below vegetation. This results in sub-optimal performance for these classes. This issue could be mitigated with the introduction of soft labels for overlapping semantic classes, or introducing additional losses on multiple BEV height slices (*c.f.*, [56]).

**Acknowledgments.** This project was partially funded by the Sofja Kovalevskaja Award of the Humboldt Foundation, the EU ISFP project PRE-CRISIS (ISFP-2022-TFI-AG-PROTECT-02-101100539), the PRIN project LEGO-AI (Prot. 2020TA3K9N) and the MUR PNRR project FAIR - Future AI Research (PE00000013) funded by the NextGenerationEU. It was carried out in the Vision and Learning joint laboratory of FBK-UNITN and used the CINECA, NVIDIA-AI TC clusters to run part of the experiments. We thank T. Meinhardt and I. Elezi for their feedback on this manuscript. The authors of this work take full responsibility for its content.

## References

- [1] I. Alonso, L. Riazuelo, L. Montesano, and A.C. Murillo. 3d-mininet: Learning a 2d representation from point clouds for fast and efficient 3d lidar semantic segmentation. In *IROS*, 2020. 2
- [2] Y. Balaji, S. Sankaranarayanan, and R. Chellappa. Metareg: Towards domain generalization using meta-regularization. *NeurIPS*, 2018. 2
- [3] Jens Behley, Martin Garbade, Andres Milioto, Jan Quenzel, Sven Behnke, Cyrill Stachniss, and Juergen Gall. SemanticKITTI: A Dataset for Semantic Scene Understanding of LiDAR Sequences. In *ICCV*, 2019. 2, 5
- [4] S. Bucci, A. D’Innocente, Y. Liao, F.M. Carlucci, B. Caputo, and T. Tommasi. Self-supervised learning across domains. *T-PAMI*, 2021. 2
- [5] F.M. Carlucci, A. D’Innocente, S. Bucci, B. Caputo, and T. Tommasi. Domain generalization by solving jigsaw puzzles. In *CVPR*, 2019. 2
- [6] M. Caron, I. Misra, J. Mairal, P. Goyal, P. Bojanowski, and A. Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *NeurIPS*, 2020. 6
- [7] S. Choi, S. Jung, H. Yun, J.T. Kim, S. Kim, and J. Choo. Robustnet: Improving domain generalization in urban-scene segmentation via instance selective whitening. In *CVPR*, 2021. 2, 3, 5, 7, 8, 9
- [8] C. Choy, J.Y. Gwak, and S. Savarese. 4D spatio-temporal convnets: Minkowski convolutional neural networks. In *CVPR*, 2019. 1, 2, 4
- [9] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A large-scale hierarchical image database. In *CVPR*, 2009. 2
- [10] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun. CARLA: An open urban driving simulator. In *ACRL*, 2017. 5
- [11] A. D’Innocente and B. Caputo. Domain generalization with domain-specific aggregation modules. In *GCPR*, 2019. 2
- [12] E. Fini, E. Sangineto, S. Lathuilière, Z. Zhong, M. Nabi, and E. Ricci. A unified objective for novel class discovery. In *ICCV*, 2021. 6
- [13] W. K. Fong, R. Mohan, J. V. Hurtado, L. Zhou, H. Caesar, O. Beijbom, and A. Valada. Panoptic nusenes: A large-scale benchmark for lidar panoptic segmentation and tracking. *RAL*, 2021. 2, 5
- [14] M. Ghifary, D. Balduzzi, B.W. Kleijn, and M. Zhang. Scatter component analysis: A unified framework for domain adaptation and domain generalization. *T-PAMI*, 2016. 2
- [15] B. Graham, M. Engelcke, and L. van der Maaten. 3d semantic segmentation with submanifold sparse convolutional networks. In *CVPR*, 2018. 2
- [16] Q. Hu, B. Yang, L. Xie, S. Rosa, Y. Guo, Z. Wang, N. Trigoni, and A. Markham. RandLA-Net: Efficient semantic segmentation of large-scale point clouds. In *CVPR*, 2020. 1, 2
- [17] Z. Huang, H. Wang, E. P. Xing, and D. Huang. Self-challenging improves cross-domain generalization. In *ECCV*, 2020. 2
- [18] S. Jadon. A survey of loss functions for semantic segmentation. In *CIBCB*, 2020. 4
- [19] M. Jaritz, T.H. Vu, R. de Charette, E. Wirbel, and P. Pérez. xmuda: Cross-modal unsupervised domain adaptation for 3d semantic segmentation. In *CVPR*, 2020. 2
- [20] M. Jaritz, T.-H. Vu, R. de Charette, E. Wirbel, and P. Pérez. xMUDA: Cross-modal unsupervised domain adaptation for 3d semantic segmentation. In *CVPR*, 2020. 2
- [21] A. Khosla, T. Zhou, T. Malisiewicz, A. Efros A, and A. Torralba. Undoing the damage of dataset bias. In *ECCV*, 2012. 2
- [22] F. Langer, A. Milioto, A. Haag, J. Behley, and C. Stachniss. Domain Transfer for Semantic Segmentation of LiDAR Data using Deep Neural Networks. In *IROS*, 2020. 2, 6
- [23] F. Langer, A. Milioto, A. Haag, J. Behley, and C. Stachniss. Domain transfer for semantic segmentation of LiDAR data using deep neural networks. In *IROS*, 2021. 2, 6, 7, 8, 9
- [24] M. Laskin, A. Srinivas, and P. Abbeel. Curl: Contrastive unsupervised representations for reinforcement learning. In *ICML*, 2020. 2
- [25] A. Lehner, S. Gasperini, A. Marcos-Ramiro, M. Schmidt, M.N. Mahani, N. Navab, B. Busam, and F. Tombari. 3d-vfield: Adversarial augmentation of point clouds for domain generalization in 3d object detection. In *CVPR*, 2022. 3
- [26] D. Li, Y. Yang, and Y.-Z. Song T. Hospedales. Learning to generalize: Meta-learning for domain generalization. In *AAAI*, 2018. 2
- [27] J. Li, X. He, Y. Wen, Y. Gao, Y. Cheng, and D. Zhang. Panoptic-phnet: Towards real-time and high-precision lidar panoptic segmentation via clustering pseudo heatmap. In *CVPR*, 2022. 2
- [28] Q. Li, Y. Wang, Y. Wang, and H. Zhao. Hdmapnet: An online hd map construction and evaluation framework. In *ICRA*, 2022. 1
- [29] Y. Li, X. Tian, M. Gong, Y. Liu, T. Liu, K. Zhang, and D. Tao. Deep domain generalization via conditional invariant adversarial networks. In *ECCV*, 2018. 2
- [30] T.Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. Lawrence Zitnick. Microsoft COCO: Common objects in context. In *ECCV*, 2014. 2
- [31] Q. Liu, Q. Dou, L. Yu, and A.P. Heng. Ms-net: multi-site network for improving prostate segmentation with heterogeneous mri data. *TMI*, 2020. 2
- [32] W.C. Ma, I. Tartavull, I. A. Bârsan, S. Wang, M. Bai, G. Matyus, N. Homayounfar, S. K. Lakshminanth, A. Pokrovsky, and R. Urtasun. Exploiting sparse semantic hd maps for self-driving vehicle localization. In *IROS*, 2019. 1
- [33] A. Milioto, I. Vizzo, J. Behley, and C. Stachniss. Rangenet++: Fast and accurate lidar semantic segmentation. In *IROS*, 2019. 2
- [34] A. Nekrasov, J. Schult, O. Litany, B. Leibe, and F. Engelmann. Mix3D: Out-of-Context Data Augmentation for 3D Scenes. In *3DV*, 2021. 2, 3, 5, 7, 8, 9
- [35] X. Pan, P. Luo, J. Shi, and X. Tang. Two at once: Enhancing learning and generalization capacities via ibn-net. In *ECCV*, 2018. 2, 3, 5, 7, 8, 9
- [36] C.R. Qi, L. Yi, H. Su, and L.J. Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *arXiv*, 2017. 2

- [37] C. Qin, H. You, L. Wang, C.J. Kuo, and Y. Fu. Pointdan: A multi-scale 3d domain adaption network for point cloud representation. In *NeurIPS*, 2019. 2
- [38] C. Saltori, Fabio Galasso, G. Fiameni, N. Sebe, E. Ricci, and F. Poiesi. Cosmix: Compositional semantic mix for domain adaptation in 3d lidar segmentation. In *ECCV*, 2022. 2, 4, 5, 7, 8, 9
- [39] C. Saltori, E. Krivosheev, S. Lathuilière, N. Sebe, F. Galasso, G. Fiameni, E. Ricci, and F. Poiesi. Gipso: Geometrically informed propagation for online adaptation in 3d lidar segmentation. In *ECCV*, 2022. 2, 5
- [40] C. Saltori, S. Lathuilière, N. Sebe, E. Ricci, and F. Galasso. Sf-uda3d: Source-free unsupervised domain adaptation for lidar-based 3d object detection. In *3DV*, 2020. 2
- [41] J. Sanchez, J.E. Deschaut, and F. Goulette. Domain generalization of 3d semantic segmentation in autonomous driving. *arXiv*, 2022. 3
- [42] S. Shankar, V. Piratla, S. Chakrabarti, S. Chaudhuri, and P. Jyothi. Generalizing across domains via cross-gradient training. *ICLR*, 2018. 2
- [43] H. Tang, Z. Liu, S. Zhao, Y. Lin, J. Lin, H. Wang, and S. Han. Searching efficient 3d architectures with sparse point-voxel convolution. In *ECCV*, 2020. 2
- [44] H. Thomas, C.R. Qi, J.-E. Deschaut, B. Marcotegui, F. Goulette, and J. L. Guibas. Kpconv: Flexible and deformable convolution for point clouds. In *ICCV*, 2019. 1, 2
- [45] S. Thrun, M. Montemerlo, H. Dahlkamp, D. Stavens, A. Aron, J. Diebel, P. Fong, J. Gale, M. Halpenny, and G. Hoffmann. Stanley: The robot that won the darpa grand challenge. *Journal of field Robotics*, 2006. 1, 2
- [46] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *J. Mach. Learn.*, 9:85, 2008. 5
- [47] G. Wang, H. Han, S. Shan, and X. Chen. Cross-domain face presentation attack detection via multi-domain disentangled representation learning. In *CVPR*, 2020. 2
- [48] H. Wang, Z. He, Z.C. Lipton, and E.P. Xing. Learning robust representations by projecting superficial statistics out. *ICLR*, 2019. 2
- [49] Y. Wang, X. Chen, Y. You, L.E. Li, B. Hariharan, M. Campbell, K.Q. Weinberger, and W.L. Chao. Train in germany, test in the usa: Making 3d object detectors generalize. In *CVPR*, 2020. 2, 6, 7, 8, 9
- [50] B. Wu, A. Wan, X. Yue, and K. Keutzer. SqueezeSeg: Convolutional neural nets with recurrent crf for real-time road-object segmentation from 3d lidar point cloud. In *ICRA*, 2018. 2
- [51] B. Wu, X. Zhou, S. Zhao, X. Yue, and K. Keutzer. SqueezeSegv2: Improved model structure and unsupervised domain adaptation for road-object segmentation from a lidar point cloud. In *ICRA*, 2019. 2
- [52] X. Wu, Y. Lao, L. Jiang, X. Liu, and H. Zhao. Point transformer v2: Grouped vector attention and partition-based pooling. In *NeurIPS*, 2022. 2
- [53] A. Xiao, J. Huang, D. Guan, K. Cui, S. Lu, and L. Shao. Polarmix: A general data augmentation technique for lidar point clouds. In *NeurIPS*, 2022. 2, 3
- [54] A. Xiao, J. Huang, D. Guan, F. Zhan, and S. Lu. Synlidar: Learning from synthetic lidar sequential point cloud for semantic segmentation. In *AAAI*, 2022. 2
- [55] A. Xiao, J. Huang, D. Guan, F. Zhan, and S. Lu. Transfer learning from synthetic to real lidar point cloud for semantic segmentation. In *AAAI*, 2022. 2
- [56] Bin Yang, Wenjie Luo, and Raquel Urtasun. Pixor: Real-time 3d object detection from point clouds. In *CVPR*, 2018. 9
- [57] J. Yang, S. Shi, Z. Wang, H. Li, and X. Qi. St3d: Self-training for unsupervised domain adaptation on 3d object detection. In *CVPR*, 2021. 2, 5
- [58] L. Yi, B. Gong, and T. Funkhouser. Complete & label: A domain adaptation approach to semantic segmentation of lidar point clouds. In *CVPR*, 2021. 2
- [59] S. Yun, D. Han, S.J. Oh, S. Chun, J. Choe, and Y. Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *ICCV*, 2019. 2, 3
- [60] H. Zhang, M. Cisse, Y.N. Dauphin, and D. Lopez-Paz. mixup: Beyond empirical risk minimization. In *ICLR*, 2018. 2, 3
- [61] J. Zhang, L. Chen, B. Ouyang, B. Liu, J. Zhu, Y. Chen, Y. Meng, and D. Wu. Pointcutmix: Regularization strategy for point cloud classification. *Neurocomputing*, 2022. 2, 3, 5, 7, 8, 9
- [62] Y. Zhang, Z. Zhou, P. David, X. Yue, Z. Xi, B. Gong, and H. Foroosh. Polarnet: An improved grid representation for online lidar point clouds semantic segmentation. In *CVPR*, 2020. 2
- [63] H. Zhao, L. Jiang, J. Jia, P. Torr, and V. Koltun. Point transformer. In *ICCV*, 2021. 2
- [64] S. Zhao, Y. Wang, B. Li, B. Wu, Y. Gao, P. Xu, T. Darrell, and K. Keutzer. epointda: An end-to-end simulation-to-real domain adaptation framework for lidar point cloud segmentation. In *AAAI*, 2021. 2
- [65] Y. Zhao, Z. Zhong, N. Zhao, N. Sebe, and G.H. Lee. Style-hallucinated dual consistency learning for domain generalized semantic segmentation. In *ECCV*, 2022. 2
- [66] K. Zhou, Z. Liu, Y. Qiao, T. Xiang, and C.C. Loy. Domain generalization in vision: A survey. *T-PAMI*, 2021. 2, 3
- [67] K. Zhou, Y. Yang, T. Hospedales, and T. Xiang. Deep domain-adversarial image generation for domain generalisation. In *AAAI*, 2020. 2
- [68] X. Zhu, H. Zhou, T. Wang, F. Hong, Y. Ma, W. Li, H. Li, and D. Lin. Cylindrical and asymmetrical 3d convolution networks for lidar segmentation. In *CVPR*, 2021. 1, 2