

# Curricular Object Manipulation in LiDAR-based Object Detection

Ziyue Zhu<sup>1\*</sup> Qiang Meng<sup>2\*</sup> Xiao Wang<sup>2</sup> Ke Wang<sup>2</sup> Liujiang Yan<sup>2</sup> Jian Yang<sup>1†</sup>

<sup>1</sup>Tianjin Key Laboratory of Visual Computing and Intelligent Perception,  
 College of Computer Science, Nankai University <sup>2</sup>Didi Chuxing

zhuziyue@mail.nankai.edu.cn irvingmeng@didiglobal.com

## Abstract

*This paper explores the potential of curriculum learning in LiDAR-based 3D object detection by proposing a curricular object manipulation (COM) framework. The framework embeds the curricular training strategy into both the loss design and the augmentation process. For the loss design, we propose the COMLoss to dynamically predict object-level difficulties and emphasize objects of different difficulties based on training stages. On top of the widely-used augmentation technique called GT-Aug in LiDAR detection tasks, we propose a novel COMAug strategy which first clusters objects in ground-truth database based on well-designed heuristics. Group-level difficulties rather than individual ones are then predicted and updated during training for stable results. Model performance and generalization capabilities can be improved by sampling and augmenting progressively more difficult objects into the training samples. Extensive experiments and ablation studies reveal the superior and generality of the proposed framework. The code is available at <https://github.com/ZZY816/COM>.*

## 1. Introduction

LiDAR sensors can provide accurate, high-definition 3D measurements of the surrounding environment. Such 3D information plays a noninterchangeable role in safety-critical applications like 3D object detection in self-driving. However, the rich 3D information from LiDAR sensors does not come without problems. Usually presented in the form of a point cloud, LiDAR data suffers from (i) non-uniformity: the point density decreases monotonically as the laser range increases; (ii) orderless: the geometry of a point cloud remains unchanged even if all of its points are randomly shuffled. (iii) sparsity: when quantized into voxel grids, a significant portion of the voxels are empty;

To build a robust and performant LiDAR object detector,

\*The first two authors have equal contribution to this work.

†Jian Yang is the corresponding author (csjyang@nankai.edu.cn).

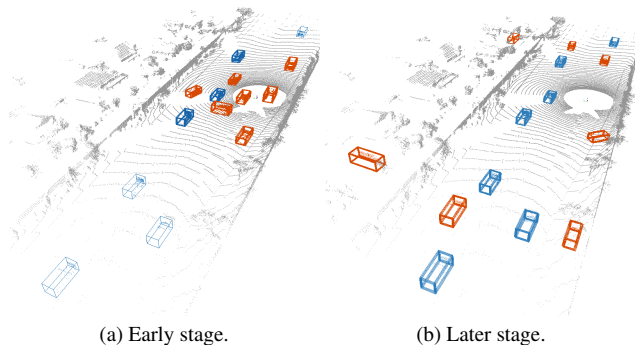


Figure 1. The proposed Curricular Object Manipulation (COM) works in an easy-to-hard manner. In early stages, COMAug constrains the augmented objects (highlighted in red) to be easy ones and COMLoss down-weights losses from difficult objects (marked in boxes with thin lines). Objects with varying degrees of difficulty are inserted into the point clouds in later stages. On the other hand, hard objects will contribute more to loss values as training progresses. **Best viewed in color.**

different data representations have been explored to alleviate the non-uniformity and orderless challenges. Feature extraction from the raw orderless point cloud can be made possible by performing radius search or nearest neighbor search in the 3D Euclidean space [6, 32, 41, 56]. Another popular solution is to quantize the input point cloud into a fixed grid of voxels [61] or pillars of voxels [21]. At the price of quantization error, later processing can be done efficiently on the regular voxel pillars or grids [52].

But these different data representations do not change the sparsity of the LiDAR point cloud data. Compared with image object detection tasks, sparse point clouds contain much less input stimuli and positive samples for neural network training, as depicted in Figure 1. Thus, effective data augmentation strategies are critical for faster model convergence and better detection performance [13, 18, 31, 50, 52]. Among them, GT-Aug [52] (see Figure 2) is widely adopted. GT-Aug first aggregates ground truth labels from the training dataset into a database. During training, randomly selected samples from the database are inserted into




Figure 2. 3D object detection mAP of the car category with hard difficulties on the KITTI dataset from 2018 to 2021. It is obvious from the figure that the GT-Aug strategy boosts the KITTI 3D object detection benchmark by a large margin since its inception [52]. GT-Aug has since become the de facto augmentation practice in popular open source toolkits [7, 48].

the point cloud to amplify the supervision signal.

Notice that GT-Aug treats all samples in the database equally, and all epochs of the training process equally. It has brought to our attention that selecting too many hard examples at early stages may overwhelm the training, while selecting too many easy samples at the later stages may slow the model convergence. Similar conclusions were also reached independently in the facial recognition field [16]. This important finding raises two questions for the widely used GT-Aug strategy: (i) at a given training stage, how to select samples that benefit the current stage the most, (ii) at different training stages, how to adjust the sampling strategies accordingly. However, solving these two questions is not yet enough as the original objects in the training sample can also be ill-suited for current training. Therefore, we raise one additional question as (iii) how to properly handle augmented objects as well as original objects can contribute to the model performance.

This work answers the above questions by leveraging curriculum learning. Curriculum learning draws inspiration from the human cognitive process, which begins with easier concepts and gradually moves on to more complicated ones [1, 44]. Enlightened by such easy-to-hard paradigm, we propose a curricular object manipulation (COM) framework for the LiDAR object detection task. Our framework consists of (i) **COMLoss** to manipulate the contributions from objects of different difficulties, and (ii) **COMAug** to manipulate the sampling process in GT-Aug.

In the COM framework, we employ the classification loss as a simple yet effective proxy for object difficulties. The COMLoss suppresses loss contributions from hard objects in earlier stages and gradually loses the suppression, as depicted in Fig. 1. Unfortunately, using classification score as the difficulty proxy can cause an inevitable paradox in COMAug. Specifically, COMAug relies on up-to-date scores of all objects to perform difficulty-adaptive augmentation. In contrast, all objects should be sampled recently for augmentation to update their scores, which is impossible because of the limited number of augmented objects in each training frame. We design a clustering based method to address such paradox: objects with similar diffi-

culties are grouped together, and the difficulty estimates are updated for the groups rather than for the individual objects. During training, hard groups will be sampled with monotonically increasing probabilities as epoch increases, while objects within each group will be sampled uniformly. In our work, objects are grouped by their geometry attributes, such as distance, dimension, angle, and occupancy ratio.

We demonstrate the efficacy of our proposed method through extensive experiments and ablation studies. In summary, our contributions include:

- We propose the COM framework which embeds the easy-to-hard training strategy into both loss design and augmentation process in LiDAR-based object detection. For the loss design, COMLoss is introduced to dynamically predict object-level difficulties, based on which we emphasize objects to different extents when the training proceeds. For the augmentation, a well-designed COMAug first clusters objects in ground-truth database with carefully-picked heuristics. During training, COMAug updates group-level difficulties and controls sampling process in augmentation in a curricular manner.
- To the best of our knowledge, COM is the first to explore the potentials of curriculum learning in conventional LiDAR-based 3D object detection task. Extensive experiments and ablation studies reveal the superiority and generality of the proposed framework.

## 2. Related Works

### 2.1. LiDAR-based 3D Object Detection

LiDAR based object detection aims at localizing objects of interest from the input point cloud. Current works in this area can be roughly classified based on their LiDAR data representation. Range view based solutions [2, 4, 12, 30, 46] have high computation efficiency due to the compactness of the 2D range view representation, but usually have inferior detection performance caused by the 2D-3D dimensional gaps. By directly extracting features from raw point clouds, point-based detectors [6, 24, 32, 33, 39, 40, 41, 55, 56, 59, 60] achieve satisfying performances but commonly suffer from high computational costs incurred by radius search or nearest neighbor queries in the 3D Euclidean space. In contrast, voxel-based detectors first transform non-uniform point clouds into regular 2D pillars or 3D voxels and employ convolutions for efficient processing in later stages. Pioneering works, including VoxelNet [61] and PointPillars [21], demonstrate great efficiency-utility trade-off, thus attract much attention from the community [9, 10, 11, 52, 58]. Without loss of generality, we primarily focus on voxel-based methods in our experiments.

## 2.2. Data Augmentation in Point Clouds

Constrained by the enormous annotation cost, public LiDAR datasets usually come in with much smaller volume compared with image datasets, *e.g.*, 15K frames in KITTI [14] compared with 328K images in MS-COCO [26]. Thus, effective data augmentation strategies are critical for the performance and the generalization capabilities of LiDAR object detection models.

In addition to simple geometry deformations, such as *random rotation*, *random flip*, and *translation*, LiDAR tasks usually employ ground-truth augmentation [52] to mitigate the sparsity issue in point clouds. Before training, ground truth objects with their corresponding point clouds are first collected into a database. During training, additional ground truth objects sampled from the database are concatenated into current training point to supplement the supervision. This strategy, termed as GT-Aug, has been widely used in current literature [37, 38, 52, 53, 55, 58] and popular open source toolkits [7, 48]. Thus, our experiments are concentrated on the most effective GT-Aug strategy.

## 2.3. Curriculum Learning

Curriculum learning improves model performance and generalization by progressively drawing harder data samples for training [1, 3, 5, 15, 17, 19, 23, 34, 44]. In spite of the fact that curriculum learning has demonstrated its effectiveness in certain classification tasks, its application in object detection tasks [3, 36], especially in LiDAR 3D detection [54], remains largely unexplored. In Saxena et al. [36], each object is assigned with a difficulty value that is updated by gradient during training. Superloss [3] proposes to find close-form solutions on the fly and down-weight contributions of samples with high variance. In our work, we simply employ the loss value as the indicator of difficulty. Without loss of generality, the module can be directly replaced by those in works[3, 36] for better performance.

In addition to training recipe, curriculum learning is also proved to be successful for data augmentation [27, 35, 47, 49]. Following the same high-level idea of presenting synthetic data with increasing difficulties, we aim at designing a curricular scheme for the GT-Aug in point clouds.

## 3. Methodology

In this section, we present details of the curriculum object manipulation (COM) framework, as illustrated in Fig. 3. Basic concepts and notations in LiDAR-based object detection are first reviewed in Sec. 3.1. COMLoss, which manipulates the losses from objects of different difficulties, is described in Sec. 3.2. COMAug, which manipulates the object sampling process in GT-Aug, is described in Sec. 3.3.

## 3.1. Notations and Definitions

The GT-Aug strategy firstly collects a database  $\mathcal{B} = \{(x_i, y_i)\}_1^n$  of ground-truth objects. For each ground-truth object, 3D coordinates and additional features (such as intensity, elongation, *etc.*) for 3D points belonging to the object are collected as  $x_i$ , while human annotations (such as bounding box center position, dimension, yaw angle, category, *etc.*) are collected as  $y_i$ . At any given training sample, assuming there are  $P$  annotated ground-truth objects in the current LiDAR point cloud, the GT-Aug strategy works by complementing the target objects up to a predefined threshold  $\Gamma$  by sampling  $Q = \min(0, \Gamma - P)$  objects from the pre-collected database  $\mathcal{B}$ .

Given the  $P$  original objects and the  $Q$  augmented objects, the optimization objective for LiDAR based 3D object detectors can be summarized into the general form below:

$$\mathcal{L} = \frac{1}{N} \left( \mathcal{L}_n + \sum_{p=1}^P (\mathcal{L}_c^p + \mathcal{L}_r^p) + \sum_{q=1}^Q (\mathcal{L}_c^q + \mathcal{L}_r^q) \right) \quad (1)$$

For original objects  $p$  and augmented objects  $q$ , their classification loss  $\mathcal{L}_c$  and regression loss  $\mathcal{L}_r$  are both taken into consideration during training.  $\mathcal{L}_n$  is introduced to handle the point cloud background. The normalization factor  $N$ , which depends on the number of positive labels, is usually customized by different detection models.

## 3.2. COMLoss

The first component of the COM framework is the COM-Loss aiming at improving the model performance and generalization capabilities by the curriculum learning scheme. Despite the demonstrated success of curriculum learning in classification tasks, applying the curriculum learning paradigm in the object detection context has rarely been researched [3, 36]. There are two potential causes of the situation. Firstly, the curriculum learning framework treats one image sample indivisibly and performs sample-wise manipulation for classification tasks. In contrast, one sample (image or point cloud) in the object detection context typically contains objects of varying difficulties, necessitating object-level difficulty predictions and pacing patterns. Secondly, modern object detectors usually have some hard-mining strategies built-in, for example, through OHEM [42], or focal loss [25] and its variants [22, 59]. Applying curriculum learning into object detection tasks needs careful balance between the easy-to-hard curriculum learning paradigm, and the built-in hard-mining scheme.

### 3.2.1 Difficulty Criterion

LiDAR-based object detectors usually assign a categorical likelihood for each predicted object of interest. Following

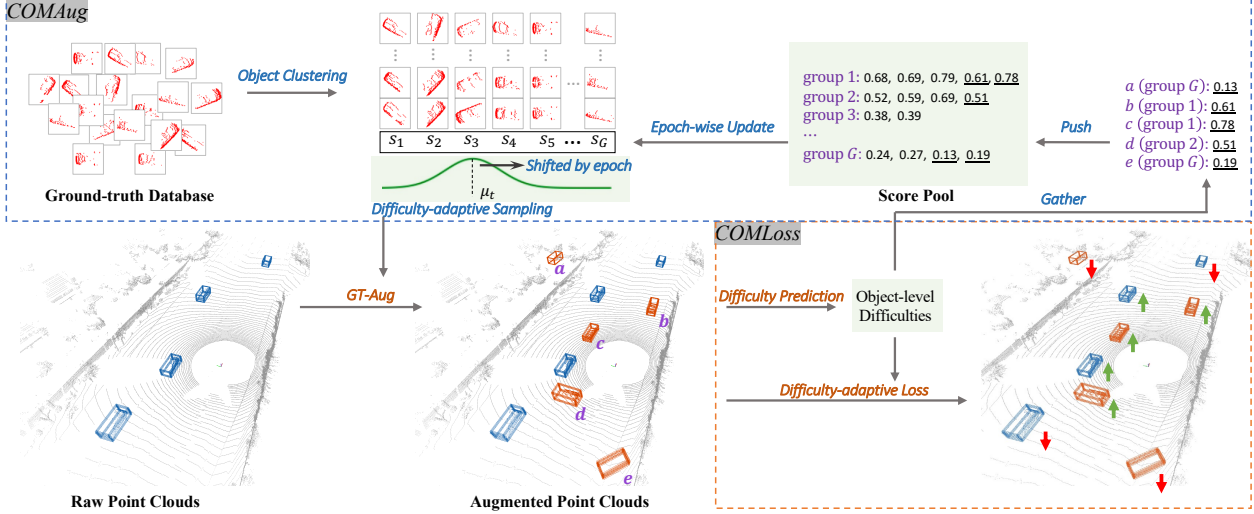


Figure 3. The proposed COM framework is mainly composed of COMAug and COMLoss. (1) COMAug firstly clusters objects in the ground-truth database into  $G$  groups by well-designed heuristics. At the beginning of each epoch, these groups are assigned with scores  $\{s_g\}_1^G$  based on a score pool storing the predicted difficulties of sampled objects in the last epoch. According to group scores  $\{s_g\}_1^G$ , the difficulty-adaptive sampling selects increasingly difficult objects for augmentation as training proceeds, which is realized by a curve with a shifted-by-epoch center  $\mu_t$ . (2) After GT-Aug inserts sampled objects into raw point clouds, COMLoss predicts difficulties for target objects in the augmented point clouds. Then object losses are re-weighted (marked by up/down arrows) with respect to object difficulties and training stages.

the idea that loss values can be used as alternative indicators [16, 19, 29, 34], we collect the predicted classification scores to gauge the object-level difficulties. Alternative solutions of object-level difficulties estimation include using handcrafted heuristics [43] and pre-trained models [27, 34]. Compared with these, our choice of using the classification scores as the difficulty criterion has clear advantages: (i) classification scores are evaluated on-the-fly as the model being trained, thus these scores align much better with the model optimization status quo, compared with pre-trained models or heuristics; (ii) classification scores are evaluated as part of the forward loss computation, thus no additional computation costs are incurred.

However, there still exists statistical divergence between the classification score distribution and the object difficulty distribution. Such divergence also changes as training progresses. For example, scores of easy objects tend to be small at early training stages, while hard object scores can be relatively large in later stages. Using  $\{s_p\}_1^P, \{s_q\}_1^Q$  to represent the collected classification scores for current original and augmented objects, we mitigate the discrepancies by introducing an adaptive threshold  $\tau$  which is essentially the running average of the original object scores:

$$\tau \leftarrow (1 - \alpha) \cdot \tau + \alpha \cdot (\sum_p s_p) / P, \quad (2)$$

where  $\tau$  is initialized as 0 and  $\alpha$  is the momentum coefficient. We do not include scores of the augmented objects  $q_i$  when estimating  $\tau$  as augmented objects can have fluctuat-

ing scores due to the sampling process randomness. Given adaptive threshold  $\tau$ , object difficulties are represented by:

$$\begin{aligned} \tilde{s}_p &= s_p - \tau, \quad \forall p \in \{1, 2, \dots, P\}, \\ \tilde{s}_q &= s_q - \tau, \quad \forall q \in \{1, 2, \dots, Q\}. \end{aligned} \quad (3)$$

It's noteworthy to mention that a small  $\tilde{s}_p$  corresponds to hard samples. Although this is slightly counterintuitive, we maintain it for the sake of clarity in the follow-up sections.

### 3.2.2 Difficulty-adaptive Loss

The goal of COMLoss is to formulate a mechanism to divert the training focus to different objects at different stages of model training. In particular, in earlier stages easier objects should be emphasized while in later stages harder objects should attract more attention. To reach this goal, we design an adaptive weighting function  $w$  to dynamically adjust the focus of the optimization objective as training proceeds:

$$\begin{aligned} w &= 1 + h_t \cdot (1 - e^{\beta \cdot \tilde{s}}) / (1 + e^{\beta \cdot \tilde{s}}), \\ \text{where } h_t &= H \cdot (t_r - t) / T. \end{aligned} \quad (4)$$

Here  $\tilde{s}$  is the object difficulty and  $\beta$  controls the curve shape.  $h_t$  is the height of the Sigmoid curve at epoch  $t$  and  $T$  is the number of total epochs.  $H$  is a parameter to control re-weighting degree. Before epoch  $t_r$ , the  $h_t$  is greater than 0 and the easy/hard samples are emphasized/suppressed respectively. For epochs  $t > t_r$ , hard samples are emphasized while easy samples become down-weighted. We call the

epoch  $t_r$  as the tipping point which is an important parameter for COMLoss. We found through experiments that the easy-to-hard strategy can still bring performance gains even after setting  $t_r$  very close to  $T$ . When  $t$  approaches  $t_r \approx T$ , the weighting function  $w$  becomes almost flat, thus treating all objects equally. We believe that the built-in hard example mining mechanisms (such as focal loss [22, 25, 59]) lead to the performance gains in such scenarios.

Given the weighting function  $w$ , the proposed COMLoss formulation becomes:

$$\mathcal{L} = \frac{1}{N} \left( \mathcal{L}_n + \sum_{p=1}^P w_p \cdot (\mathcal{L}_c^p + \mathcal{L}_r^p) + \sum_{q=1}^Q w_q \cdot (\mathcal{L}_c^q + \mathcal{L}_r^q) \right). \quad (5)$$

### 3.3. COMAug

During the training of LiDAR-based 3D object detectors, GT-Aug randomly samples objects from the ground-truth database  $\mathcal{B}$  and inserts them into current training point. Nevertheless, GT-Aug cannot foretell whether the sampled objects would benefit the training more than the left out ones. Intuitively, sampling harder samples in early stages can increase learning difficulty and consequently lead to unstable training. On the other hand, easier samples will not bring in much useful knowledge in the later stages.

One solution to this problem is the curricular scheme which samples increasingly difficult objects for augmentation. But known and up-to-date object difficulty scores are a prerequisite for such curricular scheme. To fulfill such requirement, a naive way is to initialize object difficulties with the same score value and update an object’s difficulty if only sampled for augmentation. This solution however has two significant flaws: (1) The probability of one object to be sampled is small because of the enormous volume of the ground-truth database. Thus, object difficulties are mostly not updated throughout training, or have predictions that are severely outdated. For example, the Waymo Open Dataset [45] contains 12.6 million 3D bounding box annotations in LiDAR data. The probability of one vehicle being sampled for augmentation during one epoch of training is less than 1%. (2) The sampling process can be easily trapped in a downward spiral: objects have limited chances of being selected for augmentation, thus their difficulty estimations can be inaccurate. Improperly sampled augmentation objects may introduce perturbations in model training, which further hinders the reliability of the updated difficulty estimations. These issues motivate the COMAug to cluster objects in the database (Sec. 3.3.1) and update group-level scores rather than individual ones (Sec. 3.3.2). By design, COMAug is able to perform difficulty-adaptive sampling (Sec. 3.3.3), achieving the curricular augmentation along the training process.

#### 3.3.1 Object Clustering

The core of COMAug is to assign objects with similar properties to the same cluster. Although point cloud provides rich geometry information for objects of interest, how to select effective grouping criterion is critical for COMAug. For example, extensive works have demonstrated that object distances [8, 11, 28, 37, 53], occlusion [51], object sizes and angles [58] are all influential factors for detection quality.

Assuming the bounding box for one object is  $[x, y, z, l, w, h, \alpha]$ , where  $x, y, z$  are the coordinates of the box center.  $l, w, h, \alpha$  are the length, width, height and heading of the box. Drawing inspirations from the literature and our empirical studies, the following four factors are used to perform object clustering in COMAug: (1) the distance  $f_d$  to the LiDAR sensor calculated by  $f_d = \sqrt{x^2 + y^2 + z^2}$ ; (2) the size  $f_s$  of the bounding box represented by  $f_s = \max(l, w, h)$ ; (3) the relative angle  $f_a$  between the box heading and the azimuth of the box center, *i.e.*,  $f_a = \alpha - \arctan(y/x)$ ; and (4) the occupancy ratio  $f_o$  of observed area over the entire bounding box. In our implementation, we divide the object into a particular number of equally spaced 3D voxels and assign the ratio of non-empty ones to  $f_o$ . Fig. 4 shows a visualization of clustering results from our experiments.




Figure 4. Visualization of several groups of the vehicle class. Each sub-graph shows the aggregation of point clouds from 20 random objects in one group. Top four rows show vehicles with sizes  $f_s \in [5, 8]$  meters, while bottom four rows shows vehicles with sizes no larger than 5 meters. In each four rows, the occupancy ratio  $f_o$  increases from top to bottom. Changes of angle  $f_a$  can be observed in the left-to-right direction.

#### 3.3.2 Group-level Score Update

We group the aggregated ground truth objects into  $g$  groups and each group has a difficulty score  $\{s_g\}_1^G$ . At any given

moment during training, difficulty scores of the augmented objects (*i.e.*,  $\{\tilde{s}_g\}_1^G$ ), as predicted from the COMLoss, are also gathered. For each group, we design a score pool  $\mathcal{P}_g$  which stores all the scores gathered during one epoch of training (see Fig. 3). At the end of each epoch, the score of each group  $s_g$  is updated to the average of the score pool  $\mathcal{P}_g$ . Compared with the aforementioned individual-level update paradigm, the group-level score update can greatly lessen the issues of outdated and unstable score predictions.

Also, we found that for group score  $s_g$ , epoch-wise update works better than momentum update in COMLoss. It has brought to our attention that the distribution of group score pool sizes  $|\mathcal{P}_g|$  can be highly skewed during training. For example, easy groups can have a much larger score pool, especially in early stages. Momentum update thus will make groups update in different paces, hindering the effectiveness of carefully designed curriculum learning paradigm. Epoch-wise update, on the other hand, can circumvent the issue.

### 3.3.3 Difficulty-adaptive Sampling

The last component of COMAug is a difficulty-adaptive sampler. Given the sampler, for each query, we first select one group, and then randomly sample one object from the chosen group. The sampler should further fulfill these requirements: (1) adaptive to difficulty scores; (2) adaptive to training stages; and (3) robust to the clustering results.

We introduce a Gaussian sampling curve centered at  $\mu_t$  with hyperparameter  $\sigma$ , where  $\mu_t$  shifts with epoch number  $t$ . With this curve, each group  $g$  is assigned with a probability  $p_g = \exp\{-\frac{(\tilde{s}_g - \mu_t)^2}{2\sigma^2}\}$ . By this means, the sampling is adaptive to difficulties as groups with scores close to  $\mu_t$  are of high chances to be selected. We consequently make  $\mu_t$  to change by epoch  $t$  to meet the second requirement. Without loss of generality, we assume the group scores  $\{s_i\}_1^G$  are already sorted in descending order. Then we assign  $\mu_t$  with the score from harder groups as epoch  $t$  increases. Specifically, we let

$$\mu_t = s_g, \text{ where } g = \min(\lfloor (\lambda \cdot t/T) \cdot G \rfloor, G), \quad (6)$$

where  $T$  is the number of total epochs and  $\lambda$  is a hyperparameter controlling the pacing speed.

Directly using the probability  $p_g$  to sample the groups can be problematic when the clustered groups have highly imbalanced numbers of elements. For example, with the same group difficulties, objects in tiny groups can be repeatedly sampled while those in the large group will be under-sampled. Therefore, we propose to normalize the  $p_g$  by making summations of all probabilities to be 1 as well as considering the group size  $n_g$ . In the end, the probability of

selecting group  $g$  is computed as:

$$p^g = (p_g \cdot n_g) / \left( \sum_{i=1}^G p_i \cdot n_i \right). \quad (7)$$

## 4. Experiments

In this section, we detail the experimental setup in Sec. 4.1 and present the main results in Sec. 4.2. Comprehensive ablation studies on the proposed COMLoss and COMAug are shown in Sec. 4.3 and Sec. 4.4, respectively.

### 4.1. Experimental Setup

**Implementation Details.** We implement our method on top of OpenPCDet [48] and use the default configurations. Specifically, we train all detectors for 30 epochs by an Adam one cycle optimizer. The learning rate is warmed up from 0.003 to 0.03 within the first 12 epochs before dropping. Experiments are conducted on 8 A100 GPUs and batch size in each GPU is 2 for CenterPoint [58], PointPillars [21] and SECOND [52]. Note here we use the one-stage and pillar-based version of CenterPoint. We fix random seeds for fair comparisons. For GT-Aug, the thresholds  $\Gamma$  for class vehicle, pedestrian and cyclist are 15, 10, 10 respectively.

**Dataset and Metrics.** Our experiments are conducted on the large-scale Waymo Open Dataset [45] (WOD) which contains 798 sequences for training, 202 for validation, and 150 for testing respectively. 20% training data (32k frames) are uniformly sampled for training. For evaluation, we report the 3D LEVEL\_1 average precision ( $AP_{L1}$ ) and LEVEL\_2 average precision ( $AP_{L2}$ ) on the WOD validation set, where LEVEL\_1 and LEVEL\_2 represent the level of difficulty.

**COM Configurations.** Unless stated otherwise, the parameters of COMLoss are default to  $\alpha = 0.001$  for Eq. (2) and  $t_r = 30, \beta = -5$  for Eq. (4).  $H$  in Eq. (4) is the hyperparameter in experiments. We let  $\lambda = 0.5$  and  $\sigma = 0.2$  in the sampling curve in COMAug. Details of object clustering are placed in the supplementary due to the limited space.

### 4.2. Main Results

In this part, we show the efficacy and generality of the proposed COM framework on three popular LiDAR-based 3D object detectors: SECOND [52], PointPillars[21] and CenterPoint [58]. We compare the performances of these detectors when using GT-Aug (denoted as +GT-Aug), using COMLoss and GT-Aug (denoted as +COMLoss), using COMAug (denoted as +COMAug), using both COMLoss and COMAug (denoted as +COM), and using vanilla configurations (denoted as +None). Tab. 1 and Tab. 2 present the results on the pedestrian and vehicle categories.

In Tab. 1, the improvements brought by GT-Aug on  $AP_{L1}$  are 3.03%, 1.18% and 0.83% for the three detec-

Detector	Method	Pedestrian	
		3D AP <sub>L1</sub> (%)	3D AP <sub>L2</sub> (%)
SECOND [52]	+ None	62.37	54.05
	+ GT-Aug	65.40 (+3.03)	57.07 (+3.02)
	+ COMLoss	66.23 (+3.86)	57.77 (+3.72)
	+ COMAug	66.46 (+4.09)	57.97 (+3.92)
	+ COM	66.66 (+4.29)	58.14 (+4.09)
PointPillars [21]	+ None	64.36	55.98
	+ GT-Aug	65.54 (+1.18)	57.20 (+1.22)
	+ COMLoss	66.19 (+1.83)	57.80 (+1.82)
	+ COMAug	66.75 (+2.39)	58.40 (+2.42)
	+ COM	66.81 (+2.45)	58.52 (+2.54)
CenterPoint [58]	+ None	72.79	65.02
	+ GT-Aug	73.62 (+0.83)	65.84 (+0.82)
	+ COMLoss	74.16 (+1.37)	66.36 (+1.34)
	+ COMAug	75.34 (+2.55)	67.66 (+2.64)
	+ COM	75.43 (+2.64)	67.76 (+2.74)

Table 1. Performances on class pedestrian. In green are the gaps of at least +2.0 point.

Detector	Method	Vehicle	
		3D AP <sub>L1</sub> (%)	3D AP <sub>L2</sub> (%)
SECOND [52]	+ None	71.72	63.26
	+ GT-Aug	72.40 (+0.68)	63.95 (+0.69)
	+ COMLoss	72.82 (+1.10)	64.35 (+1.09)
	+ COMAug	73.07 (+1.35)	64.57 (+1.31)
	+ COM	73.32 (+1.60)	65.34 (+2.08)
PointPillars [21]	+ None	69.28	61.04
	+ GT-Aug	70.00 (+0.72)	61.70 (+0.66)
	+ COMLoss	70.78 (+1.50)	62.43 (+1.39)
	+ COMAug	71.01 (+1.73)	62.88 (+1.84)
	+ COM	71.38 (+2.10)	62.96 (+1.92)
CenterPoint [58]	+ None	70.97	62.54
	+ GT-Aug	71.05 (+0.08)	62.66 (+0.12)
	+ COMLoss	71.57 (+0.60)	63.13 (+0.59)
	+ COMAug	72.03 (+1.06)	63.61 (+1.07)
	+ COM	72.15 (+1.18)	64.29 (+1.75)

Table 2. Performances on class vehicle. In green are the gaps of at least +1.0 point.

tors. COMLoss further increases these values into 3.86%, 1.83% and 1.37%. Finally, our COM improves the baseline detectors by 4.09%, 2.54%, 2.74% on AP<sub>L2</sub>, respectively. In Tab. 2, on strongest CenterPoint detector COMLoss and COMAug lead to the improvements of 0.60% and 1.06%, respectively. The COM framework promotes the AP<sub>L1</sub> to 72.15, which beats the baseline value of 70.97% by 1.18%.

### 4.3. Ablation Studies on COMLoss

This part examines three key parameters for COMLoss in Eq. (4): tipping point  $t_r$ , parameters  $H$ , and  $\beta$ . Using the CenterPoint [58] with GT-Aug as the baseline, we report the improvements on 3D AP<sub>L1</sub> brought by our COMLoss.

**Effects of parameter  $\beta$ .**  $\beta = 0$  degrades model to the baseline and  $\beta < 0$  leads to the traditional easy-to-hard strategy. The reverse-curricular scheme is achieved using  $\beta > 0$ , where the model emphasizes hard samples at early stages and easy samples at later stages. On the other hand, decreasing the absolute value  $|\beta|$  flattens the curve. We conduct experiments with  $\beta$  ranging from -20 to 15, and present our results in Fig. 5. We find that the reverse-curricular scheme causes performance degradation, while steeper curve worsens the situation. In contrast, the easy-to-hard scheme can benefit the detector and  $\beta = -3$  achieves the best result.




Figure 5. Effects of parameter  $\beta$  on vehicle ( $H = 1$ ).

**Effects of parameter  $H$ .** Fig. 6 shows the result of the ablation study on the parameter  $H$ , which essentially reflects the extent of re-weighting objects. We observe that the improvements brought by COMLoss increase first and then decrease as  $H$  grows larger, and the maximum performance gain ( $> 0.5\%$ ) is obtained at  $H = 0.6$ .




Figure 6. Effects of parameter  $H$  on vehicle.

**Effects of tipping point  $t_r$ .** COMLoss emphasizes easy and hard objects correspondingly before and after the epoch  $t_r$ . This is the most crucial parameter in COMLoss as it strikes a balance between the introduced easy-to-hard scheme and the built-in hard-mining strategy in detector. We conduct extensive experiments on  $t_r$  with various height  $H$  across multiple object categories, and present the results in Fig. 7. For the class of vehicle, the performances increase with  $t_r$  for both  $H = 1$  and  $H = 0.6$ . The classes of pedestrian and cyclist observe different optimal  $t_r$  in different cases. For example, the best  $t_r$  on pedestrian is 20 for  $H = 0.6$  and 30 for  $H = 1$ . On cyclist, the best  $t_r$  is 30, 20, 25 when  $H$  is 0.3, 0.6 and 1, respectively. In each

situation, the optimal  $t_r$  is near the end of the training. Additionally, the largest improvements are achieved with different  $H$ s (*i.e.*,  $H = 0.6, 1, 0.3$  for vehicle, pedestrian and cyclist) but with the same  $t_r$  of 30. This situation may be caused by the built-in hard-mining scheme in detectors. Setting  $t_r$  as 30 means that the re-weighting curve treats all objects equally in the end. However, hard objects are still been emphasized because of the built-in mechanism. The whole training therefore still follows the easy-to-hard principle.

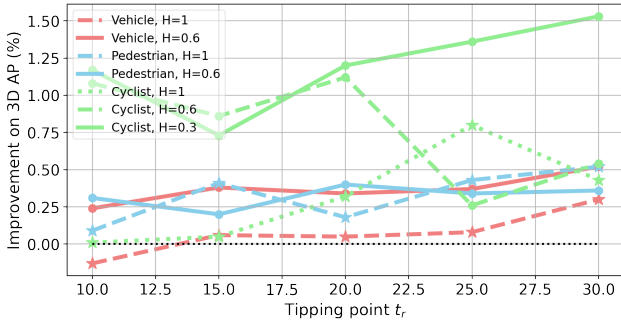


Figure 7. Effects of tipping point  $t_r$ . **Best viewed in color.**

#### 4.4. Analysis on COMAug

This part presents our analysis on COMAug including the ablation study on the object clustering as well as several visualizations. All experiments in this part are conducted with CenterPoint [58] on the vehicle category from the Waymo Open Dataset.

**Effects of clustering factors.** COMAug clusters objects by four factors including distance  $f_d$ , size  $f_s$ , angle  $f_a$  and occupancy ratio  $f_o$ . We examine the performances of COMAug with various combinations of these factors, and report the results in Tab. 3. Using no factor (A0) degrades the model to the GT-Aug baseline. Even with one factor, COMAug (A1-A4) can already result in noticeable performance gains. The best performing single factor A2 indicates that distance is the main contributor to object difficulty in LiDAR object detection. Improvements from A5 to A8 increase steadily as more factors are taken into account. Final results are of improvements 0.98% and 0.95% on  $AP_{L1}$  and  $AP_{L2}$ , respectively, when four factors are utilized.

**Visualization of sampling probability.** Based on the difficulties predicted in the last epoch, we categorize the groups into easy/medium-hard/hard parts which have top 1/3, mid 1/3 and bottom 1/3 of difficulty scores. Consequently, we visualize the mean sampling probabilities for each part as shown in Fig. 8. Note that we ignore the first epoch where all probabilities are initialized to be the same. The trends of these curves are shown as expected: easy groups are overwhelmingly sampled in early stages while rarely sampled in the end. The sampling for hard groups is in an opposite pattern. For medium groups, the sampling probabilities

increase first and then decrease, and peak in the middle of training.

	Occup.	Dist.	Angle	Size	3D $AP_{L1}$ (%)	3D $AP_{L2}$ (%)
A0					71.05	62.66
A1	✓				71.36 (+0.31)	62.93 (+0.33)
A2		✓			71.51 (+0.46)	63.11 (+0.45)
A3			✓		71.33 (+0.28)	62.92 (+0.32)
A4				✓	71.38 (+0.33)	62.98 (+0.38)
A5	✓	✓			71.68 (+0.63)	63.24 (+0.31)
A6	✓	✓	✓		71.86 (+0.81)	63.41 (+0.75)
A7	✓	✓		✓	71.89 (+0.84)	63.44 (+0.78)
A8	✓	✓	✓	✓	72.03 (+0.98)	63.61 (+0.95)

Table 3. Performances with different clustering factors. Occup. and Dist. stand for occupancy ratio and distance, respectively.




Figure 8. Visualization of sampling probability.

## 5. Conclusions and Limitations

In this work, we explore the potentials of curriculum learning in LiDAR-based object detection by proposing a novel curricular object manipulation (COM) framework. We focus on loss design in the detector and sampling strategy in GT-Aug, and correspondingly introduce the COM-Loss and COMAug modules. Extensive experiments and ablation studies on large scale benchmarks verified the efficacy of our method.

However, there is still room for future research on applying curriculum learning to 3D object detection. To inspire future work, we outline a few limitations based on our understanding: (1) our difficulty criterion depends solely on the classification loss while the regression difficulty is ignored for the sake of efficiency. Moreover, loss-based criterion is naturally affected by training noise and hysteresis effect, and therefore may not precisely reveal difficulties. Designing more accurate criterion is a promising direction. (2) we group objects by four empirically validated heuristics. More efforts are deserved for more appropriate grouping strategies as well. (3) we limit the work in the LiDAR-based object detector with GT-Aug. This work can be extended to other LiDAR-related tasks such as multi-modality 3D object detection and point cloud segmentation.



## References

- [1] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *International Conference on Machine Learning*, pages 41–48, 2009. 2, 3
- [2] Alex Bewley, Pei Sun, Thomas Mensink, Dragomir Anguelov, and Cristian Sminchisescu. Range conditioned dilated convolutions for scale invariant 3d object detection. *arXiv preprint arXiv:2005.09927*, 2020. 2
- [3] Thibault Castells, Philippe Weinzaepfel, and Jerome Revaud. Superloss: A generic loss for robust curriculum learning. *Annual Conference on Neural Information Processing Systems*, 33:4308–4319, 2020. 3
- [4] Yuning Chai, Pei Sun, Jiquan Ngiam, Weiyue Wang, Benjamin Caine, Vijay Vasudevan, Xiao Zhang, and Dragomir Anguelov. To the point: Efficient 3d object detection in the range image with graph convolution kernels. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2021. 2
- [5] Haw-Shiuan Chang, Erik Learned-Miller, and Andrew McCallum. Active bias: Training more accurate neural networks by emphasizing high variance samples. *Annual Conference on Neural Information Processing Systems*, 30, 2017. 3
- [6] Yilun Chen, Shu Liu, Xiaoyong Shen, and Jiaya Jia. Fast point r-cnn. In *International Conference on Computer Vision*, pages 9775–9784, 2019. 1, 2
- [7] MMDetection3D Contributors. MMDetection3D: OpenMMLab next-generation platform for general 3D object detection. <https://github.com/open-mmlab/mmdetection3d>, 2020. 2, 3
- [8] Jiajun Deng, Shaoshuai Shi, Peiwei Li, Wengang Zhou, Yanyong Zhang, and Houqiang Li. Voxel r-cnn: Towards high performance voxel-based 3d object detection. In *AAAI Conference on Artificial Intelligence*, volume 35, pages 1201–1209, 2021. 2, 5
- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2
- [10] Kaiwen Duan, Song Bai, Lingxi Xie, Honggang Qi, Qingming Huang, and Qi Tian. Centernet: Keypoint triplets for object detection. In *International Conference on Computer Vision*, pages 6569–6578, 2019. 2
- [11] Lue Fan, Ziqi Pang, Tianyuan Zhang, Yu-Xiong Wang, Hang Zhao, Feng Wang, Naiyan Wang, and Zhaoxiang Zhang. Embracing single stride 3d object detector with sparse transformer. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 8458–8468, 2022. 2, 5
- [12] Lue Fan, Xuan Xiong, Feng Wang, Naiyan Wang, and Zhaoxiang Zhang. Rangedet: In defense of range view for lidar-based 3d object detection. In *International Conference on Computer Vision*, pages 2918–2927, 2021. 2
- [13] Jin Fang, Xinxin Zuo, Dingfu Zhou, Shengze Jin, Sen Wang, and Liangjun Zhang. Lidar-aug: A general rendering-based augmentation framework for 3d object detection. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 4710–4720, 2021. 1
- [14] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2012. 3
- [15] Guy Hacohen and Daphna Weinshall. On the power of curriculum learning in training deep networks. In *International Conference on Machine Learning*, pages 2535–2544. PMLR, 2019. 3
- [16] Yuge Huang, Yuhan Wang, Ying Tai, Xiaoming Liu, Pengcheng Shen, Shaoxin Li, Jilin Li, and Feiyue Huang. Curricularface: adaptive curriculum learning loss for deep face recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 5901–5910, 2020. 2, 4
- [17] Lu Jiang, Zhengyuan Zhou, Thomas Leung, Li-Jia Li, and Li Fei-Fei. Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In *International Conference on Machine Learning*, pages 2304–2313. PMLR, 2018. 3
- [18] Sihyeon Kim, Sanghyeok Lee, Dasol Hwang, Jaewon Lee, Seong Jae Hwang, and Hyunwoo J Kim. Point cloud augmentation with weighted local transformations. In *International Conference on Computer Vision*, pages 548–557, 2021. 1
- [19] Yajing Kong, Liu Liu, Jun Wang, and Dacheng Tao. Adaptive curriculum learning. In *International Conference on Computer Vision*, pages 5067–5076, 2021. 3, 4
- [20] Jason Ku, Melissa Mozifian, Jungwook Lee, Ali Harakeh, and Steven L Waslander. Joint 3d proposal generation and object detection from view aggregation. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1–8. IEEE, 2018. 2
- [21] Alex H Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. Pointpillars: Fast encoders for object detection from point clouds. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 12697–12705, 2019. 1, 2, 6, 7
- [22] Xiang Li, Wenhai Wang, Lijun Wu, Shuo Chen, Xiaolin Hu, Jun Li, Jinhui Tang, and Jian Yang. Generalized focal loss: Learning qualified and distributed bounding boxes for dense object detection. *Annual Conference on Neural Information Processing Systems*, 33:21002–21012, 2020. 3, 5
- [23] Zheng Li, Xiang Li, Lingfeng Yang, Borui Zhao, Renjie Song, Lei Luo, Jun Li, and Jian Yang. Curriculum temperature for knowledge distillation. *arXiv preprint arXiv:2211.16231*, 2022. 3
- [24] Zhichao Li, Feng Wang, and Naiyan Wang. Lidar r-cnn: An efficient and universal 3d object detector. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 7546–7555, 2021. 2
- [25] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *International Conference on Computer Vision*, pages 2980–2988, 2017. 3, 5
- [26] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*, pages 740–755. Springer, 2014. 3
- [27] Hongyuan Lu and Wai Lam. Pcc: Paraphrasing with bottom-

- k sampling and cyclic learning for curriculum data augmentation. *arXiv preprint arXiv:2208.08110*, 2022. 3, 4
- [28] Jiageng Mao, Yujing Xue, Minzhe Niu, Haoyue Bai, Jiashi Feng, Xiaodan Liang, Hang Xu, and Chunjing Xu. Voxel transformer for 3d object detection. In *International Conference on Computer Vision*, pages 3164–3173, 2021. 5
- [29] Qiang Meng, Shichao Zhao, Zhida Huang, and Feng Zhou. Magface: A universal representation for face recognition and quality assessment. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 14225–14234, 2021. 4
- [30] Gregory P Meyer, Ankit Laddha, Eric Kee, Carlos Vallespi-Gonzalez, and Carl K Wellington. Lasernet: An efficient probabilistic 3d object detector for autonomous driving. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 12677–12686, 2019. 2
- [31] Alexey Nekrasov, Jonas Schult, Or Litany, Bastian Leibe, and Francis Engelmann. Mix3d: Out-of-context data augmentation for 3d scenes. In *2021 International Conference on 3D Vision (3DV)*, pages 116–125. IEEE, 2021. 1
- [32] Charles R Qi, Or Litany, Kaiming He, and Leonidas J Guibas. Deep hough voting for 3d object detection in point clouds. In *International Conference on Computer Vision*, pages 9277–9286, 2019. 1, 2
- [33] Charles R Qi, Wei Liu, Chenxia Wu, Hao Su, and Leonidas J Guibas. Frustum pointnets for 3d object detection from rgb-d data. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 918–927, 2018. 2
- [34] Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. Learning to reweight examples for robust deep learning. In *International Conference on Machine Learning*, pages 4334–4343. PMLR, 2018. 3, 4
- [35] Ramit Sawhney, Ritesh Soun, Shrey Pandit, Megh Thakkar, Sarvagya Malaviya, and Yuval Pinter. Ciaug: Equipping interpolative augmentation with curriculum learning. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1758–1764, 2022. 3
- [36] Shreyas Saxena, Oncel Tuzel, and Dennis DeCoste. Data parameters: A new family of parameters for learning a differentiable curriculum. *Annual Conference on Neural Information Processing Systems*, 32, 2019. 3
- [37] Shaoshuai Shi, Chaoxu Guo, Li Jiang, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. Pv-rcnn: Point-voxel feature set abstraction for 3d object detection. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 10529–10538, 2020. 2, 3, 5
- [38] Shaoshuai Shi, Li Jiang, Jiajun Deng, Zhe Wang, Chaoxu Guo, Jianping Shi, Xiaogang Wang, and Hongsheng Li. Pv-rcnn++: Point-voxel feature set abstraction with local vector representation for 3d object detection. *arXiv preprint arXiv:2102.00463*, 2021. 3
- [39] Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. Pointcnn: 3d object proposal generation and detection from point cloud. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–779, 2019. 2
- [40] Shaoshuai Shi, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. From points to parts: 3d object detection from point cloud with part-aware and part-aggregation network. *IEEE Trans on Pattern Analysis and Machine Intelligence*, 43(8):2647–2664, 2020. 2
- [41] Weijing Shi and Raj Rajkumar. Point-gnn: Graph neural network for 3d object detection in a point cloud. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1711–1719, 2020. 1, 2
- [42] Abhinav Shrivastava, Abhinav Gupta, and Ross Girshick. Training region-based object detectors with online hard example mining. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 761–769, 2016. 3
- [43] Petru Soviany, Radu Tudor Ionescu, Paolo Rota, and Nicu Sebe. Curriculum self-paced learning for cross-domain object detection. *Computer Vision and Image Understanding*, 204:103166, 2021. 4
- [44] Petru Soviany, Radu Tudor Ionescu, Paolo Rota, and Nicu Sebe. Curriculum learning: A survey. *International Journal of Computer Vision*, pages 1–40, 2022. 2, 3
- [45] 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 *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2446–2454, 2020. 5, 6
- [46] Pei Sun, Weiyue Wang, Yuning Chai, Gamaleldin Elsayed, Alex Bewley, Xiao Zhang, Cristian Sminchisescu, and Dragomir Anguelov. Rsn: Range sparse net for efficient, accurate lidar 3d object detection. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 5725–5734, 2021. 2
- [47] Tomoumi Takase, Ryo Karakida, and Hideki Asoh. Self-paced data augmentation for training neural networks. *Neurocomputing*, 442:296–306, 2021. 3
- [48] OpenPCDet Development Team. Openpcdet: An open-source toolbox for 3d object detection from point clouds. <https://github.com/open-mmlab/OpenPCDet>, 2020. 2, 3, 6
- [49] Jason Wei, Chengyu Huang, Soroush Vosoughi, Yu Cheng, and Shiqi Xu. Few-shot text classification with triplet networks, data augmentation, and curriculum learning. *arXiv preprint arXiv:2103.07552*, 2021. 3
- [50] Aoran Xiao, Jiaying Huang, Dayan Guan, Kaiwen Cui, Shijian Lu, and Ling Shao. Polarmix: A general data augmentation technique for lidar point clouds. *arXiv preprint arXiv:2208.00223*, 2022. 1
- [51] Qiangeng Xu, Yiqi Zhong, and Ulrich Neumann. Behind the curtain: Learning occluded shapes for 3d object detection. *arXiv preprint arXiv:2112.02205*, 2021. 5
- [52] Yan Yan, Yuxing Mao, and Bo Li. Second: Sparsely embedded convolutional detection. *Sensors*, 18(10):3337, 2018. 1, 2, 3, 6, 7
- [53] Honghui Yang, Zili Liu, Xiaopei Wu, Wenxiao Wang, Wei Qian, Xiaofei He, and Deng Cai. Graph r-cnn: Towards accurate 3d object detection with semantic-decorated local graph. *arXiv preprint arXiv:2208.03624*, 2022. 3, 5
- [54] Jihan Yang, Shaoshuai Shi, Zhe Wang, Hongsheng Li, and Xiaojuan Qi. St3d: Self-training for unsupervised domain adaptation on 3d object detection. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 10368–10378, 2021. 3
- [55] Zetong Yang, Yanan Sun, Shu Liu, and Jiaya Jia. 3dssd:

- Point-based 3d single stage object detector. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 11040–11048, 2020. [2](#), [3](#)
- [56] Zetong Yang, Yanan Sun, Shu Liu, Xiaoyong Shen, and Ji-aya Jia. Ipod: Intensive point-based object detector for point cloud. *arXiv preprint arXiv:1812.05276*, 2018. [1](#), [2](#)
- [57] Zetong Yang, Yanan Sun, Shu Liu, Xiaoyong Shen, and Ji-aya Jia. Std: Sparse-to-dense 3d object detector for point cloud. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1951–1960, 2019. [2](#)
- [58] Tianwei Yin, Xingyi Zhou, and Philipp Krahenbuhl. Center-based 3d object detection and tracking. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 11784–11793, 2021. [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [59] Haoyang Zhang, Ying Wang, Feras Dayoub, and Niko Sunderhauf. Varifocalnet: An iou-aware dense object detector. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 8514–8523, 2021. [2](#), [3](#), [5](#)
- [60] Yifan Zhang, Qingyong Hu, Guoquan Xu, Yanxin Ma, Jianwei Wan, and Yulan Guo. Not all points are equal: Learning highly efficient point-based detectors for 3d lidar point clouds. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 18953–18962, 2022. [2](#)
- [61] Yin Zhou and Oncel Tuzel. Voxelnet: End-to-end learning for point cloud based 3d object detection. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 4490–4499, 2018. [1](#), [2](#)