

# Improving 3D Object Detection through Progressive Population Based Augmentation

Shuyang Cheng<sup>1</sup>, Zhaoqi Leng<sup>\*1</sup>, Ekin Dogus Cubuk<sup>2</sup>, Barret Zoph<sup>2</sup>, Chunyan Bai<sup>1</sup>, Jiquan Ngiam<sup>2</sup>, Yang Song<sup>1</sup>, Benjamin Caine<sup>2</sup>, Vijay Vasudevan<sup>2</sup>, Congcong Li<sup>1</sup>, Quoc V. Le<sup>2</sup>, Jonathon Shlens<sup>2</sup>, and Dragomir Anguelov<sup>1</sup>

<sup>1</sup> Waymo LLC

<sup>2</sup> Google LLC

**Abstract.** Data augmentation has been widely adopted for object detection in 3D point clouds. However, all previous related efforts have focused on manually designing specific data augmentation methods for individual architectures. In this work, we present the first attempt to automate the design of data augmentation policies for 3D object detection. We introduce the Progressive Population Based Augmentation (PPBA) algorithm, which learns to optimize augmentation strategies by narrowing down the search space and adopting the best parameters discovered in previous iterations. On the KITTI 3D detection test set, PPBA improves the StarNet detector by substantial margins on the moderate difficulty category of cars, pedestrians, and cyclists, outperforming all current state-of-the-art single-stage detection models. Additional experiments on the Waymo Open Dataset indicate that PPBA continues to effectively improve the StarNet and PointPillars detectors on a 20x larger dataset compared to KITTI. The magnitude of the improvements may be comparable to advances in 3D perception architectures and the gains come without an incurred cost at inference time. In subsequent experiments, we find that PPBA may be up to 10x more data efficient than baseline 3D detection models without augmentation, highlighting that 3D detection models may achieve competitive accuracy with far fewer labeled examples.

**Keywords:** Progressive population based augmentation, data augmentation, point cloud, 3D object detection, data efficiency

## 1 Introduction

LiDAR is a prominent sensor for autonomous driving and robotics because it provides detailed 3D information critical for perceiving and tracking real-world objects [2, 29]. The 3D localization of objects within LiDAR point clouds represents one of the most important tasks in visual perception, and much effort has focused on developing novel network architectures for point clouds [1, 33, 21,

---

\* Work done while at Google LLC.

32, 37, 31, 16, 25]. Following the image classification literature, such modeling efforts have employed manually designed data augmentation schemes for boosting performance [1, 31, 16, 33, 22, 36, 25, 37].

In recent years, much work in the 2D image literature has demonstrated that investing heavily into data augmentation may lead to gains comparable to those obtained by advances in model architectures [4, 38, 20, 11, 5]. Despite this, 3D detection models have yet to significantly leverage automated data augmentation methods (but see [18]). Naively porting ideas that are effective for images to point cloud data presents numerous challenges, as the types of augmentations appropriate for point clouds differ tremendously. Transformations appropriate for point clouds are typically geometric-based and may contain a large number of parameters. Thus, the search space proposed in [4, 38] may not be naively reused for an automated search in point cloud augmentation space. Finally, because the search space is far larger, employing a more efficient search method becomes a practical necessity. Several works have attempted to significantly accelerate the search for data augmentation strategies [20, 11, 5], however it is unclear if such methods transfer successfully to point clouds.

In this work, we demonstrate that automated data augmentation significantly improves the prediction accuracy of 3D object detection models. We introduce a new search space for point cloud augmentations in 3D object detection. In this search space, we find the performance distribution of augmentation policies is quite diverse. To effectively discover good augmentation policies, we present an evolutionary search algorithm termed *Progressive Population Based Augmentation* (PPBA). PPBA works by narrowing down the search space through successive iterations of evolutionary search, and by adopting the best parameters discovered in past iterations. We demonstrate that PPBA is effective at finding good data augmentation strategies across datasets and detection architectures. Additionally, we find that a model trained with PPBA may be up to 10x more data efficient, implying reduced human labeling demands for point clouds.

Our main contributions can be summarized as follows: (1) We propose an automated data augmentation technique for localization in 3D point clouds. (2) We demonstrate that the proposed search method effectively improves point cloud 3D detection models compared to random search with less computational cost. (3) We demonstrate up to a 10x increase in data efficiency when employing PPBA. (4) Beyond 3D detection, we also demonstrate that PPBA generalizes to 2D image classification.

## 2 Related Work

Data augmentation has been an essential technique for boosting the performance of 2D image classification and object detection models. Augmentation methods typically include manually designed image transformations, to which the labels remain invariant, or distortions of the information present in the images. For example, elastic distortions, scale transformations, translations, and rotations are beneficial on models trained on MNIST [26, 3, 30, 24]. Crops, image mirror-

ing and color shifting / whitening [14] are commonly adopted on natural image datasets like CIFAR-10 and ImageNet. Recently, cutout [6] and mixup [34] have emerged as data augmentation methods that lead to good improvements in natural image datasets. For object detection in 2D images, image mirroring and multi-scale training are popular distortions [10]. Dwibedi et al. add new objects on training images by cut-and-paste [7].

While the augmentation operations mentioned above are designed by domain experts, there are also automated approaches to designing data augmentation strategies for 2D images. Early attempts include Smart Augmentation, which uses a network to generate augmented data by merging two or more image samples [17]. Ratner et al. use GANs to output sequences of data augmentation operations [23]. AutoAugment uses reinforcement learning to optimize data augmentation strategies for classification [4] and object detection [38]. More recently, improved search methods are able to find data augmentation strategies more efficiently [5, 11, 20].

While all the mentioned work so far is on 2D image classification and object detection, automated data augmentation methods have not been explored for 3D object detection tasks to the best of our knowledge. Models trained on KITTI use a wide variety of manually designed distortions. Due to the small size of the KITTI training set, data augmentation has been shown to improve performance significantly (common augmentations include horizontal flips, global scale distortions, and rotations) [1, 31, 33, 16, 25]. Yan et al. add new objects in training point clouds by pasting points inside ground truth 3D bounding boxes [31]. Despite its effectiveness for KITTI models, data augmentation was not used on some of the larger point cloud datasets [22, 36]. Very recently, an automated data augmentation approach was studied for point cloud classification [18].

Historically, 2D vision research has focused on architectural modifications to improve generalization. More recently, it was observed that improving data augmentation strategies can lead to comparable gains to a typical architectural advance [34, 8, 4, 38]. In this work, we demonstrate that a similar type of improvement can also be obtained by an effective automated data augmentation strategy for 3D object detection over point clouds.

### 3 Methods

We formulate the problem of finding the right augmentation strategy as a special case of hyperparameter schedule learning. The proposed method consists of two components: a specialized data augmentation search space for point cloud inputs and a search algorithm for the optimization of data augmentation parameters. We describe these two components below.

#### 3.1 Search Space for 3D Point Cloud Augmentation

In the proposed search space, an augmentation policy consists of  $N$  augmentation operations. Additionally, each operation is associated with a probability

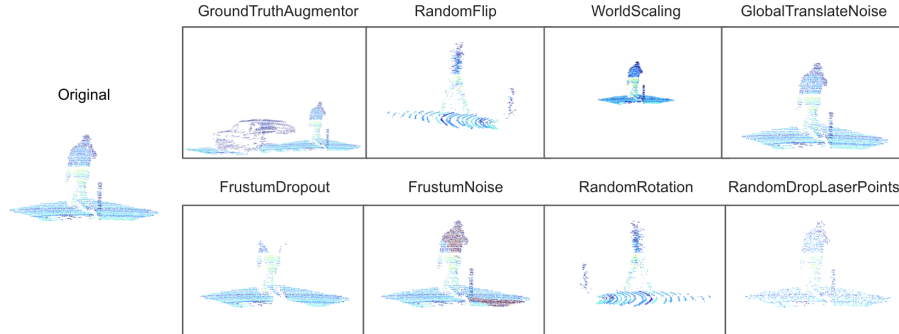


Fig. 1: **Visualization of the augmentation operations in the proposed search space.** An augmentation policy is defined by a list of distinct augmentation operations and the corresponding augmentation parameters. Details of these operations are in Table 7 and Table 8 in the Appendix.

and some specialized parameters. For example, the ground-truth augmentation operation has parameters denoting the probability for sampling vehicles, pedestrians, cyclists, etc.; the global translation noise operation has parameters for the distortion magnitude of the translation operation on x, y and z coordinates. To reduce the size of the search space and increase the diversity of the training data, these different operations are always applied according to some learned probabilities in the same, pre-determined order to point clouds during training.

The basic augmentation operations in the proposed search space fall into two main categories: *global operations*, which are applied to all points in a frame (such as rotation along Z-axis, coordinate scaling, etc.), and *local operations*, which are applied to points locally (such as dropping out points within a frustum, pasting points within bounding boxes from other frames, etc.). Our list of augmentation operations (see Fig. 1) consists of GroundTruthAugmentor, RandomFlip, WorldScaling, GlobalTranslateNoise, FrustumDropout, FrustumNoise, RandomRotation and RandomDropLaserPoints. In total, there are 8 augmentation operations and 29 operation parameters in the proposed search space.

### 3.2 Learning through Progressive Population Based Search

The proposed search process is maximizing a given metric  $\Omega$  on a model  $\theta$  by optimizing a schedule of augmentation operation parameters  $\lambda = (\lambda_t)_{t=1}^T$ , where  $t$  represents the number of iterative updates for the augmentation operation parameters during model training. For point cloud detection tasks, we use mean average precision (mAP) as the performance metric. The search process for the best augmentation schedule  $\lambda^*$  optimizes:

$$\lambda^* = \arg \max_{\lambda \in \mathcal{A}^T} \Omega(\theta) \quad (1)$$

During training, the objective function  $L$  (which is used for optimization of the model  $\theta$  given data and label pairs  $(\mathcal{X}, \mathcal{Y})$ ) is usually different from the actual performance metric  $\Omega$ , since the optimization procedure (i.e. stochastic gradient descent) requires a differentiable objective function. Therefore, at each iteration  $t$  the model  $\theta$  is optimizing:


$$\theta_t^* = \arg \min_{\theta \in \Theta} L(\mathcal{X}, \mathcal{Y}, \lambda^t) \quad (2)$$

During search, the training process of the model is split into  $\mathcal{N}$  iterations. At every iteration,  $\mathcal{M}$  models with different  $\lambda_t$  are trained in parallel and are afterwards evaluated with the metric  $\Omega$ . Models trained in all previous iterations are placed in a population  $\mathcal{P}$ . In the initial iteration, all model parameters and augmentation parameters are randomly initialized. After the first iteration, model parameters are determined through an *exploit* phase - inheriting from a better performing parent model by exploiting the rest of the population  $\mathcal{P}$ . The *exploit* phase is followed by an *exploration* phase, in which a subset of the augmentation operations will be explored for optimization by mutating the corresponding augmentation parameters in the parent model, while the remaining augmentation parameters will be directly inherited from the parent model.

Similar to Population Based Training [12], the *exploit* phase will keep the good models and replace the inferior models at the end of every iteration. In contrast with Population Based Training, the proposed method focuses only on a subset of the search space at each iteration. During the *exploration* phase, a successor might focus on a different subset of the parameters than its predecessor. In that case, the remaining parameters (parameters that the predecessor does not focus on) are mutated based on the parameters of the corresponding operations with the best overall performance. In Fig 2, we show an example of Progressive Population Based Augmentation. The complete PPBA algorithm is described in detail in Algorithm 1 in the Appendix.

### 3.3 Schedule Optimization with Historical Data

The search spaces for data augmentation are different between 2D images and 3D point clouds. For example, each operation in the AutoAugment [4] search space for 2D images has a single parameter. Furthermore, any value of this parameter within the predefined range leads to a reasonable image. For this reason, even sampling random augmentation policies for 2D images leads to some improvement in generalization [4, 5]. On the other hand, the augmentation operations for 3D point clouds are much harder to optimize. Each operation has several parameters, and a good range for these parameters is not known a priori. For example, there are five parameters – `theta_width`, `phi_width`, `distance`, `keep_prob` and `drop_type` – in the FrustumDropout operation. The analogous operation for 2D images is `cutout` [6], which has only one parameter. Therefore it is more challenging to discover optimal parameters for point cloud operations with limited resources.



**Fig 2: An example of Progressive Population Based Augmentation.** Four augmentation operations ( $a_1, a_2, a_3, a_4$ ) are applied to the input data during training; their parameter set comprises the full search space. During progressive population based search, only two augmentation operations out of the four are explored for optimization at every iteration. For example, at the beginning of iteration  $t - 1$ , augmentation parameters of ( $a_1, a_2$ ) are selected for exploration for the blue model while augmentation parameters of ( $a_3, a_4$ ) are selected for exploration for the purple model. At the end of training in iteration  $t - 1$ , an inferior model (the purple model) is *exploited* by a model with higher performance (the blue model). Afterwards, a successor will inherit both model parameters and augmentation parameters from the winner model - the blue model. During the *exploration* phase, the selected augmentation operations for exploration by the successor model are randomly sampled and become ( $a_2, a_3$ ). Since augmentation parameters of  $a_3$  have not been explored by the predecessor (the blue model), corresponding augmentation parameters of the best model (the green model), in which  $a_3$  has been selected for exploration, will be adopted for exploration by the successor model.

In order to learn the parameters for individual operations effectively, PPBA modifies only a small portion of the parameters in the search space at every iteration, and the historical information from the previous iterations are reused to optimize the augmentation schedule. By narrowing down the focus on certain

subsets of the search space, it becomes easier to distinguish the inferior augmentation parameters. To mitigate the slowing down of search speed caused by the search space shrinkage at each training iteration, the best parameters of each operation discovered in the past iterations are adopted by the successors, when their focused subsets of the search space are different from their predecessors.

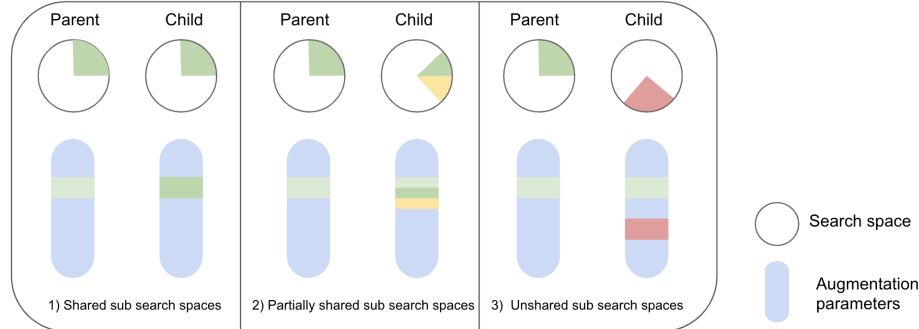


Fig. 3: **Three types of scenarios for the subsets of the search space explored by the parent and the child models.** 1) The subsets are the same. 2) The subsets are partially shared. 3) The subsets are unshared. In both 1) and 2), the overlapped augmentation parameters for exploration in the child model are mutated based on the corresponding parameters in the parent model (updating from light green to dark green). In both 2) and 3), the non-overlapped augmentation parameters for exploration in the child model are mutated based on the best augmentation parameters discovered in the past iterations (if available) or random sampling (updating from blue to yellow/red).

In Fig 3, we show three types of scenarios for the subsets of the search space explored by a successor and its predecessor. The details of the exploration phase based on historical data are described in Algorithm 2 in the Appendix.

## 4 Experiments

In this section, we empirically investigate the performance of PPBA on predictive accuracy, computational efficiency and data efficiency. We focus on single-stage detection models due to their simplicity, speed advantages and widespread adoption [31, 16, 22].

We first benchmark PPBA on the KITTI object detection benchmark [9] and the Waymo Open Dataset [27] (Sections 4.1 and 4.2). Our results show PPBA improves the baseline models and the improvement magnitude may be comparable to advances in 3D perception architectures. Next, we compare PPBA with random search and PBA [11] on the KITTI Dataset (Section 4.3) to demonstrate PPBA’s effectiveness and efficiency. In addition, we study the data efficiency of

PPBA on the Waymo Open Dataset (Section 4.4). Our experiments show that PPBA can achieve competitive accuracy with far fewer labeled examples compared with no augmentation. Finally, the PPBA algorithm was designed for, but is not limited to, 3D object detection tasks. We study its ability to generalize to 2D image classification and present results in Section 4.5.

#### 4.1 Surpassing Single-Stage Models on the KITTI Dataset

The KITTI Dataset [9] is generally recognized to be a small dataset for modern methods, and thus, data augmentation is critical to the performance of models trained on it [31, 16, 22]. We evaluate PPBA with StarNet [22] on the KITTI *test* split in Table 1. PPBA improves the detection performance of StarNet significantly, outperforming all current state-of-the-art single-stage point cloud detection models on the moderate difficulty category.

Table 1: Performance comparison of single-stage point cloud detection models on the KITTI *test* set using 3D evaluation. mAP is calculated with an IOU of 0.7, 0.5 and 0.5 for vehicles, cyclists and pedestrians, respectively

Method	Car			Pedestrian			Cyclist		
	Easy	Mod.	Hard	Easy	Mod.	Hard	Easy	Mod.	Hard
ContFuse [19]	83.68	68.78	61.67	-	-	-	-	-	-
VoxelNet [37]	77.47	65.11	57.73	39.48	33.69	31.51	61.22	48.36	44.37
SECOND [31]	83.13	73.66	66.20	51.07	42.56	37.29	70.51	53.85	46.90
3D IoU Loss [35]	<b>84.43</b>	76.28	68.22	-	-	-	-	-	-
PointPillars [16]	79.05	74.99	68.30	52.08	43.53	41.49	75.78	59.07	52.92
StarNet [22]	81.63	73.99	67.07	48.58	41.25	39.66	73.14	58.29	52.58
StarNet [22] + PPBA	84.16	<b>77.65</b>	<b>71.21</b>	<b>52.65</b>	<b>44.08</b>	<b>41.54</b>	<b>79.42</b>	<b>61.99</b>	<b>55.34</b>

During the PPBA search, 16 trials are trained to optimize the mAP for car (30 iterations) and for pedestrian/cyclist (20 iterations), respectively. The same training and inference settings<sup>3</sup> as [22] are used, while all trials are trained on the *train* split (3,712 samples) and validated on the *validation* split (3,769 samples). We train the first iteration for 3,000 steps, and all subsequent iterations for 1,000 steps with batch size 64. The search is conducted in the search space described in Section 3.1.

Manually designed augmentation policies are typically kept constant during training. In contrast, stochasticity lies at the heart of the augmentation policies in PPBA, i.e. each operation is applied stochastically and its parameters evolve as the training progresses. We have found that simply using the final parameters discovered by PPBA gets worse results than PPBA.

<sup>3</sup> <http://github.com/tensorflow/lingvo>



We use GroundTruthAugmentor to highlight the difference between the manually designed and learned augmentation policies. While training a StarNet vehicle detection model on KITTI with PPBA, the probability of applying the operation decreases from 100% to 16% and the probability of pasting vehicles reduces from 100% to 21%, while the probability of pasting pedestrians and cyclists increases from 0% to 28% and 8% respectively. This suggests that pasting the object of interest in every frame during training, as in manual designed policies, is not an optimal strategy and introducing a diverse set of objects from other classes is beneficial.

#### 4.2 Automated Data Augmentation Benefits Large-Scale Data

The Waymo Open Dataset is a recently released, large-scale dataset for 3D object detection in point clouds [27]. The dataset contains roughly 20x more scenes than KITTI, and roughly 20x more human-annotated objects per scene. This dataset presents an opportunity to ask whether data augmentations – being critical to model performance on the KITTI dataset due to the small size of the dataset – continue to provide a benefit in a large-scale training setting more reflective of the self-driving conditions in the real world.

To address this question, we evaluate the proposed method on the Waymo Open Dataset. In particular, we evaluate PPBA with StarNet [22] and PointPillars [16] on the *test* split in Table 2 and Table 3 on both LEVEL\_1 and LEVEL\_2 difficulties at different ranges. Our results indicate that PPBA notably improves the predictive accuracy of 3D detection across architectures, difficulty levels and object classes. These results indicate that data augmentation remains an important method for boosting model performance even in large-scale dataset settings. Furthermore, the gains due to PPBA may be as large as changing the underlying architecture, without any increase in inference cost.

Table 2: Performance comparison on the Waymo Open Dataset *test* set for vehicle detection. Note that the results of PointPillars [16] on the Waymo Open Dataset are reproduced by [27]

Method	Difficulty Level	3D mAP (IoU=0.7)				3D mAPH (IoU=0.7)			
		Overall	0-30m	30-50m	50m-Inf	Overall	0-30m	30-50m	50m-Inf
StarNet [22]	1	61.5	82.2	56.6	32.2	61.0	81.7	56.0	31.8
StarNet [22] + PPBA	1	<b>64.6</b>	<b>85.8</b>	<b>59.5</b>	<b>35.1</b>	<b>64.1</b>	<b>85.3</b>	<b>58.9</b>	<b>34.6</b>
StarNet [22]	2	54.9	81.3	49.5	23.0	54.5	80.8	49.0	22.7
StarNet [22] + PPBA	2	<b>56.2</b>	<b>82.8</b>	<b>54.0</b>	<b>26.8</b>	<b>55.8</b>	<b>82.3</b>	<b>53.5</b>	<b>26.4</b>
PointPillars [16]	1	63.3	82.3	59.2	35.7	62.8	81.9	58.5	35.0
PointPillars [16] + PPBA	1	<b>67.5</b>	<b>86.7</b>	<b>63.5</b>	<b>39.4</b>	<b>67.0</b>	<b>86.2</b>	<b>62.9</b>	<b>38.7</b>
PointPillars [16]	2	55.6	81.2	52.9	27.2	55.1	80.8	52.3	26.7
PointPillars [16] + PPBA	2	<b>59.6</b>	<b>85.6</b>	<b>57.6</b>	<b>30.0</b>	<b>59.1</b>	<b>85.1</b>	<b>57.0</b>	<b>29.5</b>

When performing the search with PPBA, 16 trials are trained to optimize the mAP for car and pedestrian, respectively. The list of augmentation operations

Table 3: Performance comparison on the Waymo Open Dataset *test* set for pedestrian detection. Note that the results of PointPillars [16] on the Waymo Open Dataset are reproduced by [27]

Method	Difficulty Level	3D mAP (IoU=0.5)				3D mAPH (IoU=0.5)			
		Overall	0-30m	30-50m	50m-Inf	Overall	0-30m	30-50m	50m-Inf
StarNet [22]	1	67.8	76.0	66.5	55.3	59.9	67.8	59.2	47.0
StarNet [22] + PPBA	1	<b>69.7</b>	<b>77.5</b>	<b>68.7</b>	<b>57.0</b>	<b>61.7</b>	<b>69.3</b>	<b>61.2</b>	<b>48.4</b>
StarNet [22]	2	61.1	73.1	61.2	44.5	54.0	65.2	54.5	37.8
StarNet [22] + PPBA	2	<b>63.0</b>	<b>74.8</b>	<b>63.2</b>	<b>46.5</b>	<b>55.8</b>	<b>66.8</b>	<b>56.2</b>	<b>39.4</b>
PointPillars [16]	1	62.1	71.3	60.1	47.0	50.2	59.0	48.3	35.8
PointPillars [16] + PPBA	1	<b>66.4</b>	<b>74.7</b>	<b>64.8</b>	<b>52.7</b>	<b>54.4</b>	<b>62.5</b>	<b>52.5</b>	<b>41.2</b>
PointPillars [16]	2	55.9	68.6	55.2	37.9	45.1	56.7	44.3	28.8
PointPillars [16] + PPBA	2	<b>60.1</b>	<b>72.2</b>	<b>59.7</b>	<b>42.8</b>	<b>49.2</b>	<b>60.4</b>	<b>48.2</b>	<b>33.4</b>

described in Section 3.1, except for GroundTruthAugmentor and RandomFlip, are used during search. In our experiments, we have found RandomFlip has a negative impact on heading prediction for both car and pedestrian.

For both StarNet and PointPillars on the Waymo Open Dataset, the same training and inference settings<sup>4</sup> as [22] is used. All trials are trained on the full *train* set and validated on the 10% *validation* split (4,109 samples). For StarNet, we train the first iteration for 8,000 steps and the remaining iterations for 4,000 steps with batch size 128. The training steps for PointPillars are reduced by half in each iteration with batch size 64. We perform the search for 25 iterations on StarNet and for 20 iterations on PointPillars.

Even though StarNet and PointPillars are two distinct types of detection models, we have observed similar patterns in the evolution of their augmentation schedules. For StarNet and PointPillars, the probability of FrustumDropout is reduced from 100% to 23% and 56%, and the maximum rotation angle in RandomRotation is reduced from 0.785 to 0.54 and 0.42. These examples indicate that applying weaker data augmentation towards the end of training is beneficial.

### 4.3 Better Results with Less Computation

Above, we have verified the effectiveness of PPBA on improving 3D object detection on the KITTI Dataset and the Waymo Open Dataset. In this section, we analyze the computational cost of PPBA, and compare PPBA with random search and PBA [11] on the KITTI *test* split.

All searches are performed with StarNet [22] and the search space described in Section 3.1. For Random Search<sup>5</sup>, 1,000 distinct augmentation policies are randomly sampled and trained. PBA is run with 16 total trials while training

<sup>4</sup> <http://github.com/tensorflow/lingvo>

<sup>5</sup> Our initial experiment on random search shows the performance distribution of augmentation policies is spread on the KITTI *validation* split. In order to save computation resources, the random search here is performed on a fine-grained search space.

the first iteration for 3,000 steps and the remaining iterations for 1,000 steps with batch size 64.

The baseline StarNet is trained for 8 hours with a TPU v3-32 Pod [13, 15] on vehicle detection and pedestrian/cyclist detection models. Random search requires about  $1,000 \times 8 = 8,000$  TPU hours for training. In comparison, both PBA and PPBA train with a much smaller cost of  $8 \times 16 = 128$  TPU hours, with an additional real-time computation overhead of waiting for the evaluation result for  $8 \times 16 = 128$  TPU hours. We observe that PPBA results in a more than 30x speedup compared to random search, while identifying better-performing augmentation strategies. Furthermore, PPBA outperforms PBA by a substantial margin with the same computational budget.

Table 4: Comparison of 3D mAP on StarNet on the KITTI *test* set across data augmentation methods

Method	TPU Hours	Car			Pedestrian			Cyclist		
		Easy	Mod.	Hard	Easy	Mod.	Hard	Easy	Mod.	Hard
Manual design [22]	8	81.63	73.99	67.07	48.58	41.25	39.66	73.14	58.29	52.58
Random Search	8,000	81.89	74.94	67.39	<b>52.78</b>	<b>44.71</b>	41.12	73.71	59.92	54.09
PBA	256	83.16	75.02	69.72	41.28	34.48	32.24	76.8	59.43	52.77
PPBA	256	<b>84.16</b>	<b>77.65</b>	<b>71.21</b>	52.65	44.08	<b>41.54</b>	<b>79.42</b>	<b>61.99</b>	<b>55.34</b>

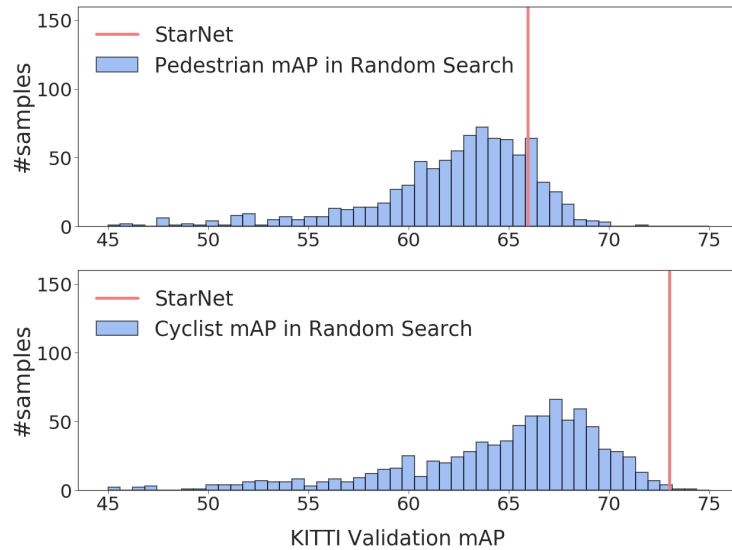


Fig. 4: 3D mAP of a population of 1,000 random augmentation policies for pedestrian and cyclist on the moderate difficulty on the KITTI *validation* split

While searching the augmentation policies randomly for pedestrian/cyclist detection, the majority of samples perform worse than the manual designed augmentation strategy on the KITTI *validation* split (see Fig. 4). Unlike image augmentation search spaces, where each operation has one parameter and even random policies lead to some improvement in generalization, point cloud augmentations are harder to optimize, with a larger number of parameters (e.g. geometric distance, operation strength, distribution of categorical sampling, etc.) and no good priors for the parameters’ ranges. Because of the complex search space, it is challenging to discover good augmentation policies with random search, especially for the cyclist category. We find that it is effective to fine tune the parameter search space of each operation to improve the overall performance of random search. However, the whole process is expensive and requires domain expertise.

We observe that PBA is not effective at discovering better augmentation policies, compared to random search or even to manual search, when the detection category is sensitive to inferior augmentation parameters. In PBA, the full search space is explored at every iteration, which is inefficient for searching parameters in a high dimensional space. To mitigate the inefficiency, PPBA progressively explores a subset of search space at every iteration, and the best parameters discovered in past iterations are adopted in the exploration phase. As in Table 4, PPBA shows much larger improvements on the car and cyclist categories, demonstrating the effectiveness of the proposed strategy.

#### 4.4 Automated Data Augmentation Improves Data Efficiency

In this section, we conduct experiments to determine how PPBA performs when the dataset size grows. The experiments are performed with PointPillars on subsets of the Waymo Open Dataset with the following number of training examples: 10%, 30%, 50%, by randomly sampling run segments and single frames of sensor data, respectively. During training, the decay interval of the learning rate is linearly decreased accordingly to the percentile of data sampled (e.g., reduce the decay interval of learning rate by 50% when sampling 50% of the training examples), while the number of training epochs is set to be inversely proportional to the percentile of data sampled. As it is commonly known that smaller datasets need more regularization, we increase weight decay from  $10^{-4}$  to  $10^{-3}$ , when training on 10% of the examples.

Compared to downsampling from single frames of sensor data, performance degradation of PointPillars models is more severe when downsampling from run segments. This phenomenon is due to the relative lack of diversity in the run segments, which tend to contain the same set of distinct vehicles and pedestrians. As in Table 5, Fig. 5 and Fig. 6, we compare the overall 3D detection mAP on the Waymo Open Dataset *validation* set for all ground truth examples with  $\geq 5$  points and rated as LEVEL.1 difficulty for 3 sets of PointPillars models: with no augmentation, random augmentation policy and PPBA. While random augmentation policy can improve the PointPillars baselines and demonstrate the effectiveness of the proposed search space, PPBA pushes the limit even further.

Table 5: Compare 3D mAP of PointPillars with no augmentation, random augmentation and PPBA on the Waymo Open Dataset *validation* set as the dataset size grows

Method	Sample Unit	10%		30%		50%		100%	
		Car	Pedestrian	Car	Pedestrian	Car	Pedestrian	Car	Pedestrian
Baseline	run segment	42.5	46.1	49.5	56.4	52.5	59.1	57.2	62.3
Random	run segment	49.5	50.6	54.1	58.8	56.1	60.5	60.9	63.5
PPBA	run segment	<b>54.2</b>	<b>55.8</b>	<b>57.6</b>	<b>63.0</b>	<b>58.7</b>	<b>65.1</b>	<b>62.4</b>	<b>66.0</b>
Baseline	single frame	52.4	56.9	55.3	60.7	56.7	61.2	57.2	62.3
Random	single frame	58.3	59.8	59.4	61.9	59.7	62.1	60.9	63.5
PPBA	single frame	<b>59.8</b>	<b>64.2</b>	<b>60.7</b>	<b>65.5</b>	<b>61.2</b>	<b>66.2</b>	<b>62.4</b>	<b>66.0</b>




Fig. 5: Vehicle detection 3D mAP for PointPillars on Waymo Open Dataset *validation* set with no augmentation, random augmentation and PPBA as the dataset size changes




Fig. 6: Pedestrian detection 3D mAP for PointPillars on Waymo Open Dataset *validation* set with no augmentation, random augmentation and PPBA as the dataset size changes

PPBA is 10x more data efficient when sampling from single frames of sensor data, and 3.3x more data efficient when sampling from run segments. As we expected, the improvement from PPBA becomes larger when the dataset size is reduced.

#### 4.5 Progressive Population Based Augmentation Generalizes on Image Classification

So far our experiments have demonstrated that PPBA consistently improves over alternatives for 3D object detection across datasets and architectures. However, PPBA is a general algorithm, and in this section we validate its true versatility by applying it to a common 2D image classification problem.

To search for augmentation policies, we use the same reduced subset of the ImageNet training set with 120 classes and 6,000 samples as [4, 20]. During the PPBA search, 16 trials are trained to optimize the Top-1 accuracy on the reduced validation set for 8 iterations while 4 operations are selected for exploration at every iteration. When replaying the learned augmentation schedule on the full training set, the ResNet-50 model is trained for 180 epochs with a batch size of 4096, a weight decay of  $10^{-4}$  and a cosine decay learning rate schedule with learning rate of 1.6. The results on the ImageNet validation set, shown in Table 6, confirm that PPBA can be used as a highly efficient auto augmentation algorithm for tasks other than 3D object detection.

Table 6: Comparison of Top-1 accuracy (%) and computational cost across augmentation methods on the ImageNet validation set for ResNet-50. Note that the baseline results with Inception-style Pre-processing is reproduced by [4]

Method	Accuracy	GPU Hours	Hardware
Inception-style Pre-processing [28]	76.3	-	-
AutoAugment [4]	77.6	15000	GPU P100
Fast AutoAugment [20]	77.6	450	GPU V100
PPBA	77.5	16	GPU V100

## 5 Conclusion

We have presented Progressive Population Based Augmentation, a novel automated augmentation algorithm for point clouds. PPBA optimizes the augmentation schedule via narrowing down the search space and adopting the best parameters from past iterations. Compared with random search and PBA, PPBA can more effectively and more efficiently discover good augmentation policies in a rich search space for 3D object detection. Experimental results on the KITTI dataset and the Waymo Open Dataset demonstrate that the proposed method can significantly improve 3D object detection in terms of performance and data efficiency. While we have also validated the effectiveness of PPBA on a common task such as image classification, exploring the potential applications of the algorithm to more tasks and models remains an exciting direction of future work.

## References

1. Chen, X., Ma, H., Wan, J., Li, B., Xia, T.: Multi-view 3d object detection network for autonomous driving. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1907–1915 (2017)
2. Cho, H., Seo, Y.W., Kumar, B.V., Rajkumar, R.R.: A multi-sensor fusion system for moving object detection and tracking in urban driving environments. In: 2014 IEEE International Conference on Robotics and Automation (ICRA). pp. 1836–1843. IEEE (2014)
3. Ciregan, D., Meier, U., Schmidhuber, J.: Multi-column deep neural networks for image classification. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition. pp. 3642–3649. IEEE (2012)
4. Cubuk, E.D., Zoph, B., Mane, D., Vasudevan, V., Le, Q.V.: Autoaugment: Learning augmentation policies from data. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2019)
5. Cubuk, E.D., Zoph, B., Shlens, J., Le, Q.V.: Randaugment: Practical data augmentation with no separate search. arXiv preprint arXiv:1909.13719 (2019)
6. DeVries, T., Taylor, G.W.: Improved regularization of convolutional neural networks with cutout. arXiv preprint arXiv:1708.04552 (2017)
7. Dwibedi, D., Misra, I., Hebert, M.: Cut, paste and learn: Surprisingly easy synthesis for instance detection. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1301–1310 (2017)
8. Fang, H.S., Sun, J., Wang, R., Gou, M., Li, Y.L., Lu, C.: Instaboost: Boosting instance segmentation via probability map guided copy-pasting. In: The IEEE International Conference on Computer Vision (2019)
9. Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Are we ready for autonomous driving? the kitti vision benchmark suite. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2012)
10. Girshick, R., Radosavovic, I., Gkioxari, G., Dollár, P., He, K.: Detectron (2018)
11. Ho, D., Liang, E., Stoica, I., Abbeel, P., Chen, X.: Population based augmentation: Efficient learning of augmentation policy schedules. In: International Conference on Machine Learning. pp. 2731–2741 (2019)
12. Jaderberg, M., Dalibard, V., Osindero, S., Czarnnecki, W.M., Donahue, J., Razavi, A., Vinyals, O., Green, T., Dunning, I., Simonyan, K., Fernando, C., Kavukcuoglu, K.: Population based training of neural networks. arXiv preprint arXiv:1711.09846 (2017)
13. Jouppi, N.P., Young, C., Patil, N., Patterson, D., Agrawal, G., Bajwa, R., Bates, S., Bhatia, S., Boden, N., Borchers, A., et al.: In-datacenter performance analysis of a tensor processing unit. In: Proceedings of the 44th Annual International Symposium on Computer Architecture. pp. 1–12 (2017)
14. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: Advances in Neural Information Processing Systems (2012)
15. Kumar, S., Bitorff, V., Chen, D., Chou, C., Hechtman, B., Lee, H., Kumar, N., Mattson, P., Wang, S., Wang, T., Xu, Y., Zhou, Z.: Scale mlperf-0.6 models on google tpu-v3 pods. arXiv preprint arXiv:1909.09756 (2019)
16. Lang, A.H., Vora, S., Caesar, H., Zhou, L., Yang, J., Beijbom, O.: Pointpillars: Fast encoders for object detection from point clouds. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 12697–12705 (2019)

17. Lemley, J., Bazrafkan, S., Corcoran, P.: Smart augmentation learning an optimal data augmentation strategy. *IEEE Access* **5**, 5858–5869 (2017)
18. Li, R., Li, X., Heng, P.A., Fu, C.W.: Pointaugmt: an auto-augmentation framework for point cloud classification. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 6378–6387 (2020)
19. Liang, M., Yang, B., Wang, S., Urtasun, R.: Deep continuous fusion for multi-sensor 3d object detection. In: *Proceedings of the European Conference on Computer Vision*. pp. 641–656 (2018)
20. Lim, S., Kim, I., Kim, T., Kim, C., Kim, S.: Fast autoaugment. In: *Advances in Neural Information Processing Systems* (2019)
21. Luo, W., Yang, B., Urtasun, R.: Fast and furious: Real time end-to-end 3d detection, tracking and motion forecasting with a single convolutional net. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 3569–3577 (2018)
22. Ngiam, J., Caine, B., Han, W., Yang, B., Chai, Y., Sun, P., Zhou, Y., Yi, X., Alsharif, O., Nguyen, P., et al.: Starnet: Targeted computation for object detection in point clouds. *arXiv preprint arXiv:1908.11069* (2019)
23. Ratner, A.J., Ehrenberg, H., Hussain, Z., Dunnmon, J., Ré, C.: Learning to compose domain-specific transformations for data augmentation. In: *Advances in Neural Information Processing Systems*. pp. 3239–3249 (2017)
24. Sato, I., Nishimura, H., Yokoi, K.: Apac: Augmented pattern classification with neural networks. *arXiv preprint arXiv:1505.03229* (2015)
25. Shi, S., Guo, C., Jiang, L., Wang, Z., Shi, J., Wang, X., Li, H.: Pv-rcnn: Point-voxel feature set abstraction for 3d object detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10529–10538 (2020)
26. Simard, P.Y., Steinkraus, D., Platt, J.C., et al.: Best practices for convolutional neural networks applied to visual document analysis. In: *Proceedings of International Conference on Document Analysis and Recognition* (2003)
27. Sun, P., Kretschmar, H., Dotiwala, X., Chouard, A., Patnaik, V., Tsui, P., Guo, J., Zhou, Y., Chai, Y., Caine, B., Vasudevan, V., Han, W., Ngiam, J., Zhao, H., Timofeev, A., Ettinger, S., Krivokon, M., Gao, A., Joshi, A., Zhang, Y., Shlens, J., Chen, Z., Anguelov, D.: Scalability in perception for autonomous driving: Waymo open dataset. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* pp. 2446–2454 (2020)
28. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A., et al.: Going deeper with convolutions. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 1–9 (2015)
29. Thrun, S., Montemerlo, M., Dahlkamp, H., Stavens, D., Aron, A., Diebel, J., Fong, P., Gale, J., Halpenny, M., Hoffmann, G., et al.: Stanley: The robot that won the darpa grand challenge. *Journal of field Robotics* **23**(9), 661–692 (2006)
30. Wan, L., Zeiler, M., Zhang, S., Le Cun, Y., Fergus, R.: Regularization of neural networks using dropconnect. In: *International Conference on Machine Learning*. pp. 1058–1066 (2013)
31. Yan, Y., Mao, Y., Li, B.: Second: Sparsely embedded convolutional detection. *Sensors* **18**(10), 3337 (2018)
32. Yang, B., Liang, M., Urtasun, R.: Hdnet: Exploiting HD maps for 3d object detection. In: *Proceedings of The 2nd Conference on Robot Learning*. pp. 146–155 (2018)



33. Yang, B., Luo, W., Urtasun, R.: Pixor: Real-time 3d object detection from point clouds. In: Proceedings of the IEEE conference on Computer Vision and Pattern Recognition. pp. 7652–7660 (2018)
34. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. In: International Conference on Learning Representations (2018)
35. Zhou, D., Fang, J., Song, X., Guan, C., Yin, J., Dai, Y., Yang, R.: Iou loss for 2d/3d object detection. In: International Conference on 3D Vision (3DV). IEEE (2019)
36. Zhou, Y., Sun, P., Zhang, Y., Anguelov, D., Gao, J., Ouyang, T., Guo, J., Ngiam, J., Vasudevan, V.: End-to-end multi-view fusion for 3d object detection in lidar point clouds. In: Proceedings of the Conference on Robot Learning (2019)
37. Zhou, Y., Tuzel, O.: Voxelnet: End-to-end learning for point cloud based 3d object detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4490–4499 (2018)
38. Zoph, B., Cubuk, E.D., Ghiasi, G., Lin, T.Y., Shlens, J., Le, Q.V.: Learning data augmentation strategies for object detection. arXiv preprint arXiv:1906.11172 (2019)