

UniDistill: A Universal Cross-Modality Knowledge Distillation Framework for 3D Object Detection in Bird’s-Eye View

Shengchao Zhou^{1*} Weizhou Liu^{1*} Chen Hu^{1†} Shuchang Zhou¹ Chao Ma²

¹ MEGVII Technology

² MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University

{zhoushengchao, liuweizhou, huchen, zsc}@megvii.com, chaoma@sjtu.edu.cn

Abstract

In the field of 3D object detection for autonomous driving, the sensor portfolio including multi-modality and single-modality is diverse and complex. Since the multi-modal methods have system complexity while the accuracy of single-modal ones is relatively low, how to make a trade-off between them is difficult. In this work, we propose a universal cross-modality knowledge distillation framework (UniDistill) to improve the performance of single-modality detectors. Specifically, during training, UniDistill projects the features of both the teacher and the student detector into Bird’s-Eye-View (BEV), which is a friendly representation for different modalities. Then, three distillation losses are calculated to sparsely align the foreground features, helping the student learn from the teacher without introducing additional cost during inference. Taking advantage of the similar detection paradigm of different detectors in BEV, UniDistill easily supports LiDAR-to-camera, camera-to-LiDAR, fusion-to-LiDAR and fusion-to-camera distillation paths. Furthermore, the three distillation losses can filter the effect of misaligned background information and balance between objects of different sizes, improving the distillation effectiveness. Extensive experiments on nuScenes demonstrate that UniDistill effectively improves the mAP and NDS of student detectors by 2.0%~3.2%.

1. Introduction

3D object detection plays a critical role in autonomous driving and robotic navigation. Generally, the popular 3D detectors can be categorized into (1) single-modality detectors that are based on LiDAR [18, 33, 34, 42, 43] or camera [1, 13, 20, 24] and (2) multi-modality detectors [22, 30, 36, 37] that are based on both modalities. By fusing the complementary knowledge of two modalities,

*Equal Contribution

†Corresponding Author

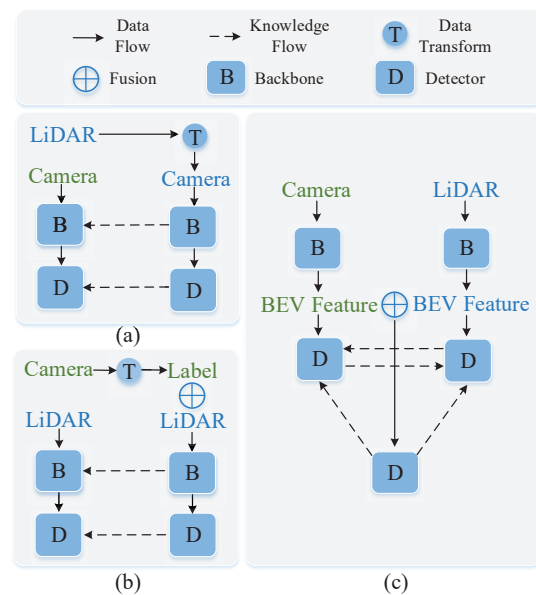


Figure 1. Illustration of our proposed UniDistill. The characters in green and blue represent the data process of camera and LiDAR respectively. (a) and (b) show the procedure of two previous knowledge distillation methods, where the modalities of the teacher and the student are restricted. By contrast, our proposed UniDistill in (c) supports four distillation paths.

multi-modality detectors outperform their single-modality counterparts. Nevertheless, simultaneously processing the data of two modalities unavoidably introduces extra network designs and computational overhead. Worse still, the breakdown of any modality directly fails the detection, hindering the application of these detectors.

As a solution, some recent works introduced knowledge distillation to transfer complementary knowledge of other modalities to a single-modality detector. In [6, 15, 46], as illustrated in Figure 1(a) and 1(b), for a single-modality student detector, the authors first performed data transformation of different modalities to train a structurally identical

teacher. The teacher was then leveraged to transfer knowledge by instructing the student to produce similar features and prediction results. In this way, the single-modality student obtains multi-modality knowledge and improves performance, without additional cost during inference.

Despite their effectiveness to transfer cross-modality knowledge, the application of existing methods is limited since the modalities of both the teacher and the student are restricted. In [6], the modalities of the teacher and student are fixed to be LiDAR and camera while in [15, 46], they are determined to be LiDAR-camera and LiDAR. However, the sensor portfolio in the field of 3D object detection results in a diverse and complex application of different detectors. With restricted modalities of both the teacher and student, these methods are difficult to be applied in more situations, *e.g.*, the method in [6] is not suitable to transfer knowledge from a camera based teacher to a LiDAR based student.

To solve the above problems, we propose a universal cross-modality knowledge distillation framework (UniDistill) that helps single-modality detectors improve performance. Our motivation is based on the observation that the detectors of different modalities adopt a similar detection paradigm in bird’s-eye view (BEV), where after transforming the low-level features to BEV, a BEV encoder follows to further encode high-level features and a detection head produces response features to perform final prediction.

UniDistill takes advantage of the similarity to construct the universal knowledge distillation framework. As in Figure 1(c), during training, UniDistill projects the features of both the teacher and the student detector into the unified BEV domain. Then for each ground truth bounding box, three distillation losses are calculated to transfer knowledge: (1) A feature distillation loss that transfers the semantic knowledge by aligning the low-level features of 9 crucial points. (2) A relation distillation loss that transfers the structural knowledge by aligning the relationship between the high-level features of 9 crucial points. (3) A response distillation loss that closes the prediction gap by aligning the response features in a Gaussian-like mask. Since the aligned features are commonly produced by different detectors, UniDistill easily supports LiDAR-to-camera, camera-to-LiDAR, fusion-to-LiDAR and fusion-to-camera distillation paths. Furthermore, the three losses sparsely align the foreground features to filter the effect of misaligned background information and balance between objects of different scales, improving the distillation effectiveness.

In summary, our contributions are three-fold:

- We propose a universal cross-modality knowledge distillation framework (UniDistill) in the friendly BEV domain for single-modality 3D object detectors. With the transferred knowledge of different modalities, the performance of single-modality detectors is improved without additional cost during inference.
- Benefiting from the similar detection paradigm in BEV, UniDistill supports LiDAR-to-camera, camera-to-LiDAR, fusion-to-LiDAR and fusion-to-camera distillation paths. Moreover, three distillation losses are designed to sparsely align foreground features, filtering the effect of background information misalignment and balance between objects of different sizes.
- Extensive experiments on nuScenes demonstrate that UniDistill can effectively improve the mAP and NDS of student detectors by 2.0%~3.2%.

2. Related Work

2.1. 3D Object Detection

Recent mainstream 3D object detectors can be generally divided into two categories: (1) Single-modality detectors based on LiDAR [23, 25, 27, 32, 43] or camera [13, 20, 28, 38–40] (2) Multi-modality detectors [14, 22, 30, 44] with both types of data as input. One category of LiDAR based detectors are grid-based [17, 26, 41, 43, 47], where the unstructured LiDAR points are first distributed to regular grids. As the seminal work, VoxelNet [47] voxelizes the point clouds, performs 3D convolution, reshapes the features into BEV features and then proposes bounding boxes. PointPillars [17] substitutes the voxels with pillars to encode point clouds, avoiding time-consuming 3D convolution operations. CenterPoint [43] proposes an anchor-free detection head and obtains better performance.

Most camera based detectors perform detection in perspective view like 2D detection. In [38], the authors proposed FCOS3D, which is an extension of FCOS [35] to the field of 3D object detection. Similarly in [40], following DETR [2], the authors proposed DETR3D to perform detection in an attention pattern. Recently, some works [13, 20] are proposed to detect objects in BEV by applying Lift-Splat-Shoot [29] to transform image features from perspective view to BEV and achieve satisfactory improvement.

LiDAR-camera based detectors outperform the above single-modality counterparts by fusing the complementary knowledge of LiDAR points and images. With 2D detection results, F-PointNet [30] obtains candidate object areas, gathers the points inside and then performs LiDAR based detection. AVOD [16] and MV3D [5] perform modality fusion with object proposals via ROI pooling. UVTR [19] unifies multi-modality representations in the voxel space for transformer-based 3D object detection. A recent state-of-the-art detector BEVFusion [22] proposes to transform both the image and the LiDAR features to BEV for modality fusion and result prediction.

Our work is inspired by that the recent detectors adopt a similar detection paradigm in BEV. After transforming the features to BEV, a similar procedure follows, where a BEV




Figure 2. The overview of our proposed UniDistill. The branches on the top and bottom are the teacher detector and student detector respectively. For a teacher and a student, three distillation losses to align specific foreground features are calculated after extracting the low-level features and transforming them to BEV, which are feature distillation, relation distillation and response distillation.

encoder further encodes high-level features and a detection head produces prediction results. The similarity of the detection paradigm makes cross-modality knowledge distillation in a universal framework possible.

2.2. Knowledge Distillation

Knowledge distillation is initially proposed in [10] for model compression and the main idea is transferring the learned knowledge from a teacher network to a student. Different works have different interpretations of the knowledge, which include the soft targets of the output layer [10] and the intermediate feature map [31]. Because of the effectiveness of knowledge distillation, it has been widely investigated in a variety of computer vision tasks, such as 2D object detection [7, 9, 45] and semantic segmentation [11, 21].

Recently, it is introduced into 3D object detection for knowledge transfer to single-modality detectors. In [6], the authors proposed to transfer the depth knowledge of LiDAR points to a camera based student detector by training another camera based teacher with LiDAR projected to perspective view. In [15, 46], a PointPainting [36] teacher is leveraged to instruct a CenterPoint student to produce similar features and responses.

Although these methods are effective for knowledge transfer, the modalities of the teacher and the student are restricted. However, the diverse and complex application of different detectors will restrict their application. Instead, our proposed UniDistill projects the features of detectors

to BEV and supports LiDAR-to-camera, camera-to-LiDAR, fusion-to-LiDAR and fusion-to-camera distillation paths.

We note that a recent work BEVDistill also performs cross-modality distillation in BEV. However, it also imposes restrictions on the modalities of the teacher and the student, resulting in a limited application. Moreover, their main distillation losses are designed for transformer based detectors. Differently, our UniDistill aims to improve the performance of CNN based detectors.

3. Methodology

In this section, we describe our proposed UniDistill in detail. In Section 3.1, we briefly introduce the similar detection paradigm of different detectors in BEV, where some low-level BEV features are obtained via view-transform and further encoded to be high-level features and response features. In Section 3.2, we depict the proposed UniDistill framework for cross-modality knowledge distillation. For a teacher and a student detector, after obtaining their features in BEV, three distillation losses aligning foreground features are calculated to perform knowledge transfer.

3.1. Preliminary

The recent detectors adopt a similar detection paradigm in BEV, where the major difference is the approach to get the low-level BEV features. For the camera based detectors proposed in [13, 20], after pixel-wise depth estimation, the image features are projected from the perspective view into

BEV via the view-transform operation proposed in [29] to form low-level BEV features $\mathbf{F}_{\text{cam}}^{\text{low}}$.

With respect to LiDAR based detectors presented in [17, 43, 47], the unstructured point clouds are first distributed to regular voxels or pillars. The features of voxels or pillars are then extracted and reshaped into low-level BEV features $\mathbf{F}_{\text{ldr}}^{\text{low}}$ by concatenating the voxel features in the same column.

In [22], the authors proposed a method to construct a LiDAR-camera based detector simply by fusing the low-level features $\mathbf{F}_{\text{ldr}}^{\text{low}}$ and $\mathbf{F}_{\text{cam}}^{\text{low}}$ of a LiDAR and camera based detector. It first concatenates $\mathbf{F}_{\text{ldr}}^{\text{low}}$ and $\mathbf{F}_{\text{cam}}^{\text{low}}$ and then processes the result with a fully convolutional network to produce the fused features $\mathbf{F}_{\text{fuse}}^{\text{low}}$.

The following steps are the same for different detectors. For a detector of any modality, a BEV encoder (BEVEnc) first takes its low-level BEV features $\mathbf{F}_{\text{mod}}^{\text{low}}$ as input to further encode the high-level features $\mathbf{F}_{\text{mod}}^{\text{high}}$:

$$\mathbf{F}_{\text{mod}}^{\text{high}} = \text{BEVEnc}(\mathbf{F}_{\text{mod}}^{\text{low}}), \quad (1)$$

where mod is the detector modality and is in {ldr,cam,fuse}. Then a detection head (DetHead) produces the classification and regression heatmaps $\mathbf{F}_{\text{mod}}^{\text{cls}}$ and $\mathbf{F}_{\text{mod}}^{\text{reg}}$, based on which the final predictions are generated:

$$\mathbf{F}_{\text{mod}}^{\text{cls}}, \mathbf{F}_{\text{mod}}^{\text{reg}} = \text{DetHead}(\mathbf{F}_{\text{mod}}^{\text{high}}). \quad (2)$$

Therefore, regardless of the modality, these detectors will consistently produce \mathbf{F}^{low} , \mathbf{F}^{high} , \mathbf{F}^{cls} and \mathbf{F}^{reg} during the procedure of detection.

3.2. UniDistill

Taking advantage of the above similar detection paradigm in BEV, we propose a universal cross-modality knowledge distillation framework (UniDistill), which easily supports LiDAR-to-camera, camera-to-LiDAR, fusion-to-LiDAR and fusion-to-camera distillation paths. The overview of UniDistill is illustrated in Figure 2. During training, after transforming the low-level features of the teacher and the student of different modalities to BEV, three distillation losses are then calculated. The losses are finally combined with the original detection loss to train the student. In this way, the student can mimic the teacher to learn cross-modal knowledge and thus gain better detection results, without introducing additional cost during inference.

Denoting the modalities of the teacher and the student as MT and MS respectively, the three distillation losses can be interpreted as follows: (1) The first loss ‘‘feature distillation’’ transfers the semantic knowledge in the low-level BEV features $\mathbf{F}_{\text{MT}}^{\text{low}}$ to $\mathbf{F}_{\text{MS}}^{\text{low}}$ by point-wisely aligning the features of 9 crucial points of each ground truth bounding box. (2) The second loss ‘‘relation distillation’’ transfers the structural knowledge in the high-level BEV features $\mathbf{F}_{\text{MT}}^{\text{high}}$

to $\mathbf{F}_{\text{MS}}^{\text{high}}$ by group-wisely aligning the relationship between 9 crucial points of each ground truth bounding box. (3) The third loss ‘‘response distillation’’ closes the prediction gap by aligning the heatmaps $(\mathbf{F}_{\text{MS}}^{\text{cls}}, \mathbf{F}_{\text{MS}}^{\text{reg}})$ with $(\mathbf{F}_{\text{MT}}^{\text{cls}}, \mathbf{F}_{\text{MT}}^{\text{reg}})$ in a Gaussian-like mask for each ground truth bounding box.

3.2.1 Feature Distillation

Since the low-level BEV features provide semantic knowledge for further process, we propose feature distillation to align $\mathbf{F}_{\text{MS}}^{\text{low}}$ with $\mathbf{F}_{\text{MT}}^{\text{low}}$. One intuitive method is completely aligning $\mathbf{F}_{\text{MS}}^{\text{low}}$ with $\mathbf{F}_{\text{MT}}^{\text{low}}$, however, the background information misalignment between different modalities will decrease the effectiveness. Worse still, because different objects occupy areas of different sizes, it will focus more on aligning the features of large objects than small objects.

To mitigate the above effects, in feature distillation, we only align the features of foreground objects and equally select 9 crucial points for each of them for alignment. Specifically, in BEV, the bounding box of one foreground object can be regarded as a 2D rotated box and described by the coordinates of its corners $\{(x_i, y_i)\}_{i=1}^4$. Based on the four corners, the coordinates of midpoints of 4 edges and the center of the box can be calculated. We collect these 9 points together as $\{(x_i, y_i)\}_{i=1}^9$ and regard them as crucial points for the bounding box. For each point $\mathbf{p}_i=(x_i, y_i)$, we calculate the difference between its features on $\mathbf{F}_{\text{MT}}^{\text{low}}$ and $\mathbf{F}_{\text{MS}}^{\text{low}}$ to form the feature distillation loss \mathcal{L}_{Fea} :

$$\mathcal{L}_{\text{Fea}} = \left(\sum_{i=1}^9 |\mathbf{F}_{\text{MT}}^{\text{low}}(x_i, y_i) - \mathbf{F}_{\text{MS}}^{\text{low}}(x_i, y_i)| \right) / 9. \quad (3)$$

In addition, we note that when the performance of the teacher detector is worse than the student, *e.g.*, the modalities of the teacher and the student are camera and LiDAR, using the vanilla feature distillation can even degrade the final performance. Inspired by the feature adaptation operation in [3, 31], in this situation, we also introduce an adaptive layer Adapt_1 , which is a one-layer convolutional network, after $\mathbf{F}_{\text{MS}}^{\text{low}}$ to produce new features $\hat{\mathbf{F}}_{\text{MS}}^{\text{low}}$ and calculate feature distillation instead with $\hat{\mathbf{F}}_{\text{MS}}^{\text{low}}$ and $\mathbf{F}_{\text{MT}}^{\text{low}}$. During inference, the adaptive layer is removed and the original low-level feature $\mathbf{F}_{\text{MS}}^{\text{low}}$ is further processed to generate predictions. Therefore, there is no modification to the structure of the student detector.

3.2.2 Relation Distillation

The high-level features can provide knowledge about the structure of the scene, which is important for the final prediction. To transfer the structural knowledge from the teacher to the student, we propose relation distillation to align the relationship between the different parts of an object in $\mathbf{F}_{\text{MS}}^{\text{high}}$ and $\mathbf{F}_{\text{MT}}^{\text{high}}$. Specifically, for one ground

truth bounding box, we also consider its 9 crucial points $\{(x_i, y_i)\}_{i=1}^9$ the same as those in feature distillation. We first gather their features on $\mathbf{F}_{\text{MS}}^{\text{high}}$ and calculate the relationship between them to form a relation matrix $\text{RelMat}^{\text{MS}}$ with the size of 9×9 :

$$\text{RelMat}_{i,j}^{\text{MS}} = \Phi(\mathbf{F}_{\text{MS}}^{\text{high}}(x_i, y_i), \mathbf{F}_{\text{MS}}^{\text{high}}(x_j, y_j)), \quad (4)$$

where $1 \leq i, j \leq 9$ and Φ represents the cosine similarity function. With the same operation, another relation matrix $\text{RelMat}^{\text{MT}}$ can be calculated based on $\mathbf{F}_{\text{MT}}^{\text{high}}$ and then the relation distillation loss \mathcal{L}_{Rel} is calculated to completely align $\text{RelMat}^{\text{MS}}$ with $\text{RelMat}^{\text{MT}}$:

$$\mathcal{L}_{\text{Rel}} = \left(\sum_{1 \leq i, j \leq 9} |\text{RelMat}_{i,j}^{\text{MT}} - \text{RelMat}_{i,j}^{\text{MS}}| \right) / 81. \quad (5)$$

Furthermore, the same as that in feature distillation, when the teacher performs worse than student, we introduce another one-layer convolutional network as an adaptive layer Adapt_2 after $\mathbf{F}_{\text{MS}}^{\text{high}}$ to produce new features $\hat{\mathbf{F}}_{\text{MS}}^{\text{high}}$ and calculate relation distillation with $\hat{\mathbf{F}}_{\text{MS}}^{\text{high}}$ and $\mathbf{F}_{\text{MT}}^{\text{high}}$.

3.2.3 Response Distillation

To make the final prediction of the student similar to the teacher, we further propose response distillation. Based on the high-level BEV features, the detection head will produce a classification heatmap $\mathbf{F}^{\text{cls}} \in R^{H \times W \times C}$ and a regression heatmap $\mathbf{F}^{\text{reg}} \in R^{H \times W \times T}$, where H and W are the height and width of the heatmap, C is the number of classes to be predicted and T is the number of regression targets. We further gather the max value of each position in \mathbf{F}^{cls} to form a new heatmap $\mathbf{F}_{\text{max}}^{\text{cls}}$:

$$\mathbf{F}_{\text{max}}^{\text{cls}}(i, j) = \max_{1 \leq k \leq C} \mathbf{F}^{\text{cls}}(i, j, k), \quad (6)$$

where $1 \leq i \leq H$ and $1 \leq j \leq W$. Then $\mathbf{F}_{\text{max}}^{\text{cls}}$ is concatenated with \mathbf{F}^{reg} to be the response features \mathbf{F}^{resp} . In this way, we obtain the response features of the teacher and the student, which are $\mathbf{F}_{\text{MT}}^{\text{resp}}$ and $\mathbf{F}_{\text{MS}}^{\text{resp}}$ respectively, and calculate response distillation loss to align $\mathbf{F}_{\text{MS}}^{\text{resp}}$ with $\mathbf{F}_{\text{MT}}^{\text{resp}}$. To mitigate the effect of background information misalignment, we only align the part of response features near the foreground objects. Instead of selecting 9 crucial points, we find that the quality of the response values in $\mathbf{F}_{\text{MT}}^{\text{resp}}$ near the center of a ground truth bounding box is good enough to guide $\mathbf{F}_{\text{MS}}^{\text{resp}}$. Therefore, for one ground truth bounding box, we generate a Gaussian-like mask with the same method in [6], gather the response values inside the mask and calculate response distillation loss $\mathcal{L}_{\text{Resp}}$ to align them:

$$\mathcal{L}_{\text{Resp}} = \sum_{1 \leq i \leq H, 1 \leq j \leq W} |\mathbf{F}_{\text{MT}}^{\text{resp}}(i, j) - \mathbf{F}_{\text{MS}}^{\text{resp}}(i, j)| \times \text{Mask}(i, j), \quad (7)$$

where Mask is the calculated Gaussian-like mask and only the area near the ground truth bounding box is non-zero.

3.2.4 Total Objectives

After being calculated for each ground truth bounding box, every distillation loss is then averaged over all bounding boxes. Finally, we combine the detection loss \mathcal{L}_{Det} of the student with the distillation losses as the total loss $\mathcal{L}_{\text{Total}}$:

$$\mathcal{L}_{\text{Total}} = \mathcal{L}_{\text{Det}} + \lambda_1 \cdot \mathcal{L}_{\text{Fea}} + \lambda_2 \cdot \mathcal{L}_{\text{Rel}} + \lambda_3 \cdot \mathcal{L}_{\text{Resp}}, \quad (8)$$

where λ_1 , λ_2 and λ_3 are hyperparameters used to balance the scale of different losses. The student is then optimized by $\mathcal{L}_{\text{Total}}$, achieving better performance.

4. Experiments

4.1. Experimental Setup

4.1.1 Dataset

We choose the popular large-scale dataset nuScenes for evaluation in our experiments. This dataset consists of 1,000 driving sequences (700/150/150 for train/val/test) with images from 6 cameras, points from 5 Radars and 1 LiDAR. Each sequence is 20 seconds long and every 10 frames are fully annotated with 3D object bounding boxes. For each frame, there are 6 RGB camera images and 20 LiDAR scanning frames. Following BEVFusion [22], we set the area inside $[-54.0\text{m}, 54.0\text{m}] \times [-54.0\text{m}, 54.0\text{m}] \times [-5\text{m}, 3\text{m}]$ as our region of interest.

4.1.2 Evaluation Metrics

We use the official evaluation metrics of nuScenes, which are the mean Average Precision(mAP), mean Average Translation Error(mATE), mean Average Scale Error(mASE), mean Average Orientation Error(mAOE), mean Average Velocity Error(mAVE), mean Average Attribute Error(mAAE) and the nuScenes detection score(NDS). The mAP measures the recall and precision of predicted bounding boxes. NDS is the composition of other metrics to comprehensively judge the detection capacity. The remaining metrics calculate the positive results of prediction on the corresponding aspects. The data unit of results is “%”.

4.1.3 Models and Training

We adopt BEVDet [13] with ResNet-50 as our camera based detector and CenterPoint [43] as our LiDAR based detector. The LiDAR-camera detector is constructed by fusing the low-level BEV features of BEVDet and CenterPoint, following the method proposed in BEVFusion [22]. The size of voxel to produce BEV features is $0.075\text{m} \times 0.075\text{m} \times 0.2\text{m}$.

The detectors are trained by an AdamW optimizer with the learning rate to be $1\text{e-}4$ for LiDAR and LiDAR-camera

Table 1. Performance analysis of UniDistill in four distillation paths on the nuScenes test dataset. “**” indicates re-implementation on our new student detector. “L” and “C” represent the LiDAR and camera.

Method	Modality	Teacher Modality	mAP \uparrow	NDS \uparrow	mATE \downarrow	mASE \downarrow	mAOE \downarrow	mAVE \downarrow	mAAE \downarrow
CVCNet [4]	L	-	55.8	64.2	30.0	24.8	43.1	26.9	11.9
Guided 3DOD [8]	L	-	60.9	67.3	28.8	24.5	40.0	25.3	12.8
AFDetV2 [12]	L	-	62.4	68.5	25.7	23.4	34.1	29.9	13.7
S2M2-SSD* [46]	L	L+C	63.6	69.6	25.4	23.9	34.3	25.7	12.8
UniDistill	L+C	-	65.4	70.6	25.1	23.8	32.5	25.6	12.8
UniDistill	L	-	61.4	67.8	26.8	25.1	33.6	27.8	14.8
UniDistill	L	C	63.4(+2.0)	69.8(+2.0)	24.9(-1.9)	23.7(-1.4)	32.1(-1.5)	24.7(-3.1)	13.1(-1.7)
UniDistill	L	L+C	63.9(+2.5)	70.1(+2.3)	25.0(-1.8)	23.8(-1.3)	32.8(-0.8)	24.5(-3.3)	12.7(-2.1)
UniDistill	C	-	26.4	36.1	69.7	26.6	55.8	117.3	17.8
UniDistill	C	L	28.9(+2.5)	38.4(+2.3)	65.9(-3.8)	25.9(-0.7)	51.4(-4.4)	106.4(-10.9)	17.0(-0.8)
UniDistill	C	L+C	29.6(+3.2)	39.3(+3.2)	63.7(-6.0)	25.7(-0.9)	49.2(-6.6)	108.4(-8.9)	16.7(-1.1)

based detectors and $2e-4$ for the camera based detector. The training batch size is 20 and the training epoch is 20 for all detectors. More details are supplemented in the appendix.

UniDistill is evaluated to see whether it can transfer knowledge in four distillation paths: (1) From the LiDAR-camera based teacher detector to the LiDAR based student. (2) From the LiDAR-camera based teacher detector to the camera based student. (3) From the camera based teacher detector to the LiDAR based student. (4) From the LiDAR based teacher detector to the camera based student. The hyperparameters λ_1, λ_2 and λ_3 in each path are: (1) 10, 1, 10. (2) 10, 5, 10. (3) 10, 5, 1. (4) 100, 40, 10. The adaptive layers for feature distillation and relation distillation are introduced when evaluating in path (3).

4.2. Comparison with the State-of-the-Arts

We first evaluate the performance of UniDistill on the test dataset of nuScenes and Table 1 reports the results. It is revealed that in all of the four distillation paths, UniDistill helps transfer knowledge from the teacher detector to student and improve its performance. Besides, with the LiDAR-camera based detector as the teacher, the LiDAR based student obtains better performance than other state-of-the-art LiDAR based detectors, proving the effectiveness of UniDistill. We also compare UniDistill with S2M2-SSD [46], which performs cross-modality knowledge distillation from a PointPainting [36] teacher detector also to a CenterPoint student detector. The result shows that UniDistill helps the student obtain better performance.

4.3. Ablation Studies

In this section, some experiments are conducted on the validation dataset to show the effect of each distillation loss and the rationality of specific designs. For efficiency, we turn off the auto-scaling between classification and regression loss in \mathcal{L}_{Det} . We report the overall results in Table 2 to show the effect of each loss.

4.3.1 Effect of Feature Distillation

As in the second setting of Table 2, using the feature distillation improves the NDS and mAP of the student in four paths. Moreover, the fifth, seventh and eighth settings prove that it is complementary to other distillation losses.

Another experiment in path (4) is conducted to show the rationality of feature distillation to align the features of 9 crucial points. We compare the original feature distillation with two modified ones which align the low-level BEV features (1) completely or (2) inside a Gaussian-like mask like the response distillation. The results in Table 3 show that feature distillation performs better when selecting 9 crucial points for alignment. Moreover, in this situation, the AP of small objects (pedestrians and motors) improves a lot while the influence on large objects (cars and trucks) is minor.

We further compare the original feature distillation with a modified one which aligns the high-level BEV features of 9 crucial points in path (4). The results illustrated in Table 4 reveal that calculating feature distillation with the low-level BEV features obtains better performance.

4.3.2 Effect of Relation Distillation

The comparison between the third setting and first setting of Table 2 shows that using the relation distillation improves the NDS and mAP of the student in all paths. And the results in the fifth, sixth and eighth settings further prove that it is complementary to other distillation losses.

We conduct another experiment in path (4) to show the rationality of selecting 9 crucial points for relation distillation calculation. We compare the original relation distillation with two modified ones which align the relationship between (1) all of the high-level BEV features or (2) the features inside a Gaussian-like mask like the response distillation. The results illustrated in Table 5 show that the relation distillation obtains the best performance when calculating the relationship between 9 crucial points for alignment.

Table 2. Ablation study of three proposed distillation losses on the nuScenes validation dataset. “T” and “S” represent the teacher detector and student detector respectively and the mAP/NDS of the student is reported.

Setting	Loss			Modality			
	\mathcal{L}_{Fea}	\mathcal{L}_{Rel}	$\mathcal{L}_{\text{Resp}}$	T:LiDAR+Camera S:LiDAR	T:LiDAR+Camera S:Camera	T:Camera S:LiDAR	T:LiDAR S:Camera
1				53.5/63.9	20.3/33.1	53.5/63.9	20.3/33.1
2	✓			56.1/65.5	21.6/34.5	54.3/64.6	21.1/34.3
3		✓		54.1/64.0	22.3/35.7	55.2/65.3	21.7/35.0
4			✓	58.7/66.7	25.7/37.1	55.7/65.6	24.9/36.3
5	✓	✓		-	-	-	23.5/35.4
6		✓	✓	-	-	-	25.3/36.7
7	✓		✓	-	-	-	25.3/37.0
8	✓	✓	✓	59.7/67.5	26.5/37.8	57.0/66.3	26.0/37.3

Table 3. Ablation study in path (4) to show that feature distillation performs better when selecting crucial points for alignment.

Method	AP					NDS
	car	truck	ped	motor	mean	
Baseline	38.5	20.1	9.4	18.5	20.3	33.1
Complete	38.0	13.1	14.2	22.0	20.3	32.6
Gaussian	45.3	21.4	10.3	16.8	20.6	32.8
Crucial	44.0	14.9	21.9	22.4	21.1	34.3

Table 4. Ablation study in path (4) to show that feature distillation performs better when aligning the low-level BEV features.

Method	mAP	mASE	mAOE	mAAE	NDS
Baseline	20.3	27.9	46.6	21.9	33.1
High-Level	20.6	28.1	46.9	23.2	32.3
Low-Level	21.1	27.8	46.3	21.9	34.3

We further compare the original relation distillation with the modified one, which aligns the relationship between the low-level BEV features of 9 crucial points. The results in path (4) are illustrated in Table 6 and reveal that calculating relation distillation with the high-level BEV features obtains better performance.

4.3.3 Effect of Response Distillation

As the fourth setting of Table 2 shows, using the response distillation improves the NDS and mAP of the student in all paths. Furthermore, the sixth, seventh and eighth settings prove that it is complementary to other distillation losses.

Another experiment in path (4) is conducted to prove the rationality of aligning the response features inside the Gaussian-like mask for response distillation. The original response distillation is compared with two modified ones which align the response features (1) completely or (2) of 9 crucial points like those in feature distillation. The results illustrated in Table 7 show that the response distillation obtains better performance when selecting the response fea-

Table 5. Ablation study in path (4) to show that relation distillation performs better when selecting crucial points for alignment.

Method	AP					NDS
	car	truck	ped	motor	mean	
Baseline	38.5	20.1	9.4	18.5	20.3	33.1
Complete	44.0	18.3	15.9	17.9	20.4	33.4
Gaussian	44.9	18.3	15.3	18.1	20.8	34.0
Crucial	44.1	18.3	19.2	18.8	21.7	35.0

Table 6. Ablation study in path (4) to show that relation distillation performs better when aligning the high-level BEV features.

Method	mAP	mASE	mAOE	mAAE	NDS
Baseline	20.3	27.9	46.6	21.9	33.1
Low-Level	19.7	27.7	54.4	24.3	32.3
High-Level	21.7	28.1	46.2	21.2	35.0

tures inside the Gaussian-like mask for alignment.

To form the response features, we first gather the max value of each position in the classification heatmap to obtain a new heatmap. In path (4), we further compare the original response distillation with a modified one based on the response features formed by concatenating the original classification and regression heatmaps. The results in Table 8 show that gathering the max value to form the response features helps response distillation perform better.

4.3.4 Effect of Adaptive Layers

When evaluating in path (3), we introduce two adaptive layers Adapt_1 and Adapt_2 after the low- and high-level BEV features to avoid performance degradation. Here we conduct experiments to show the indispensability of these adaptive layers by removing them and re-evaluating the performance. The results in Table 9 show that without the adaptive layers, the feature distillation and relation distillation even worsen the performance of the student detector. Therefore, in this situation, we keep the two adaptive layers for help.




Figure 3. Illustration of detection results. The boxes in red and green are the predicted and ground truth bounding boxes respectively. (a), (b) and (c) show the results of the LiDAR-camera based teacher detector, the LiDAR based student detector without UniDistill and the LiDAR based student detector with UniDistill respectively. The results show that UniDistill helps the student detector generate more accurate predictions and fewer false positive results.

Table 7. Ablation study in path (4) to show that response distillation performs better when aligning the features in Gaussian mask.

Method	AP					NDS
	car	truck	ped	motor	mean	
Baseline	38.5	20.1	9.4	18.5	20.3	33.1
Complete	45.3	18.8	13.6	22.1	23.4	34.9
Crucial	46.7	17.4	14.6	22.5	23.4	35.3
Gaussian	47.8	21.1	26.5	23.4	24.9	36.3

Table 8. Ablation study in path (4) to show that gathering the max value of each position in the classification heatmap helps response distillation perform better.

Method	mAP	mASE	mAOE	mAAE	NDS
Baseline	20.3	27.9	46.6	21.9	33.1
W/O Max	24.3	27.5	46.7	22.5	35.4
Max	24.9	27.1	42.6	24.0	36.3

Table 9. Ablation study in path (3) to show that the adaptive layers help feature distillation and relation distillation perform better.

Method	\mathcal{L}_{Fea}			\mathcal{L}_{Rel}		
	mAP	mAVE	NDS	mAP	mAVE	NDS
Baseline	53.5	22.5	63.9	53.5	22.5	63.9
W/O Adapt	53.3	25.3	63.5	53.1	24.3	63.2
With Adapt	54.3	21.3	64.6	55.2	21.4	65.3

We further compare the detection loss \mathcal{L}_{Det} with/without the adaptive layers and the baseline is the student without UniDistill. The results in Figure 4 show that, with the adaptive layers, the detection loss gradually becomes lower than the baseline. However, without the adaptive layers, the detection loss is always larger than the baseline, leading to worse performance. We think the problem results from that the feature quality of camera based teacher is worse than the LiDAR based student. Without the adaptive layers, aligning the features of student with teacher can decrease their quality. However, with the adaptive layers, the student can decide whether to learn from the teacher so that the performance is improved.

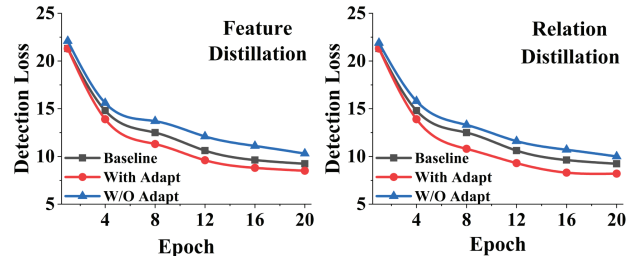


Figure 4. Illustration to show that the adaptive layers help feature distillation and relation distillation decrease the detection loss.

4.4. Visualization

In this section, we visualize the 3D object detection results to qualitatively show the effectiveness of UniDistill. The results are illustrated in Figure 3, where the teacher and the student are LiDAR-camera based and LiDAR based. From the results, the student detector can localize objects better with the help of UniDistill. Moreover, due to the balance between objects of different sizes, there are fewer false positive predictions on small objects.

5. Conclusion

In this work, we propose a universal cross-modality knowledge distillation framework (UniDistill) to improve the performance of single-modality 3D object detectors in BEV. UniDistill projects the features of both the teacher and student into a unified BEV domain and then calculates three distillation losses to align the features for knowledge transfer. Taking advantage of the similar detection paradigm in BEV, UniDistill supports LiDAR-to-camera, camera-to-LiDAR, fusion-to-LiDAR and fusion-to-camera distillation paths. Furthermore, the proposed three distillation losses sparsely align foreground features to filter the misaligned background information and balance between objects of different sizes. Extensive experiments demonstrate that UniDistill is effective to improve the performance of student detectors. Inspired by the merits of block-wise distillation, in the future, we plan to leverage the distillation losses in a block-wise manner for acceleration, to further explore the potential of UniDistill.

References

- [1] Garrick Brazil and Xiaoming Liu. M3d-rpn: Monocular 3d region proposal network for object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9287–9296, 2019. 1
- [2] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229, 2020. 2
- [3] Guobin Chen, Wongun Choi, Xiang Yu, Tony Han, and Manmohan Chandraker. Learning efficient object detection models with knowledge distillation. *Advances in neural information processing systems*, 30, 2017. 4
- [4] Qi Chen, Lin Sun, Ernest Cheung, and Alan L Yuille. Every view counts: Cross-view consistency in 3d object detection with hybrid-cylindrical-spherical voxelization. *Advances in Neural Information Processing Systems*, 33:21224–21235, 2020. 6
- [5] Xiaozhi Chen, Huimin Ma, Ji Wan, Bo Li, and Tian Xia. Multi-view 3d object detection network for autonomous driving. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 1907–1915, 2017. 2
- [6] Zhiyu Chong, Xinzhu Ma, Hong Zhang, Yuxin Yue, Haojie Li, Zhihui Wang, and Wanli Ouyang. Monodistill: Learning spatial features for monocular 3d object detection. In *International Conference on Learning Representations*, 2021. 1, 2, 3, 5
- [7] 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. 3
- [8] Hamidreza Fazlali, Yixuan Xu, Yuan Ren, and Bingbing Liu. A versatile multi-view framework for lidar-based 3d object detection with guidance from panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17192–17201, 2022. 6
- [9] Jianyuan Guo, Kai Han, Yunhe Wang, Han Wu, Xinghao Chen, Chunjing Xu, and Chang Xu. Distilling object detectors via decoupled features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2154–2164, 2021. 3
- [10] Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2(7), 2015. 3
- [11] Yuenan Hou, Zheng Ma, Chunxiao Liu, Tak-Wai Hui, and Chen Change Loy. Inter-region affinity distillation for road marking segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12486–12495, 2020. 3
- [12] Yihan Hu, Zhuangzhuang Ding, Runzhou Ge, Wenxin Shao, Li Huang, Kun Li, and Qiang Liu. Afdetv2: Rethinking the necessity of the second stage for object detection from point clouds. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 969–979, 2022. 6
- [13] Junjie Huang, Guan Huang, Zheng Zhu, and Dalong Du. Bevdet: High-performance multi-camera 3d object detection in bird-eye-view. *arXiv preprint arXiv:2112.11790*, 2021. 1, 2, 3, 5
- [14] Tengpeng Huang, Zhe Liu, Xiwu Chen, and Xiang Bai. Eppnet: Enhancing point features with image semantics for 3d object detection. In *European Conference on Computer Vision*, pages 35–52, 2020. 2
- [15] Bo Ju, Zhikang Zou, Xiaoqing Ye, Minyue Jiang, Xiao Tan, Errui Ding, and Jingdong Wang. Paint and distill: Boosting 3d object detection with semantic passing network. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 5639–5648, 2022. 1, 2, 3
- [16] Jason Ku, Melissa Mozifian, Jungwook Lee, Ali Harakeh, and Steven L Waslander. Joint 3d proposal generation and object detection from view aggregation. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1–8. IEEE, 2018. 2
- [17] 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/CVF conference on computer vision and pattern recognition*, pages 12697–12705, 2019. 2, 4
- [18] Bo Li. 3d fully convolutional network for vehicle detection in point cloud. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1513–1518. IEEE, 2017. 1
- [19] Yanwei Li, Yilun Chen, Xiaojuan Qi, Zeming Li, Jian Sun, and Jiaya Jia. Unifying voxel-based representation with transformer for 3d object detection. In *Advances in Neural Information Processing Systems*, 2022. 2
- [20] Yinhao Li, Zheng Ge, Guanyi Yu, Jinrong Yang, Zengran Wang, Yukang Shi, Jianjian Sun, and Zeming Li. Bevdepth: Acquisition of reliable depth for multi-view 3d object detection. *arXiv preprint arXiv:2206.10092*, 2022. 1, 2, 3
- [21] 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. 3
- [22] Zhijian Liu, Haotian Tang, Alexander Amini, Xinyu Yang, Huizi Mao, Daniela Rus, and Song Han. Bevfusion: Multi-task multi-sensor fusion with unified bird’s-eye view representation. *arXiv preprint arXiv:2205.13542*, 2022. 1, 2, 4, 5
- [23] Zhe Liu, Xin Zhao, Tengpeng Huang, Ruolan Hu, Yu Zhou, and Xiang Bai. Tanet: Robust 3d object detection from point clouds with triple attention. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 11677–11684, 2020. 2
- [24] Shujie Luo, Hang Dai, Ling Shao, and Yong Ding. M3dssd: Monocular 3d single stage object detector. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6145–6154, 2021. 1
- [25] Jiageng Mao, Minzhe Niu, Haoyue Bai, Xiaodan Liang, Hang Xu, and Chunjing Xu. Pyramid r-cnn: Towards better performance and adaptability for 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2723–2732, 2021. 2

- [26] Jiageng Mao, Yujing Xue, Minzhe Niu, Haoyue Bai, Jiashi Feng, Xiaodan Liang, Hang Xu, and Chunjing Xu. Voxel transformer for 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3164–3173, 2021. 2
- [27] Xuran Pan, Zhuofan Xia, Shiji Song, Li Erran Li, and Gao Huang. 3d object detection with pointformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7463–7472, 2021. 2
- [28] Dennis Park, Rares Ambrus, Vitor Guizilini, Jie Li, and Adrien Gaidon. Is pseudo-lidar needed for monocular 3d object detection? In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3142–3152, 2021. 2
- [29] Jonah Philion and Sanja Fidler. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In *European Conference on Computer Vision*, pages 194–210, 2020. 2, 4
- [30] Charles R Qi, Wei Liu, Chenxia Wu, Hao Su, and Leonidas J Guibas. Frustum pointnets for 3d object detection from rgb-d data. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 918–927, 2018. 1, 2
- [31] 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. 3, 4
- [32] 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. 2
- [33] Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. Point-rcnn: 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
- [34] Shaoshuai Shi, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. From points to parts: 3d object detection from point cloud with part-aware and part-aggregation network. *IEEE transactions on pattern analysis and machine intelligence*, 43(8):2647–2664, 2020. 1
- [35] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9627–9636, 2019. 2
- [36] Sourabh Vora, Alex H Lang, Bassam Helou, and Oscar Beijbom. Pointpainting: Sequential fusion for 3d object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4604–4612, 2020. 1, 3, 6
- [37] Chunwei Wang, Chao Ma, Ming Zhu, and Xiaokang Yang. Pointaugmenting: Cross-modal augmentation for 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11794–11803, 2021. 1
- [38] Tai Wang, Xinge Zhu, Jiangmiao Pang, and Dahua Lin. Fcos3d: Fully convolutional one-stage monocular 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 913–922, 2021. 2
- [39] Yan Wang, Wei-Lun Chao, Divyansh Garg, Bharath Hariharan, Mark Campbell, and Kilian Q Weinberger. Pseudo-lidar from visual depth estimation: Bridging the gap in 3d object detection for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8445–8453, 2019. 2
- [40] Yue Wang, Vitor Campagnolo Guizilini, Tianyuan Zhang, Yilun Wang, Hang Zhao, and Justin Solomon. Detr3d: 3d object detection from multi-view images via 3d-to-2d queries. In *Conference on Robot Learning*, pages 180–191. PMLR, 2022. 2
- [41] Yan Yan, Yuxing Mao, and Bo Li. Second: Sparsely embedded convolutional detection. *Sensors*, 18(10):3337, 2018. 2
- [42] Zetong Yang, Yanan Sun, Shu Liu, and Jiaya Jia. 3dssd: Point-based 3d single stage object detector. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11040–11048, 2020. 1
- [43] 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, 2, 4, 5
- [44] Jin Hyeok Yoo, Yecheol Kim, Jisong Kim, and Jun Won Choi. 3d-cvf: Generating joint camera and lidar features using cross-view spatial feature fusion for 3d object detection. In *European Conference on Computer Vision*, pages 720–736, 2020. 2
- [45] Linfeng Zhang and Kaisheng Ma. Improve object detection with feature-based knowledge distillation: Towards accurate and efficient detectors. In *International Conference on Learning Representations*, 2020. 3
- [46] Wu Zheng, Mingxuan Hong, Li Jiang, and Chi-Wing Fu. Boosting 3d object detection by simulating multimodality on point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13638–13647, 2022. 1, 2, 3, 6
- [47] Yin Zhou and Oncel Tuzel. Voxelnet: 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. 2, 4