

Towards Efficient 3D Object Detection with Knowledge Distillation

Jihan Yang¹ Shaoshuai Shi² Runyu Ding¹ Zhe Wang³ Xiaojuan Qi¹

¹The University of Hong Kong ²Max Planck Institute for Informatics ³SenseTime Research
{jhyang, ryding, xjq}@eee.hku.hk, {shaoshuaics, wzlewis16}@gmail.com

Abstract

Despite substantial progress in 3D object detection, advanced 3D detectors often suffer from heavy computation overheads. To this end, we explore the potential of knowledge distillation (KD) for developing efficient 3D object detectors, focusing on popular pillar- and voxel-based detectors. In the absence of well-developed teacher-student pairs, we first study how to obtain student models with good trade offs between accuracy and efficiency from the perspectives of model compression and input resolution reduction. Then, we build a benchmark to assess existing KD methods developed in the 2D domain for 3D object detection upon six well-constructed teacher-student pairs. Further, we propose an improved KD pipeline incorporating an enhanced logit KD method that performs KD on only a few pivotal positions determined by teacher classification response, and a teacher-guided student model initialization to facilitate transferring teacher model’s feature extraction ability to students through weight inheritance. Finally, we conduct extensive experiments on the Waymo and KITTI dataset. Our best performing model achieves 65.75% LEVEL 2 mAPH, surpassing its teacher model and requiring only 44% of teacher flops on Waymo. Our most efficient model runs 51 FPS on an NVIDIA A100, which is $2.2\times$ faster than PointPillar with even higher accuracy on Waymo. Code is available at <https://github.com/CVMI-Lab/SparseKD>.

1 Introduction

3D object detection from point clouds is a fundamental perception task with broad applications on autonomous driving, robotics and smart city, etc. Recently, benefited from large-scale 3D perception datasets [9, 2, 37] and advanced point- [29], pillar- [21, 43] and voxel-based [10, 51] representations of sparse and irregular LiDAR point cloud scenes, 3D detection has achieved remarkable progress [46, 35, 33, 1, 49]. However, stronger performance is often accompanied with heavier computation burden (see Figure 1), rendering their adoption in real-world applications still a challenging problem.

Recent attempts to improve efficiency focus on developing specified architectures for point-based 3D object detectors [5, 50], not generalizable to a wide spectrum of pillar/voxel-based methods [51, 21, 46, 33, 49, 7]. Here, we aim at a model-agnostic framework for obtaining efficient and accurate 3D object detectors with knowledge distillation (KD). Due to its effectiveness, generality and simplicity, KD has become a de facto strategy to develop efficient models in a variety of 2D tasks [15, 24, 6, 16, 45]. It facilitates

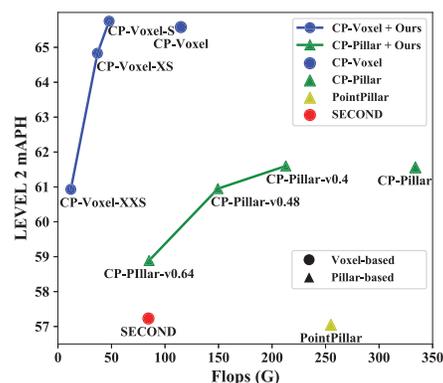


Figure 1: Performance and flops comparison of single-stage detectors on Waymo: CP-Pillar [49], CP-Voxel [49], SECOND [46], PointPillar [21] and Ours.

improving the performance of a lightweight and efficient student model by harvesting knowledge learned by an accurate yet computationally heavy teacher model. Despite of numerous studies in 2D tasks, the investigation of KD for efficient 3D object detection has largely escaped research attention with unresolved research challenges. In this paper, we conduct the first systematic study on knowledge distillation for developing high-performance and efficient 3D LiDAR-based detectors.

First, we study *how to obtain lightweight student detectors with satisfactory efficiency and accuracy trade offs given a pre-trained teacher 3D object detector*. Different from the 2D domain where well-developed backbone architectures with different model efficiencies are readily available (e.g., ResNet 18 vs. ResNet 50) [36, 39, 13], such scalable backbones in 3D scenes are still under-explored. This makes the design of suitable student models a non-trivial problem. Intuitively, a good student model should achieve a good compromise between accuracy and efficiency as a poor student model may have inferior capabilities or architectural level shortcomings, causing difficulties for further knowledge distillation. In this perspective, we first propose the Cost Performance Ratio (CPR) to fairly evaluate a model in terms of both efficiency and capability. Then, we study different factors including model width (*i.e.* number of channels), depth (*i.e.* number of layers) and input resolution on both pillar-based and voxel-based architectures. Specifically, we find that pillar-based architecture favors input-level compression (*i.e.* reduce input resolution) while voxel-based detector prefers width-level compression, due to less spatial redundancy of voxel-based detectors.

Second, we empirically investigate *the effectiveness of existing knowledge distillation methods on this new setting* upon accurate teacher models and efficient student models. As there is no prior research for this problem, we benchmark seven existing knowledge distillation methods on top of six teacher-student pairs covering voxel- and pillar-based architectures. Specifically, we evaluate the following major streams: **logit KD** distilling on model outputs (KD [15] and GID-L [6]), **feature KD** mimicking intermediate features (FitNet [32], Mimic [22], FG [42] and GID-F [6]) and **label KD** leveraging teacher's predictions for label assignment [28]. Further, we also study their synergy effects and empirically find that feature KD outperforms others individually, but fails to cooperate with other KD manners to further boost the performance in 3D detection. This is potentially caused by the fact that logit and label KD can add an implicit regularization to the intermediate features. The best strategy that we obtain through empirical studies is to use FG [42] or Mimic [22] individually.

Third, we propose *simple, general, and effective strategies to improve knowledge distillation on 3D object detection* upon the strong KD baseline derived as above. Motivated by the extreme imbalance between small informative areas containing 3D objects and large redundant background areas in 3D scenes, we design a modified logit KD method, namely pivotal position logit KD, enforcing imitation on only locations with highly confident or top-ranked teacher predictions. These areas are shown to be near instance centers or error-prone positions. Besides, to facilitate effective knowledge transfer from teacher to student, we develop Teacher Guided Initialization (TGI) which remaps pre-trained teacher parameters to initialize a student model. This is shown to be effective in inheriting teacher models' feature extraction abilities while collaborating well with logit and label KD techniques.

Finally, our empirical studies on efficient model design and knowledge distillation methods yield superior performance in delivering efficient and effective pillar- and voxel-based 3D detectors. This is extensively verified on the largest annotated 3D object detection dataset – Waymo [37]. As shown in Figure 1, our best performing model, *i.e.* CP-Voxel-S, even outperforms its teacher model (*i.e.* CP-Voxel) while has $2.4\times$ fewer flops. Moreover, our most efficient pillar-based model (*i.e.* CP-Pillar-v0.64) can run 51 FPS with 58.89% mAPH while previous fastest voxel/pillar-based detector – PointPillar runs 23 FPS with 57.03% mAPH on an NVIDIA A100 GPU (see Suppl.). Our method is also shown to be effective in knowledge transfer from heavy two-stage object detectors to lightweight single-stage detectors. In addition, our method can generalize well to other settings such as KITTI dataset with SECOND as well as advance compression methods in Sec. 5.5, other detectors, and even 3D semantic segmentation in Suppl..

2 Related Work

3D LiDAR-based Object Detection targets to localize and classify 3D objects from point clouds. Point-based methods [35, 5, 47, 50] took raw point clouds and leveraged PointNet++ [29] to extract sparse point features and generate point-wise 3D proposals. Pillar-based works [21, 43] finished voxelization in bird eye's view and extracted pillar-wise features with PointNet++. Voxel-based methods [51, 46, 33] voxelized point clouds and obtained voxel-wise features with 3D sparse convolutional networks, which is the most popular data treatment. Besides, range-based works [1, 38]

were proposed for long-range and fast detection. Recently, designing efficient 3D detectors has drawn some attentions [5, 50] with raw point data treatment. In this work, we focus on exploring model-agnostic knowledge distillation methods to boost the efficiency of 3D detectors.

Knowledge Distillation aims to transfer knowledge from a large teacher model to a lightweight student network, which is a thriving area in efficient deep learning. Hinton *et al.* proposed the seminal concept of knowledge distillation (KD) [15], which distilled knowledge between teacher and student on the output level (*i.e.* prediction logits). Another line of research proposed to help student's optimization with hints stored in informative intermediate features from teacher [32, 17, 20, 14, 19, 4]. In addition, some works attempted distillation techniques in 2D object detection [22, 42, 6, 30, 48, 28] by emphasizing instance-wise distillation and feature knowledge. Mimic [22], FG [42] and GID [6] sampled local region features with box proposals or custom indicators for foreground-aware feature imitation. Label KD [28] utilized teacher's information for label assignment of student. Recently, knowledge distillation has also been leveraged to transfer knowledge in multi-modality setup [11, 25] or multi-frame to single-frame setup [44] in 3D detection area. However, to the best of our knowledge, we are the first to explore knowledge distillation in the most popular setup: single-frame 3D LiDAR-based object detection. In this work, we propose an enhanced 3D detection KD pipeline with our designed efficient 3D detectors on the popular voxel/pillar data representation. Our lightweight detectors CP-Voxel-S and CP-Pillar-v0.4 slightly outperform their state-of-the-art teacher detectors separately while requiring much less computation overhead.

3 Designing Efficient Student Networks

As there are no readily available lightweight backbone architectures for constructing student networks, we carry out an empirical study on how to obtain an efficient model with satisfactory efficiency and accuracy trade offs to facilitate further knowledge distillation. In this section, we will first describe our experimental setups and model evaluation metrics. Then, we study different strategies to obtain an efficient model and conduct an in-depth analysis on how to achieve good trade offs for pillar- and voxel-based architectures.

3.1 Basic Setups and Evaluation Metrics

Basic setups. For detector architectures, we focus on two variants of the state-of-the-art model CenterPoint [49]: CenterPoint-Pillar (CP-Pillar) and CenterPoint-Voxel (CP-Voxel), covering the most popular pillar- and voxel-based 3D detectors [46, 51, 21, 43, 33, 34, 7]. For dataset, we perform all experiments on the largest annotated 3D LiDAR perception dataset Waymo Open Dataset (WOD) [37] with 20% training samples for fast verification. For model training, we follow the training scheme of popular 3D detection codebase OpenPCDet [41] to ensure fair comparisons and standardization. Note that there is no any knowledge distillation method engaged in this stage.

Evaluation metrics. Following [31, 23], we employ number of parameters, flops, activations, latency (*i.e.* test time) and peak GPU training memory (batch size 1) as quantitative indicators to evaluate model efficiency from parameter, computation and memory throughout aspects. Note that activations, flops and latency are averaged over 99 frames with a GTX-1060 GPU. We present latency more for reference since it largely depends on hardware devices as well as operation-level optimizations (see Sec. ??). We use LEVEL 2 mAPH as the performance evaluation metric following WOD [37].

Since our major target here is to design student networks with favorable trade offs between performance and efficiency, we propose a quantitative indicator, namely Cost Performance Ratio (CPR), to directly measure a model in this respective. To construct CPR, we use activations (acts) as the metric to evaluate efficiency since it strongly correlates with the runtime on hardware accelerators such as GPUs as shown in [31]. Our CPR finally combines the activation ratio and mAPH as follows:

$$\text{CPR} = 0.5 \times \left(1 - \frac{\text{acts}_s}{\text{acts}_t}\right) + 0.5 \times \left(\frac{\text{mAPH}_s}{\text{mAPH}_t}\right)^3, \quad (1)$$

where the subscripts s and t represent student and teacher model respectively. CPR is normalized to $[0, 1]$ by weighting the relative activation decrease and performance drop ratio of a student network compared to the teacher. Notice that the third power is used for the performance degradation term to penalize acceleration methods that result in drastic performance degradation more severely. We argue that it is necessary to ensure a relative good performance of efficient detectors as models with poor accuracy may suffer from architectural level problems which can cause difficulties in developing knowledge distillation techniques and prohibiting obtaining accurate and efficient detectors.

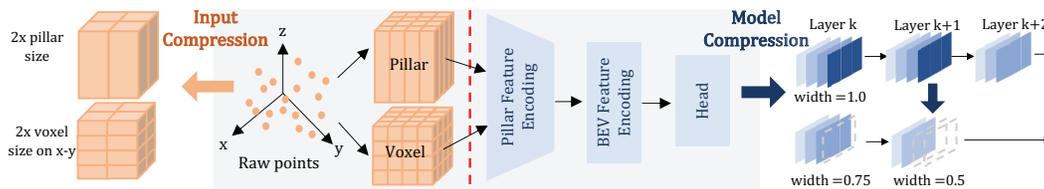


Figure 2: Architecture of detector and illustration of model (right) and input (left) compression.

3.2 Acceleration Strategy

Here, to obtain efficient models, we investigate model and input resolution compression techniques on pillar- and voxel-based 3D detectors as shown in Figure 2 and study their impacts on model accuracy.

Model Compression. Here, we compress a given teacher backbone by trimming it along depth (*i.e.* number of layers) or width (number of channels) as shown in the right of Figure 2. To relieve the burden of exhaustive layer-wise studies, we leverage the functional similarity among different layers to group them into three major modules: Pillar Feature Encoding (PFE) module, Bird eye’s view Feature Encoding (BFE) module and detection head (see Figure 2: right), and conduct analysis on module-level. Specifically, PFE corresponds to network components before projecting features to bird eye’s view (BEV) grid, including sparse 3D convolutional backbone [10] to extract voxel-wise features in voxel-based architectures or PointNet++ [29] to encode pillar-wise features from points in pillar-based detectors. After PFE, the features are aggregated to pillar features and mapped to 2D BEV grid. Then, the BFE consisting of 2D convolutional layers is adopted to extract final detection feature map on the BEV grid. Finally, the detection head takes outputs from BFE to produce final prediction results. For width pruning, we consider the width of teacher detectors as 1.0, and slim each module of the student by decreasing its number of channels with a given width. Depth trimming also follows this paradigm but still maintains the minimal structure of each module (*e.g.* downsample and upsample layers) to ensure basic detection ability. Model trimming results are shown in Table 1.

Table 1: Model compression results. Teacher models are marked in gray. See text for details.

Detector	Architecture					Efficiency					LEVEL 2 mAPH	CPR
	Width			Depth		Params (M)	Flops (G)	Acts (M)	Latency (ms)	Mem. (G)		
CP-Pillar	1.00	1.00	1.00	1.00	1.00	5.2	333.9	303.0	157.9	5.2	59.09	-
	1.00	0.50	0.50	1.00	1.00	1.3	87.6	161.8	78.5	3.2	54.50	0.63
	0.50	0.50	0.50	1.00	1.00	1.3	85.0	152.7	74.0	2.9	52.33	0.60
	1.00	1.00	1.00	1.00	0.50	2.2	258.5	234.1	118.5	4.3	55.24	0.52
	1.00	1.00	1.00	1.00	0.33	1.4	234.6	210.0	107.8	4.0	47.97	0.42
CP-Voxel	1.00	1.00	1.00	1.00	1.00	7.8	114.7	101.9	125.7	2.8	64.29	-
	1.00	0.50	0.50	1.00	1.00	4.0	47.8	65.7	98.0	2.1	62.23	0.63
	0.75	0.50	0.50	1.00	1.00	2.8	36.9	58.4	88.2	1.9	61.16	0.64
	0.50	0.50	0.50	1.00	1.00	1.9	28.8	51.2	75.1	1.7	59.47	0.64
	0.50	0.25	0.25	1.00	1.00	1.0	12.0	33.1	70.4	1.3	56.26	0.67
	0.25	0.25	0.25	1.00	1.00	0.5	7.3	25.8	66.0	1.1	49.84	0.61
	1.00	1.00	1.00	0.50	0.50	3.0	63.9	65.2	73.0	1.9	60.95	0.61
	1.00	1.00	1.00	0.33	0.33	1.8	47.9	52.2	59.0	1.6	55.78	0.57

Input Compression. Besides model complexities, input resolution also has impacts on model efficiency [40, 30]. For instance, by halving the input resolution, the computation overhead for 2D convolution layers in the BFE module and detection head will be reduced to $\frac{1}{4}$ (see Figure 2: left). Besides, there are large background areas in the sparse and large-scale 3D scenes, which naturally has redundancies and offers the possibility of processing the data on a coarser resolution for input compression. Specifically, input compression is realized by increasing the voxel/pillar size on the x-y plane when constructing voxel/pillar. As shown in Table 2, we gradually increase students’ voxel size with 25% of teachers’ voxel size, and record their efficiency and accuracy metrics.

3.3 Conclusion and Analysis

By analyzing the experimental results shown in Table 1 and Table 2, we draw the following conclusions on efficient model design for pillar- and voxel-based detectors.

Width vs. depth compression: width-level pruning is preferred. As shown in Table 1, trimming networks on width generally achieves higher CPR than on depth for both CP-Pillar and CP-Voxel. For instance, with stronger performance, CP-Voxel (d) needs $1.5\times$ fewer acts and $4\times$ fewer flops compared to CP-Voxel (g). As backbones of 3D detectors are much shallower than their 2D counterparts

Table 2: Input compression results. Teacher models are marked in gray. See text for details.

Architecture		Efficiency					LEVEL 2	CPR
Detector	Voxel Size (m)	Params (M)	Flops (G)	Acts (M)	Latency (ms)	Mem. (G)	mAPH	
CP-Pillar	0.32	5.2	333.9	303.0	157.9	5.2	59.09	-
	0.40	5.2	212.9	197.7	103.4	3.8	57.55	0.64
	0.48	5.2	149.4	142.3	81.9	3.0	56.27	0.70
	0.56	5.2	109.9	109.0	66.3	2.6	54.45	0.71
	0.64	5.2	85.1	88.0	54.5	2.1	52.81	0.71
CP-Voxel	0.100	7.8	114.8	101.9	125.7	2.8	64.29	-
	0.125	7.8	77.5	70.1	99.9	2.2	61.55	0.59
	0.150	7.8	53.9	50.0	84.3	1.8	58.14	0.62
	0.175	7.8	44.3	41.2	74.1	1.5	55.99	0.63
	0.200	7.8	32.9	31.4	67.5	1.3	52.80	0.62

(e.g. only 19 convolution layers in CP-Pillar and 35 convolution layers in CP-Voxel), this renders depth compression more challenging and less scalable than width compression in 3D detection.

Module-wise pruning selection: PFE module has the least redundancy to be reduced. Since PFE, BFE and detection head perform different functions, they might have their own redundancies in 3D detection. Comparing CP-Voxel (c), (d) and (e) as well as CP-Pillar (a) and (b) in Table 1, we find that less pruning on PFE achieves significantly higher CPR, demonstrating that network parameters in the PFE module is crucial for high performance and it is necessary to maintain network complexities in PFE modules.

Favorable compression strategies for different detection architectures. Comparing results of different compression strategies in Table 1 and 2, we find that pillar-based architecture (*i.e.* CP-Pillar) is more suitable for input compression while voxel-based architecture (*i.e.* CP-Voxel) prefers width-level compression. This difference mainly lies in different spatial redundancy of BEV features on these two detectors. Since the PFE module of voxel-based detector downsamples input resolution to $1/8$, the $0.1m$ input voxel size will be magnified to $0.8m$ on BEV features. Hence, less spatial redundancy could be further compressed for CP-Voxel. However, CP-Pillar often keeps the same resolution between input and BEV features on Waymo [41, 49], and thus has larger resolution redundancy on BEV grid, which facilitates designing student network with coarser input resolution.

Summarized student networks. Driven by the above analysis, we finally derive the following compressed student models with good trade offs between efficiency and performance. For CP-Pillar, as it is friendly to input compression, we adopt the input-compressed models with voxel size 0.40, 0.48 and 0.64 in Table 2, named **CP-Pillar-v0.4**, **CP-Pillar-v0.48** and **CP-Pillar-v0.64**, respectively. As for CP-Voxel, we choose its width compressed models (a) (b) and (d) in Table 1, named **CP-Voxel-S**, **CP-Voxel-XS** and **CP-Voxel-XXS**, respectively.

4 Benchmark Knowledge Distillation for 3D Object Detection

With student networks derived in Section 3, we are now ready to conduct our empirical study on knowledge distillation for 3D object detection. Here, we benchmark seven popular 2D KD methods including logit KD (*i.e.* KD [15] and GID-L [6]), feature KD (*i.e.* FitNet [32], Mimic [22], FG [42] and GID-F [6]) and label KD [28] on six teacher-student pairs with comprehensive analysis. Notice that GID and label KD are state-of-the-art 2D KD detection methods, and GID is divided into GID-L and GID-F to investigate logit KD and feature KD separately. Implementation details and the value of hyper-parameters are described in the Suppl..

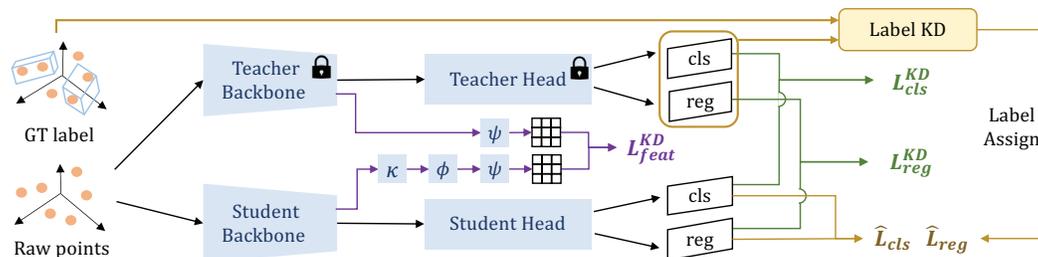


Figure 3: Overall KD paradigm for 3D detection. Teacher weights are frozen during the whole distillation procedure. Logit, feature and label KD are colored by green, purple and yellow, respectively.

4.1 Paradigm

We first evaluate existing 2D distillation methods for 3D detection. As shown in Figure 3, our overall KD paradigm for 3D detection contains three parts: logit, feature and label KD to leverage teacher guidance in response, intermediate feature and label assignment levels, respectively.

Logit KD is the most classical distillation approach introduced by [15]. It takes teacher model’s final response as guidance for training a student network and is closely related to the specific task. In 3D object detection, we calculate the logit KD loss between teacher and student outputs as follows:

$$\mathcal{L}_{\text{cls}}^{\text{KD}} = \mathbb{E}[m_{\text{cls}} \|\kappa(p_{\text{cls}}^s) - p_{\text{cls}}^t\|_2], \quad \mathcal{L}_{\text{reg}}^{\text{KD}} = \mathcal{L}_{\text{reg}}(p_{\text{reg}}^s, p_{\text{reg}}^t), \quad (2)$$

where superscripts s and t indicate student and teacher, p_{cls} and p_{reg} represent the classification response after the sigmoid and bounding box regression prediction of detector separately, κ is the bilinear interpolation to match student output resolutions towards teacher, \mathcal{L}_{reg} is the regression loss function of 3D detector and m_{cls} is a mask ranged in $[0, 1]$ to indicate important regions in p_{cls} (see green parts in Figure 3). Compared to vanilla KD [15] which mimics all teacher outputs, GID-L only focuses on some local regions covered by selected box proposals and results in different m_{cls} .

Feature KD is the major stream of work in KD for 2D object detection. It enforces student models to mimic teacher models’ intermediate feature maps. Specifically, we construct feature mimicking on the last layer of BFE between student and teacher network as follows:

$$\mathcal{L}_{\text{feat}}^{\text{KD}} = \mathbb{E}[m_{\text{feat}} \|\psi(\phi(\kappa(f^s)), y) - \psi(f^t, y)\|_2], \quad (3)$$

where y is the ground truth, ψ indicates the RoI Align [12], ϕ is a 1×1 convolution with batch normalization [18] and ReLU [27] block to align channel-wise discrepancy between teacher feature f^t and student feature f^s , m_{feat} is the mask to indicate critical regions ranged in $[0, 1]$ (see purple parts in Figure 3). Considering the imbalance between foreground and background regions in the detection, a prevailing solution is to emphasize near object regions to perform distillation [22, 42, 6]. Different investigated feature KD methods mainly vary in the selection of such critical regions m_{feat} .

Label KD is a newly proposed distillation strategy, which leverages teacher predictions in the label assignment stage of student [28]. Motivated by the simple and general label KD, we also employ it as a KD baseline method. Specifically, given a point cloud scene x and its corresponding ground truth (GT) set y , label KD first obtains the prediction y^t and confidence score o^t from the pretrained teacher detector. After filtering o^t with a given threshold τ , it then obtains a high-quality teacher prediction set \hat{y}^t . It generates a teacher assisted GT set $\hat{y}^{\text{KD}} = \{y, \hat{y}^t\}$ by combining GTs as well as confident teacher predictions and carries on label assignment for student with \hat{y}^{KD} . The final classification and regression losses with label KD on student detectors are $\hat{\mathcal{L}}_{\text{cls}}$ and $\hat{\mathcal{L}}_{\text{reg}}$.

Training Objective is the combination of three stream KD techniques as follows:

$$\mathcal{L} = \hat{\mathcal{L}}_{\text{cls}} + \lambda \hat{\mathcal{L}}_{\text{reg}} + \alpha_1 \mathcal{L}_{\text{cls}}^{\text{KD}} + \alpha_2 \mathcal{L}_{\text{reg}}^{\text{KD}} + \alpha_3 \mathcal{L}_{\text{feat}}^{\text{KD}}, \quad (4)$$

where λ , α_1 , α_2 and α_3 are trade-off parameters between different objectives. Note that the existence and implementation of each operation (e.g. ψ , κ , m_{feat} , m_{cls} , etc) varies among different KD methods.

Table 3: Knowledge distillation benchmark for 3D detection on Waymo. Performance are measured in LEVEL 2 mAPH. Best and second-best methods are noted by **bold** and underline, respectively. “Ours” indicates our proposed improved knowledge distillation method introduced in Sec. 5

Detector	No Distill	Logit KD		Feature KD				Label KD	Ours	Flops (G)	Acts (M)
		KD	GID-L	FitNet	Mimic	FG	GID-F				
CP-Pillar	59.09	-	-	-	-	-	-	-	-	333.9	303.0
CP-Pillar-v0.4	57.55	57.51	57.54	57.89	<u>58.57</u>	58.44	58.26	58.10	59.24	212.9	197.7
CP-Pillar-v0.48	56.27	55.76	56.29	55.82	57.26	57.26	57.23	<u>57.54</u>	58.53	149.4	142.3
CP-Pillar-v0.64	52.81	53.13	50.78	51.79	<u>53.83</u>	53.37	53.18	<u>53.78</u>	55.82	85.1	88.0
CP-Voxel	64.29	-	-	-	-	-	-	-	-	114.7	101.9
CP-Voxel-S	62.23	62.81	62.89	60.51	<u>63.35</u>	63.33	62.75	63.31	64.25	47.8	65.7
CP-Voxel-XS	61.16	61.30	62.25	58.94	62.23	<u>62.48</u>	62.42	61.81	63.53	36.9	58.4
CP-Voxel-XXS	56.26	56.11	57.19	52.24	57.00	<u>57.92</u>	57.16	57.02	59.28	12.0	33.1

4.2 Results and Analysis

Benchmark Analysis. As shown in Table 3, compared to the no distillation baseline, all three streams of KD methods obtain performance improvements on six teacher-student pairs. Among seven KD baseline strategies, feature-based KD methods (i.e. Mimic and FG) achieve prominent performance, which demonstrates the strong potential of learning from teacher’s hints on feature

extraction. Furthermore, we find that instance-aware local region imitation is important in distillation for 3D detection, as enormous background regions overwhelm the supervision of sparse instances. For instance, with instance-aware imitation, Mimic, FG and GID-F consistently outperform FitNet which fully imitates all spatial positions of teacher feature maps. Similar conclusions can also be drawn in logit KD by comparing the results of instance-aware GID-L and vanilla KD.

Synergy Analysis. While benchmark results in Table 3 mainly focus on the individual effectiveness of each KD manner, their synergy effect is also an important consideration, which has the potential to further improve student performance. As shown in Table 4, although feature KD itself achieves the highest performance on CP-Voxel-XXS compared to logit KD and label KD techniques, it can hardly achieve improvements or even suffers from performance degradation when combined with other KD methods. On the contrary, logit KD and label KD can cooperate well with each other to further improve student’s capability. This is potentially caused by logit KD and label KD implicitly enforcing regularization on the feature, which can be conflicted with the optimization direction of feature KD and results in a poor synergy effect.

Table 4: Synergy investigation based on CP-Voxel-XXS.

GID-L	Label KD	FG	mAPH
			56.26
✓			57.19
	✓		57.02
		✓	57.92
✓	✓		57.60
✓		✓	58.01
	✓	✓	57.06
✓	✓	✓	57.62

5 Improved Knowledge Distillation for 3D Object Detection

As shown in Table 4, the basic knowledge distillation pipeline fails to achieve remarkable results by combining the best method of three KD streams (*i.e.* GID-L [6], label KD [28] and FG [42]). In the following, we propose an improved KD pipeline for 3D object detection, including pivotal position logit KD to alleviate the extreme imbalance between foreground and background regions through only imitating response on sparse pivotal positions, label KD, and teacher guided initialization scheme to further facilitate transferring teacher’s feature extraction ability to student models.

5.1 Method

Pivotal Position Logit KD. Motivated by the imbalance of foreground and background regions, previous 2D methods attempt to only enforce output-level imitation on pixels near or covering instances [42, 6]. However, we find that it is sub-optimal in 3D scenarios given more extreme imbalance between small informative instances and large redundant background areas. For example, based on CP-Pillar, even a vehicle with 10m length and 4m width occupies only 32×13 pixels in the final 468×468 response map. Such small instances and large perception ranges in 3D detection requires more sophisticated imitation region selection than previous coarse instance-wise masking manners in 2D detection. Hence, we propose Pivotal Position (PP) logit KD which leverages cues in teacher classification response or label assignment to determine the important areas for distillation.

Specifically, pivotal position selection can be formulated as finding suitable m_{cls} in Eq. (2). Here, we show three variants to obtain it. First, confidence of high-performance teacher prediction can serve as a valuable indicator to figure out pivotal positions for student (*i.e.* confidence PP). By filtering teacher confidence o^t with a threshold τ_{pp} , we then set i, j positions of m_{cls} with $o_{i,j}^t \geq \tau_{pp}$ to one and otherwise zero. With similar spirit, we can also select top-ranked K positions (*i.e.* rank PP) in teacher classification response P_{cls}^t as pivotal positions and convert them to a one-hot embedded m_{cls} . These confident or top-ranked positions are shown to be near object centers or error-prone regions. Last, inspired by the Gaussian label assignment in CenterPoint [49], we can define pivotal positions in a soft way with center-peak Gaussian distribution for each instance as m_{cls} (*i.e.* Gaussian PP). We empirically show that all three variants of PP logit KD achieve promising gains in 3D detection.

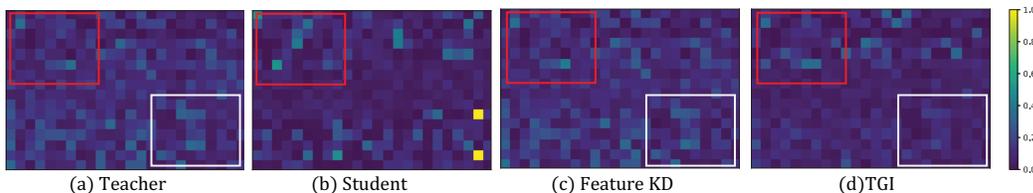


Figure 4: Visualization of channel-wise L_1 norm for distillation from CP-Pillar to CP-Pillar-v0.64.

Teacher Guided Initialization. As mentioned in Section 4.2, feature KD itself outperforms other KD methods, but its deficient synergy results hinder 3D detection KD pipeline from achieving promising performance. This might be caused by the conflict between different optimization directions on

intermediate features from five KD loss terms in Eq. (4). Hence, we explore whether there is a substitute manner which can also leverage teacher’s guidance on the feature extraction aspect. Since the value of feature map is determined by model’s weights, finding hints from teacher weights becomes an alternative to guide student’s feature extraction.

Hence, we propose to directly use the trained weights of teacher to serve as the initialization of student network, named teacher guided initialization (**TGI**) to enhance student model’s feature extraction abilities by inheriting it from a teacher model. Although such initialization is intuitive for input compressed students (e.g. CP-Pillar-v0.64), it cannot be directly applied to pruned students (e.g. CP-Voxel-S). Therefore, we employ the parameter remapping strategy FNA [8] to project teacher weights on student parameters. Take 2D convolution layer as an example, teacher and student parameters are represented with $W^t \in \mathbb{R}^{r \times q \times h \times w}$ and $W^s \in \mathbb{R}^{v \times u \times h \times w}$, where $v \leq r$ and $u \leq q$, respectively. The first v and u channels of teacher parameters are directly assigned to the slimmed student network in FNA. We visualize the channel-wise L_1 norm of backbone feature to compare TGI extracted features with other strategies. As shown in Figure 4, our TGI extracts feature similar to teacher’s (see red box region), but not naively mimics all teacher channels as feature KD does (see white box region). We argue that fully imitating all teacher features might penalize student optimization regarding the architecture or input resolution discrepancy between teacher-student pairs. We empirically show that our simple TGI strategy achieves comparable or even better performance individually among four feature KD methods and shows prominent synergy effects on six teacher-student pairs (See Suppl.).

5.2 Main Results and Comparison

To verify our improved KD pipeline, we conduct experiments on all six teacher-student pairs. As shown in Table 3, our improved KD pipeline consistently surpasses previous KD strategies on all settings with 0.7% ~ 1.9% improvement, thanks to our enhanced logit KD and more collaborative TGI. Equipping our lightweight detectors with the improved KD pipeline, CP-Voxel-S obtains comparable performance to CP-Voxel in terms of mAPH with around $2.4\times$ fewer flops and $1.6\times$ fewer activations. Furthermore, with $3.1\times$ fewer flops, CP-Voxel-XS only suffers 0.76% performance degradation. In addition, with around $1.6\times$ fewer flops and $1.5\times$ fewer activations, our distilled CP-Pillar-v0.4 even slightly outperforms CP-Pillar. CP-Pillar-v0.64 requires only 25% flops and 29% activations of teacher model, while achieving 55.82% mAPH, only 3.27% performance drop compared to CP-Pillar. These experimental results demonstrate that knowledge distillation is a promising technique to improve the performance of efficient 3D detectors.

5.3 Comparison with Other Detectors

To further demonstrate the efficiency and effectiveness of our designed detectors and KD pipeline, we also compare our distilled students with other detectors on the full WOD. As shown in Table 5, our CP-Voxel-S even outperforms its teacher CP-Voxel with around $2.0\times$ fewer parameters, $2.4\times$ fewer flops and $1.6\times$ activations. With similar latency and much fewer parameters, flops as well as activations, CP-Voxel-XS outperforms SECOND by 7.6%. Our CP-Pillar-v0.64, is $2.4\times$ faster than PointPillar on a GTX-1060 while achieves 1.8% higher performance.

5.4 Cross Stage Distillation

A prevailing strategy to improve state-of-the-art 3D object detectors is to adopt a object proposal refinement head for two stage detection [33, 7, 34]. However, despite performance improvements by 3.5%, PVRCNN++ requires around $2.2\times$ parameters, $1.8\times$ activations and $3.5\times$ latency compared to CP-Voxel (see Table 6). Such computation and parameter overheads hinder the real-world applications of state-of-the-art two-stage 3D detectors. Therefore, here we also investigate whether the knowledge of the two-stage detector can help the learning of single-stage detector. It is noteworthy that this is the first attempt at cross stage distillation in both 2D and 3D object detection. As shown in Table 6, leveraging hints

Table 5: Comparison with other detectors on full WOD. † indicates results re-implemented by us.

Detector	Params (M)	Flops (G)	Acts (M)	Latency (ms)	LEVEL 2 mAPH
PointPillar† [21]	4.8	255.0	233.5	129.1	57.05
SECOND† [46]	5.3	84.5	76.4	84.6	57.23
CP-Pillar† [49]	5.2	333.9	303.0	157.9	61.56
CP-Voxel† [49]	7.8	114.8	101.9	125.7	65.58
PV-RCNN [33]	13.1	117.7	399.4	623.2	63.33
PV-RCNN++† [34]	16.1	123.5	179.7	435.9	69.46
CP-Voxel-S + Ours	4.0	47.8	65.7	98.0	65.75
CP-Voxel-XS + Ours	2.8	36.9	58.4	88.1	64.83
CP-Voxel-XXS + Ours	1.0	12.0	33.1	70.4	60.93
CP-Pillar-v0.4 + Ours	5.2	212.9	197.7	103.4	61.60
CP-Pillar-v0.48 + Ours	5.2	149.4	142.3	81.9	60.95
CP-Pillar-v0.64 + Ours	5.2	85.1	88.0	54.5	58.89

Table 6: Cross stage 3D detector distillation.

Detector	Params (M)	Flops (G)	Acts (M)	Latency (ms)	Mem. (G)	LEVEL 2 mAPH
PV-RCNN++	16.1	123.5	179.7	435.9	4.2	67.80
CP-Voxel	7.8	114.8	101.9	125.7	2.8	64.29
CP-Voxel + Ours	7.8	114.8	101.9	125.7	2.8	65.27

from pretrained PVRCNN++, our distilled CP-Voxel achieves around 1% performance gains without any extra computation and parameter overheads during inference.

5.5 Generalization to More Scenarios

To demonstrate the generality of our compression and knowledge distillation manners, we provide experiments on other dataset, detector and compression manners. Besides, we construct experiment and discussion on extending our method to other task and more detectors in the Suppl..

Generality on Other Dataset and Detector. As shown in Table 7 and Table 8, both our compression conclusion in Sec. 3.3 and the improved KD method can generalize well to KITTI [9] dataset with a new anchor-based detector SECOND. Especially, the SECOND (a) surpasses teacher performance by around 0.5% with $3.5\times$ fewer flops, showing the generality of our conclusion and methods.

Table 7: Model and input compression results on KITTI. Teacher models are marked in gray.

Detector	Architecture			Efficiency				Moderate mAP@R40	CPR
	Width		Voxel Size (m)	Params (M)	Flops (G)	Acts (M)	Latency (ms)		
	PFE	BFE							
SECOND	1.00	1.00	0.05	5.3	80.5	69.3	77.4	67.24	-
	0.75	0.50	0.05	1.6	23.0	38.0	51.8	65.62	0.69
	0.50	0.50	0.05	1.4	20.5	35.9	46.1	64.21	0.68
	0.50	1.00	0.05	4.6	72.4	65.2	70.6	65.70	0.50
(d)	1.00	1.00	0.10	5.3	21.2	19.4	34.2	54.32	0.62

Table 8: Knowledge distillation results for 3D detection on KITTI. Performance are measured in moderate mAP over 40 recall positions. Best method is noted by **bold**.

Detector	No Distill	Logit KD		Feature KD				Label KD	Ours	Flops (G)	Acts (M)
		KD	GID-L	FitNet	Mimic	FG	GID-F				
SECOND	67.24	-	-	-	-	-	-	-	-	80.5	69.3
SECOND (a)	65.62	66.06	66.34	66.00	66.37	66.58	66.75	67.03	67.70	23.0	38.0

Generality on Advance Compression Manner. As shown in Table 9, since the coarser-resolution detector has more architecture-level redundancy with less input information, CP-Pillar (f) achieves higher CPR by combining the model and input compression strategies. Besides, after applying KD methods to CP-Pillar (f) as Table 10, it is still more efficient though fewer improvements from KD strategies with less redundancy. This demonstrates that our pipeline can substantially obtain a more efficient detector by harvesting the progress of advanced compression and distillation manners.

Table 9: Integrating different compression manners on Waymo. Teacher model is marked in gray.

Detector	Architecture			Efficiency				LEVEL 2 mAPH	CPR	
	Width			Voxel Size (m)	Params (M)	Flops (G)	Acts (M)			Latency (ms)
	PFE	BFE	Head							
CP-Pillar	1.00	1.00	1.00	0.32	5.2	333.9	303.0	157.9	59.09	-
CP-Pillar-v0.4	1.00	1.00	1.00	0.40	5.2	212.9	197.7	103.4	57.55	0.64
CP-Pillar (e)	1.00	0.875	0.875	0.32	4.0	260.1	267.7	134.7	58.53	0.54
CP-Pillar (f)	1.00	0.875	0.875	0.40	4.0	163.9	175.5	92.1	57.36	0.67

6 Ablation Studies

In this section, we conduct extensive ablation experiments to in-depth investigate the effectiveness of each component in our improved KD pipeline.

Component Ablation. Here, we investigate each component of our KD method and their synergy results. As shown in Table 11, based on CP-Voxel-XXS, both PP logit KD and TGI can obtain around 1.4% improvements separately. When incorporating with each other and label KD, they further obtain around 1.6% gains, showing the prominent synergy impact of our components. On the contrary, feature KD even suffers $0.2 \sim 0.6\%$ performance drop when combined with other KD techniques.

Investigation of TGI. We study how different parameter remapping manners and teachers influence the TGI. Besides FNA, we also attempt other parameter remapping strategies for TGI, including OFA [3] and Slim [26] by selecting important channels with designed indicators. As shown in Table 12, based on CP-Voxel-XS, OFA and Slim obtain inferior results compared to FNA, since indicator guided channel selection cannot determine consistent channel mapping for skip connection [13], while FNA survives by simply selecting beginning channels of all layers. In addition, CP-Voxel-XXS

Table 10: Knowledge distillation results for more sophisticated compressed 3D detector on Waymo. Teacher model is marked by gray Best method is noted by **bold**. § indicates the CPR is calculated according to the performance of best distilled student.

Detector	No Distill	Logit KD		Feature KD				Label KD	Ours	Flops (G)	Acts (M)	CPR [§]
		KD	GID-L	FitNet	Mimic	FG	GID-F					
CP-Pillar	59.09	-	-	-	-	-	-	-	-	333.9	303.0	-
CP-Pillar-v0.4	57.55	57.51	57.54	57.89	58.57	58.44	58.26	58.10	59.24	212.9	197.7	0.68
CP-Pillar (f)	57.36	56.93	56.70	57.15	57.81	57.48	57.77	57.57	58.62	163.9	175.5	0.70

Table 11: Component ablation study based on CP-Voxel-XXS.

PP logit KD	Label KD	TGI	Feature KD	mAPH
				56.26
✓				57.68
		✓		57.61
✓		✓		58.83
✓	✓			58.49
✓	✓	✓		59.28
✓	✓		✓	58.28
✓	✓	✓	✓	58.67

Table 12: Different remap manners studies of TGI based on CP-Voxel-XS and CP-Voxel-XXS.

Student	Teacher	Remap	mAPH
	-	-	61.16
CP-Voxel-XS	CP-Voxel	FNA	62.43
	CP-Voxel	OFA	60.06
	CP-Voxel	Slim	61.74
	-	-	56.26
CP-Voxel-XXS	CP-Voxel	FNA	54.92
	CP-Voxel-XS	FNA	57.61

only improves by inheriting parameters from CP-Voxel-XS but not CP-Voxel, which indicates a large architecture discrepancy between teacher and student hinders the effectiveness of TGI.

Different Variants of PP Logit KD. Here, we analyze different variants of PP logit KD (*i.e.* pivotal position selection according to teacher confidence, response ranking or soft Gaussian instance mask) in Section 5.1 based on CP-Voxel-S. While previous coarse instance-aware response imitation method GID-L obtains around 0.7% improvements, all three PP logit KD variants achieve around 1.6% ~ 1.9% gains. This demonstrates the significance of focusing on only vital positions in output-level distillation for 3D detection.

Table 13: Analysis of different PP logit KD variants based on CP-Voxel-S.

Logit KD Manner	mAPH
None	62.23
GID-L [6]	62.89
Confidence PP	63.84
Rank PP	63.85
Gaussian PP	64.16

Table 14: Analysis of label KD based on CP-Pillar-v0.48.

	Ground Truth		Teacher Pred		mAPH
	Cls	Reg	Cls	Reg	
(a)	all	all	no	no	56.24
(b)	all	all	all	no	56.47
(c)	all	non-overlap	all	all	56.99
(d)	all	all	all	non-duplicate	56.74
(e)	all	all	all	all	57.54

Analysis of Label KD. Although label KD has been proposed by the recent work [28], it just attributed its improvement to the “dark knowledge” from teacher. Here, we attempt to analyze how teacher predictions help student based on CP-Pillar-v0.48, where label KD performs best. As shown in Table 14, only integrating teacher prediction for classification label assignment, (b) achieves minor gains, which indicates that label KD mainly benefits the regression objective of student. When removing GTs highly overlapped with teacher predictions, (c) still achieves around 0.8% gains, which demonstrates that more achievable regression targets from teacher prediction facilitate student optimization. Last, by removing teacher predictions with object centers in the same BEV grid (*i.e.* duplicate objectives for single position), we find that (d) suffers performance drop compared to (e), which shows that duplicate regression objectives can also boost student regression ability.

7 Conclusion

We have examined the potential of knowledge distillation to serve as a generic method for obtaining efficient 3D detectors with extensive experimental results and analysis. We found that pillar-based detector prefers input compression while voxel-based detector is more suitable for width compression in designing efficient student models. Besides, we proposed pivotal position logit KD and teacher guided initialization for enhancing the 3D KD pipeline. Our best performing detector outperforms its teacher with $2.4\times$ fewer flops and our most efficient detector is $2.2\times$ faster than previous fastest voxel/pillar-based detector PointPillars on an NVIDIA A100 with higher performance. We hope our benchmark and analysis could inspire future investigations on this problem.

Acknowledgement. This work has been supported by Hong Kong Research Grant Council - Early Career Scheme (Grant No. 27209621), HKU Startup Fund, and HKU Seed Fund for Basic Research.

References

- [1] 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. [1](#), [2](#)
- [2] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020. [1](#)
- [3] Han Cai, Chuang Gan, Tianzhe Wang, Zhekai Zhang, and Song Han. Once-for-all: Train one network and specialize it for efficient deployment. *arXiv preprint arXiv:1908.09791*, 2019. [9](#)
- [4] Defang Chen, Jian-Ping Mei, Yuan Zhang, Can Wang, Zhe Wang, Yan Feng, and Chun Chen. Cross-layer distillation with semantic calibration. In *AAAI*, 2021. [3](#)
- [5] Yilun Chen, Shu Liu, Xiaoyong Shen, and Jiaya Jia. Fast point r-cnn. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9775–9784, 2019. [1](#), [2](#), [3](#)
- [6] Xing Dai, Zeren Jiang, Zhao Wu, Yiping Bao, Zhicheng Wang, Si Liu, and Erjin Zhou. General instance distillation for object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7842–7851, 2021. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [10](#)
- [7] Jiajun Deng, Shaoshuai Shi, Peiwei Li, Wengang Zhou, Yanyong Zhang, and Houqiang Li. Voxel r-cnn: Towards high performance voxel-based 3d object detection. *arXiv preprint arXiv:2012.15712*, 1(2):4, 2020. [1](#), [3](#), [8](#)
- [8] Jiemin Fang, Yuzhu Sun, Kangjian Peng, Qian Zhang, Yuan Li, Wenyu Liu, and Xinggang Wang. Fast neural network adaptation via parameter remapping and architecture search. *arXiv preprint arXiv:2001.02525*, 2020. [8](#)
- [9] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3354–3361. IEEE, 2012. [1](#), [9](#)
- [10] Benjamin Graham and Laurens van der Maaten. Submanifold sparse convolutional networks. *arXiv preprint arXiv:1706.01307*, 2017. [1](#), [4](#)
- [11] Xiaoyang Guo, Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. Liga-stereo: Learning lidar geometry aware representations for stereo-based 3d detector. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3153–3163, 2021. [3](#)
- [12] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. [6](#)
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [2](#), [9](#)
- [14] Byeongho Heo, Minsik Lee, Sangdoon Yun, and Jin Young Choi. Knowledge transfer via distillation of activation boundaries formed by hidden neurons. In *AAAI*, 2019. [3](#)
- [15] Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2(7), 2015. [1](#), [2](#), [3](#), [5](#), [6](#)
- [16] Yuenan Hou, Zheng Ma, Chunxiao Liu, and Chen Change Loy. Learning lightweight lane detection cnns by self attention distillation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1013–1021, 2019. [1](#)
- [17] Zehao Huang and Naiyan Wang. Like what you like: Knowledge distill via neuron selectivity transfer. *arXiv:1707.01219*, 2017. [3](#)

- [18] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. PMLR, 2015. 6
- [19] Xiao Jin, Baoyun Peng, Yichao Wu, Yu Liu, Jiaheng Liu, Ding Liang, Junjie Yan, and Xiaolin Hu. Knowledge distillation via route constrained optimization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019. 3
- [20] Nikos Komodakis and Sergey Zagoruyko. Paying more attention to attention: improving the performance of convolutional neural networks via attention transfer. In *International Conference on Learning Representations*, 2017. 3
- [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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 12697–12705, 2019. 1, 2, 3, 8
- [22] Quanquan Li, Shengying Jin, and Junjie Yan. Mimicking very efficient network for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6356–6364, 2017. 2, 3, 5, 6
- [23] Yanghao Li, Saining Xie, Xinlei Chen, Piotr Dollar, Kaiming He, and Ross Girshick. Benchmarking detection transfer learning with vision transformers. *arXiv preprint arXiv:2111.11429*, 2021. 3
- [24] Yifan Liu, Ke Chen, Chris Liu, Zengchang Qin, Zhenbo Luo, and Jingdong Wang. Structured knowledge distillation for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2604–2613, 2019. 1
- [25] Zhengzhe Liu, Xiaojuan Qi, and Chi-Wing Fu. 3d-to-2d distillation for indoor scene parsing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4464–4474, 2021. 3
- [26] Zhuang Liu, Jianguo Li, Zhiqiang Shen, Gao Huang, Shoumeng Yan, and Changshui Zhang. Learning efficient convolutional networks through network slimming. In *Proceedings of the IEEE international conference on computer vision*, pages 2736–2744, 2017. 9
- [27] Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines. In *International conference on machine learning*, 2010. 6
- [28] Chuong H Nguyen, Thuy C Nguyen, Tuan N Tang, and Nam LH Phan. Improving object detection by label assignment distillation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1005–1014, 2022. 2, 3, 5, 6, 7, 10
- [29] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Advances in Neural Information Processing Systems*, pages 5099–5108, 2017. 1, 2, 4
- [30] Lu Qi, Jason Kuen, Jiuxiang Gu, Zhe Lin, Yi Wang, Yukang Chen, Yanwei Li, and Jiaya Jia. Multi-scale aligned distillation for low-resolution detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14443–14453, 2021. 3, 4
- [31] Ilija Radosavovic, Raj Prateek Kosaraju, Ross Girshick, Kaiming He, and Piotr Dollár. Designing network design spaces. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10428–10436, 2020. 3
- [32] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. *arXiv preprint arXiv:1412.6550*, 2014. 2, 3, 5
- [33] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10529–10538, 2020. 1, 2, 3, 8

- [34] 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, 8
- [35] Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. Pointcnn: 3d object proposal generation and detection from point cloud. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 770–779, 2019. 1, 2
- [36] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 2
- [37] Pei Sun, Henrik Kretschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2446–2454, 2020. 1, 2, 3
- [38] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5725–5734, 2021. 2
- [39] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015. 2
- [40] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pages 6105–6114. PMLR, 2019. 4
- [41] OpenPCDet Development Team. Openpcdet: An open-source toolbox for 3d object detection from point clouds. <https://github.com/open-mmlab/OpenPCDet>, 2020. 3, 5
- [42] Tao Wang, Li Yuan, Xiaopeng Zhang, and Jiashi Feng. Distilling object detectors with fine-grained feature imitation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4933–4942, 2019. 2, 3, 5, 6, 7
- [43] Yue Wang, Alireza Fathi, Abhijit Kundu, David A Ross, Caroline Pantofaru, Tom Funkhouser, and Justin Solomon. Pillar-based object detection for autonomous driving. In *European Conference on Computer Vision*, pages 18–34. Springer, 2020. 1, 2, 3
- [44] Yue Wang, Alireza Fathi, Jiajun Wu, Thomas Funkhouser, and Justin Solomon. Multi-frame to single-frame: Knowledge distillation for 3d object detection. *arXiv preprint arXiv:2009.11859*, 2020. 3
- [45] Xiaohan Xing, Yuenan Hou, Hang Li, Yixuan Yuan, Hongsheng Li, and Max Q-H Meng. Categorical relation-preserving contrastive knowledge distillation for medical image classification. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 163–173. Springer, 2021. 1
- [46] Yan Yan, Yuxing Mao, and Bo Li. Second: Sparsely embedded convolutional detection. *Sensors*, 18(10):3337, 2018. 1, 2, 3, 8
- [47] Zetong Yang, Yanan Sun, Shu Liu, Xiaoyong Shen, and Jiaya Jia. Std: Sparse-to-dense 3d object detector for point cloud. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1951–1960, 2019. 2
- [48] Zhendong Yang, Zhe Li, Xiaohu Jiang, Yuan Gong, Zehuan Yuan, Danpei Zhao, and Chun Yuan. Focal and global knowledge distillation for detectors. *arXiv preprint arXiv:2111.11837*, 2021. 3
- [49] Tianwei Yin, Xingyi Zhou, and Philipp Krahenbuhl. Center-based 3d object detection and tracking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11784–11793, 2021. 1, 3, 5, 7, 8

- [50] 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. *arXiv preprint arXiv:2203.11139*, 2022. 1, 2, 3
- [51] Yin Zhou and Oncel Tuzel. Voxnet: End-to-end learning for point cloud based 3d object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4490–4499, 2018. 1, 2, 3

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes] We do not consider sophisticated layer-wise model compression strategies (see section 3). More details are described in supplemental file.
 - (c) Did you discuss any potential negative societal impacts of your work? [N/A] Our work is only for academic research purpose.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [N/A]
 - (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] Our implementation details are posted in supplemental materials and code will be available.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] We follow the popular codebase OpenPCDet (see section 3).
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] It is posted in supplemental materials
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See supplemental materials. Most of experiments are trained with 8 NVIDIA 1080Ti. Few experiments are trained with 8 NVIDIA V100 or NVIDIA A100. Full set results on Waymo are trained with 16 NVIDIA 1080ti.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [Yes] See our reference.
 - (b) Did you mention the license of the assets? [No] They are public released for academic utilization.
 - (c) Did you include any new assets either in the supplemental material or as a URL? [No]
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? [Yes] They are public released.
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [No] We only use public dataset.
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]