

SEED: A Simple and Effective 3D DETR in Point Clouds

Zhe Liu^{1*}, Jinghua Hou^{1*}, Xiaoqing Ye², Tong Wang², Jingdong Wang², and Xiang Bai^{1†}

¹ Huazhong University of Science and Technology
{zhe1iu1994,jhhou,xbai}@hust.edu.cn

² Baidu Inc., China

Abstract. Recently, detection transformers (DETRs) have gradually taken a dominant position in 2D detection thanks to their elegant framework. However, DETR-based detectors for 3D point clouds are still difficult to achieve satisfactory performance. We argue that the main challenges are twofold: 1) How to obtain the appropriate object queries is challenging due to the high sparsity and uneven distribution of point clouds; 2) How to implement an effective query interaction by exploiting the rich geometric structure of point clouds is not fully explored. To this end, we propose a **S**imple and **E**ffective 3D **DE**TR method (**SEED**) for detecting 3D objects from point clouds, which involves a dual query selection (DQS) module and a deformable grid attention (DGA) module. More concretely, to obtain appropriate queries, DQS first ensures a high recall to retain a large number of queries by the predicted confidence scores and then further picks out high-quality queries according to the estimated quality scores. DGA uniformly divides each reference box into grids as the reference points and then utilizes the predicted offsets to achieve a flexible receptive field, allowing the network to focus on relevant regions and capture more informative features. Extensive ablation studies on DQS and DGA demonstrate its effectiveness. Furthermore, our SEED achieves state-of-the-art detection performance on both the large-scale Waymo and nuScenes datasets, illustrating the superiority of our proposed method. *The code is available at <https://github.com/happinesslz/SEED>.*

Keywords: Point Clouds · 3D object detection · Detection transformers

1 Introduction

DEtection TRansformer (DETR) [3] is the pioneering end-to-end transformer-based detector, which redefines object detection as a set prediction problem and eliminates hand-crafted anchors and non-maximum suppression (NMS) post-processing. These superior characteristics make the DETR paradigm [3, 19, 23, 56] become the mainstream method for 2D object detection tasks.

* Equal contribution.

† Corresponding author.

However, although many efforts have been made towards the DETR paradigm for 3D object detection [1, 7, 29, 30, 45, 46, 61] in point clouds, they have not demonstrated stunning performance similar to the 2D domain and still fall behind state-of-the-art 3D detectors [9, 11, 36, 43]. The main reason is the huge gap between 2D images and 3D points (*i.e.*, dense and regular 2D images *v.s.* sparse and irregular 3D points clouds), which requires us to carry out special designs for two critical components (*i.e.*, query selection and query interaction) in the DETR paradigm. For query selection, some methods [1, 7, 61] mainly select Top-N (*e.g.*, N=200, 300 or 1000) features as queries from the score map. Although effective, these methods do not consider the quality of the selected queries for box localization. For query interaction, some works [1, 46] perform several attention operations to achieve sufficient feature interaction. However, these approaches do not sufficiently take advantage of geometric information of 3D objects from point clouds.

In this paper, to alleviate the above challenges, we propose a **S**imple and **E**ffEctive 3D **D**ETR method (**SEED**) for detecting 3D objects from point clouds. The first key design in our SEED is the proposed dual query selection (DQS) to pick out high-quality queries in a coarse-to-fine manner, which includes a foreground query selection and a quality query selection. This manner is different from existing methods by one-step query selection [1, 61]. More concretely, to ensure a high recall, we first retain a large number of foreground queries in the foreground query selection according to the estimated confidence scores from a mask predictor. Then, we employ a SEED decoder layer to allow these queries to effectively interact with Bird’s Eye View (BEV) features. The enhanced queries are fed into the stage of quality query selection to pick out high-quality queries.

The second core design in our SEED is the proposed deformable grid attention (DGA) to make the network focus on relevant regions and achieve more effective feature interaction. Specifically, to exploit the rich geometric information in point clouds, we first divide the referenced box estimated by a regression branch into uniform grids, whose corresponding features can be easily collected to describe the geometric structures of 3D objects. To alleviate the strong depen-




Fig. 1: Comparison with DETR-based detectors [1, 7, 61] and other representative methods [8, 36] on Waymo *validation* dataset [39] in terms of detection performance and running speed. For a fair comparison, we evaluate the running speed of all approaches on an NVIDIA GeForce RTX 3090 with a batch size of 1. *-S*, *-B* and *-L* means the small, base, large versions of our SEED, respectively.

dence on the high-accuracy reference box, we further use these sampling grids as reference points and apply the predicted offsets to obtain flexible receptive fields. This enables the network to focus on surrounding regions of interest, even for less precise reference boxes.

As shown in Figure 1, we compare our SEED with the existing DETR-based 3D detection methods and other representative methods [8, 36] on the Waymo *validation* dataset [39] in terms of performance and running speed. It can be clearly observed that our SEED-S (*i.e.*, the small version) not only surpasses existing DETR-based approaches in detection performance but also maintains a superior running speed. In summary, our contributions are as follows:

- We introduce a novel dual query selection module, producing high-quality queries in a coarse-to-fine manner.
- We adopt an effective deformable grid attention module, which adaptively aggregates crucial regions and performs informative query interaction by properly leveraging the geometric information of point clouds.
- The proposed SEED achieves state-of-the-art performance for 3D object detection on both the large-scale Waymo [39] and nuScenes [2] datasets.

2 Related Work

2D Object Detection with DETR. DETR [3] is an end-to-end object detector that takes objects as queries and utilizes the transformer to interact queries with image features. Besides, DETR abandons many hand-crafted operations (*e.g.*, Anchor, NMS) and utilizes Hungarian Matching to achieve the ground-truth assignment. The elegant architecture proposed by DETR brings a new insight into the research of object detection, and many works [12, 19, 23, 28, 44, 56, 62] improve DETR from different perspectives. Deformable DETR [62] introduces deformable attention into DETR and greatly improves the convergence speed of DETR. DN-DETR [19] proposes the denoising training strategy, which effectively reduces the learning difficulty of bipartite graph matching. DINO [56] utilizes contrastive learning in denoising training to achieve better performance.

LiDAR-based 3D Object Detection. 3D object detectors in point clouds can be categorized into point-based and voxel-based categories. For point-based, most methods [4, 14, 16, 24, 31, 37, 50–52, 57] directly utilize a PointNet-like backbone [32, 33] to extract point features, which can keep precise geometric structure information. However, these methods usually need to sample points to reduce computational costs, which may lose some important information in point clouds. For voxel-based, most methods [9, 13, 18, 25, 26, 35, 38, 47–49, 53, 55, 59] quantify point clouds into regular grids and utilize a 3D sparse convolution backbone to extract grid features (*e.g.* Voxel and Pillar) efficiently.

3D Object Detection with DETR. Due to the powerful feature representation of transformer, many works [1, 7, 45, 46, 61] have been explored to utilize DETR for 3D object detection in point clouds, especially for the design of two key components (query selection and query interaction) in DETR. Specifically, TransFusion [1] selects the local maximum feature in BEV features as queries

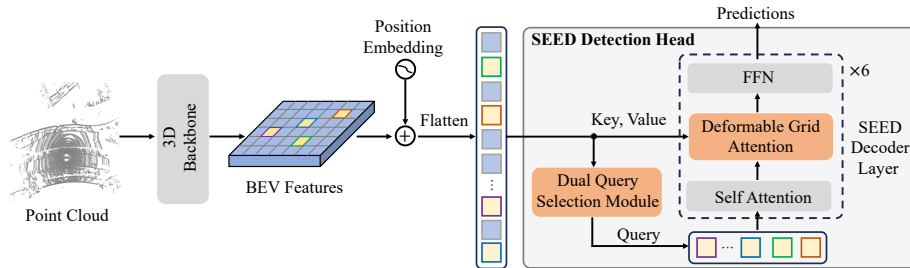


Fig. 2: Overall architecture of SEED, which consists of a 3D backbone and a SEED detection head. Specifically, the proposed SEED detection head mainly includes a dual query selection (DQS) module and a transformer decoder. The DQS utilizes a coarse-to-fine query selection strategy to select high-quality queries. The transformer decoder, including six SEED decoder layers, takes these queries as inputs and then iteratively performs a self-attention operation for inter-query interaction and a proposed deformable grid attention (DGA) for feature interaction between query and BEV features, generating final detection results.

based on the heatmap. CMT [46] adopts learnable queries initialized by 3D grids and utilizes global attention to interact queries with BEV features. Con-QueR [61] proposes a query contrast mechanism to reduce false positives. FocalFormer3D [7] utilizes a multi-stage heatmap for better query selection. Besides, FocalFormer3D [7] adopts deformable attention for efficient query interaction. Although the above DETR-based methods have made some progress, they are still inferior to some advanced methods [8, 43] that do not belong to the DETR-based paradigm. In this paper, we propose a simple and effective 3D DETR named SEED, involving a novel dual query selection module for picking out high-quality queries and a deformable grid attention module to make effective query interaction by leveraging the rich geometric information of point clouds.

3 Method

Although many attempts have been made on DETR-based 3D object detection, there is still a certain performance gap with existing advanced LiDAR-based 3D detectors [8, 43]. We argue that the main challenges come from two aspects. On the one hand, selecting superior queries from the high sparsity and uneven distribution of point clouds is not trivial. On the other hand, exploring how to make use of the rich geometric structure information from point clouds to perform effective query interaction is still challenging.

To mitigate these issues, we propose a **S**imple and **E**ffective 3D **DETR** method (**SEED**) for detecting 3D objects from point clouds. As shown in Figure 2, we present the overall pipeline of SEED. Specifically, we first feed point clouds into a classic voxel-based 3D backbone [47] to extract 3D voxel features and further convert them to BEV features. To retain their position information,

we add position embedding to the BEV features. Then, the BEV features are flattened for subsequent query selection. As for query selection, we propose a novel dual query selection (DQS), which adopts a coarse-to-fine manner to obtain high-quality queries. Finally, the transformer decoder, including six SEED decoder layers, is adopted to achieve feature interaction between the high-quality queries and the flattened BEV features, producing final detection results. In particular, our SEED decoder layer leverages an effective deformable grid attention (DGA) for query interaction instead of the cross attention operation in classic DETR decoder [3]. In the following, we will introduce the details of the proposed DQS and DGA in SEED.

3.1 Dual Query Selection Module




Fig. 3: Illustration of dual query selection (DQS). DQS adopts a coarse-to-fine manner, which consists of a foreground query selection and a quality query selection. S_c , S_l , and B_c are the predicted classification score, localization score, and regression for proposal boxes through three feed-forward networks (FFN) branches, respectively.

A proper query selection has demonstrated its importance in DETR-based 2D object detectors [5, 23, 56] to ensure accurate object localization and accelerate model convergence. However, due to the huge difference in data format between 2D images and 3D point clouds, it is necessary to consider some characteristics of point clouds, such as high sparsity and uneven distribution, for the query selection. Toward this goal, we propose a novel dual query selection (DQS) module whose main purpose is to obtain high-quality queries in a coarse-to-fine manner. We present the detailed structure of DQS in Figure 3, which involves a foreground query selection and a quality query selection.

Foreground Query Selection. First, for the foreground query selection, we utilize a binary classification predictor to distinguish backgrounds and foregrounds on the BEV features. Simultaneously, we add BEV position embedding to BEV features and flatten them along the spatial dimension to generate all queries (also named flattened BEV features). For the convenience of description, we define the flattened BEV features as $\mathbf{F}_{bev} \in \mathbb{R}^{(H \times W) \times C}$, where H , W and C are the height, width, and channel dimension of BEV features, respectively. Then, we select these queries with proportion r among the top confidence scores of BEV features \mathbf{S}_{bev} from a mask predictor as the coarse queries, which can

remain as many potential foreground queries as possible to ensure a high recall rate. Finally, the foreground query selection is formulated as:

$$\mathbf{Q}_c = \text{Top}_{N_c}(\mathbf{F}_{bev}, \mathbf{S}_{bev}), \quad (1)$$

where $N_c = H \times W \times r$ and $\text{Top}_{N_c}(x, y)$ means select top N_c queries from x according to y , and N_c is the number of coarse queries.

After getting coarse queries, we further feed them to a SEED decoder layer to achieve sufficient feature interaction between queries and flattened BEV features, producing the enhanced queries \mathbf{Q}'_c , which can be computed as:

$$\mathbf{Q}'_c = \text{Decoder}(\mathbf{Q}_c, \mathbf{F}_{bev}), \quad (2)$$

where Decoder is our SEED Decoder layer, which will be introduced in detail in the Section 3.2.

Quality Query Selection. We first feed the coarse queries \mathbf{Q}'_c into three feed-forward networks (FFN) branches to produce the classification score \mathbf{S}_c , the localization scores \mathbf{S}_l and the regression \mathbf{B}_c for coarse proposal boxes, whose corresponding ground truths are assigned based on the proposed quality-aware Hungarian Matching (refer to Section 3.3). Here, the classification score is the probability of recognizing 3D object proposals, and the localization score is defined as the 3D IoU of proposal boxes and the ground truths. Considering that the localization score is mainly for foreground objects, we set a proper classification score threshold τ to distinguish the foreground objects. Therefore, the quality scores \mathbf{S}_q by combining these two indicators can be formulated as:

$$\mathbf{S}_q^i = \begin{cases} (\mathbf{S}_c^i)^{1-\beta} \cdot (\mathbf{S}_l^i)^\beta, & \text{if } \mathbf{S}_c^i > \tau, \\ \mathbf{S}_c^i, & \text{otherwise,} \end{cases} \quad (3)$$

where $\beta \in (0, 1)$ is a hyper-parameter and is applied to control the importance of classification score and localization score, and $i = 0, 1, \dots, N_c$. Then, we select the top N_f fine proposal boxes \mathbf{B}_f according to the quality scores \mathbf{S}_q and concatenate them with the corresponding box quality scores \mathbf{S}_f . Next, we feed the concatenated features in a multi-layer perceptron (MLP) to generate the geometric-aware high-quality queries. These steps can be formulated as:

$$\mathbf{B}_f = \text{Top}_{N_f}(\mathbf{B}_c, \mathbf{S}_q), \quad (4)$$

$$\mathbf{Q}_f = \text{MLP}(\text{Concat}(\mathbf{B}_f, \mathbf{S}_f)). \quad (5)$$

Finally, the output queries \mathbf{Q}_f of DQS will be input to the subsequent SEED decoder.

3.2 SEED Decoder Layer

The proposed SEED decoder layer is applied to further enhance query feature representation by a self-attention operation and a cross attention operation and




Fig. 4: Illustration of deformable grid attention (DGA). DGA first uniformly divides each reference proposal into grids as the reference points and then utilizes the predicted offsets to achieve a flexible receptive field.

then maps the enhanced queries to task-specified outputs by an FFN. Different from existing DETR-based methods [7, 46, 61], we adopt a new cross attention operation in our SEED decoder layer, namely deformable grid attention (DGA). Next, we will introduce the reasons and the details of DGA.

Why need DGA? An effective query interaction designed for point clouds is necessary to further explore the potential of the DETR paradigm in 3D detection. First of all, unlike 2D images, a nearby object may occupy most of the whole image, which even requires a global receptive field to detect the object well. However, a 3D object usually only occupies a small local area (also be mentioned in SST [10]), which is much smaller than the range of the entire point clouds. Thus, the local attention may be enough for query interaction in point clouds. Second, point clouds possess rich geometric structures, especially for regular vehicles. Therefore, it is important to rationally utilize the geometric information of 3D objects. Third, although an accurate 3D proposal box can describe the geometric information of 3D objects, it is sub-optimal to capture some irregular objects or some hard objects. This indicates that a flexible receptive field is desired. Towards this goal, we propose a deformable grid attention (DGA), which is a new local attention but adopts a flexible receptive field to effectively leverage the geometric information of 3D objects for query interaction.

Details of DGA. As shown in Figure 4, we present the detailed structure of DGA. Specifically, we first regard the estimated proposal boxes B_f from DQS as the reference boxes and uniformly divide each reference box into $k \times k$ grids g_k (*i.e.*, yellow points in Figure 4). Then, we feed the corresponding selected queries Q_f (*i.e.*, red points in Figure 4) from DQS into a linear function, producing the predicted offsets Δg . Next, we add the offsets to the grids g to generate the final sampling positions, which can capture the geometric information of 3D objects in a flexible receptive field. Meanwhile, the attention weight A is predicted by feeding Q_f into a linear function and a softmax function. Finally, the sampled features are multiplied with A to obtain the enhanced queries. We can formulate the above process of DGA as follows:

$$\text{DGA}(\mathbf{g}, \mathbf{F}_{bev}) = \sum_{j=1}^K \mathbf{A}_j \cdot \phi(\mathbf{F}_{bev}(\mathbf{g}_j + \Delta \mathbf{g}_j)), \quad (6)$$

where $K = k \times k$ and ϕ is a linear function for transforming sampled features to the attention space. \mathbf{F}_{bev}^* denotes sampling the corresponding features of the positions $*$ on the BEV features \mathbf{F}_{bev} by a bilinear interpolation operation. Besides, we provide more comparisons with different attention operations [30,62] in our supplemental materials.

3.3 Quality-aware Hungarian Matching

Different from the traditional Hungarian Matching [3], we introduce a quality-aware Hungarian Matching (QHM) to effectively assign the ground truth. Specifically, QHM adopts the quality scores \mathbf{S}_f instead of the classic classification scores in computing classification cost. Thus, our classification cost \mathcal{C}_{cls} can be formulated as:

$$\mathcal{C}_{pos} = -(1 - \alpha) \cdot (\mathbf{S}_f)^\gamma \cdot \log(1 - \mathbf{S}_f), \quad (7)$$

$$\mathcal{C}_{neg} = -\alpha \cdot (1 - \mathbf{S}_f)^\gamma \cdot \log \mathbf{S}_f, \quad (8)$$

$$\mathcal{C}_{cls} = \mathcal{C}_{pos} - \mathcal{C}_{neg}, \quad (9)$$

where α, γ are hyper-parameters. Finally, the total matching cost in Hungarian Matching can be computed as:

$$\mathcal{C}_{match} = \lambda_{cls} \cdot \mathcal{C}_{cls} + \lambda_{reg} \cdot \mathcal{C}_{reg} + \lambda_{giou} \cdot \mathcal{C}_{giou}, \quad (10)$$

where \mathcal{C}_{reg} and \mathcal{C}_{giou} denote the regression cost and the GIoU cost, which have the same formulation in the traditional Hungarian Matching [3]. λ_{cls} , λ_{reg} , and λ_{giou} are the balanced weights.

4 Experiments

4.1 Datasets and Evaluation Metrics

Waymo Open Dataset. Waymo Open Dataset (WOD) [39] is a well-known and challenging large-scale outdoor 3D object detection benchmark. WOD consists of 1150 scenes (more than 200K frames), which are divided into three parts: 798 for training, 202 for validation, and 150 for testing. Each scene provides point clouds acquired by 64-beam LiDAR and covers a perception range with a size of $150\text{m} \times 150\text{m}$. Besides, the evaluation of WOD is divided into two levels (LEVEL 1 and LEVEL 2) according to the number of points of the object. For all experimental results, we follow the standard protocol as the evaluation metric, which adopts 3D mean Average Precision (mAP) [22] and its weighted variant by heading accuracy (mAPH) for three categories: *Vehicle*, *Pedestrian* and *Cyclist*.

nuScenes Dataset. nuScenes [2] is another widely used autonomous driving dataset for LiDAR-based 3D object detection. The nuScenes dataset consists of 1000 scenes, which are divided into three parts: 750 for training, 150 for validation, and 150 for testing. Each scene is roughly 20s long, annotated at 2Hz, and

provides point clouds collected by 32-beam LiDAR. Besides, the evaluation of the nuScenes dataset adopts the Mean Average Precision (mAP) and nuScenes detection score (NDS) to evaluate the performance of 3D detectors for 10 foreground classes.

4.2 Implementation Details

Network Architecture. In SEED, we provide three versions, namely small (*i.e.*, SEED-S), base (*i.e.*, SEED-B), and large (*i.e.*, SEED-L). In this paper, SEED-S adopts the same 3D backbone as CenterPoint [54] for a fair comparison with most existing methods. For SEED-B, we follow VoxelNext [8], which introduces multi-scale voxel feature extraction and doubles the channel of the 3D backbone based on SEED-S to improve the feature representation. For SEED-L, we utilize a smaller voxel size (0.08, 0.08, 0.15) instead of (0.1, 0.1, 0.15) to enlarge the BEV resolution based on SEED-B, which further boosts detection performance specifically for small objects such as *Pedestrian*. In DQS, we set $r = 0.3$ in foreground query selection to ensure a high recall and set $N_f = 1000, \tau = 0.2$ in quality query selection for effective query interaction in the subsequent transformer decoder. In DGA, we set $k = 5$ to divide grids as default. In the transformer decoder, we adopt six SEED decoder layers to iteratively perform query interaction and then generate the final 3D detection results.

Training. The final loss includes the DETR-head loss (similar to ConQueR [61]) and our DQS loss. In DQS loss, we supervise classification score with binary cross-entropy loss, localization score with IoU loss, and regression with Smooth-L1 loss, respectively. On WOD, we adopt the same point cloud range and data augmentations as CenterPoint [54]. Our model is optimized by AdamW optimizer [27] with the initial learning rate, weight decay, and momentum factor set to 0.001, 0.01, and 0.9, respectively. We train our model with a batch size of 24 on 8 NVIDIA Tesla V100 GPUs. We run 24 epochs for the 20% training set and only 12 epochs for the 100% training set. Besides, we utilize the fade strategy [42] to avoid over-fitting in the last epoch and the query contrast strategy [61] to achieve better performance. For quality-aware Hungarian Matching, we set α to 0.25 and γ to 2.0. In quality query selection, the hyper-parameter β is set to 0.68, 0.71, and 0.65 for *Vehicle*, *Pedestrian*, and *Cyclist*, respectively. This is a common setting in some non-DETR methods (*e.g.*, AFDetv2 [17], PillarNet [34]) for IoU-rectification. For matching cost, we set λ_{cls} , λ_{reg} , and λ_{giou} to 1, 2, and 4. For nuScenes, our SEED simply follows the settings of [1, 7], including the range of point clouds, the voxel size, data augmentations, and the training strategy.

4.3 Main Results

Results on WOD. We present the comparison with existing DETR-based methods (bottom) and other representative methods (top) on the WOD in Table 1. Here, we provide three versions of SEED, including the small (SEED-S), base (SEED-B) and large (SEED-L). Compared with DETR-based 3D detectors,

Table 1: Performances on the Waymo Open Dataset *validation* split (train with 100% training data). ‡ denotes the two-stage method. ‘SEED-S’, ‘SEED-B’, and ‘SEED-L’ mean the small, base, and large versions of SEED, respectively. **Bold** denotes the best performance in DETR-based methods. All results are presented with single-frame input, no test-time augmentation, and no model ensembling.

Methods	Present at	DETR	Vehicle 3D AP/APH		Pedestrian 3D AP/APH		Cyclist 3D AP/APH		mAP/mAPH L2
			L1	L2	L1	L2	L1	L2	
SECOND [47]	Sensors 18		72.3/71.7	63.9/63.3	68.7/58.2	60.7/51.3	60.6/59.3	58.3/57.0	61.0/57.2
PointPillars [18]	CVPR 19		72.1/71.5	63.6/63.1	70.6/56.7	62.8/50.3	64.4/62.3	61.9/59.9	62.8/57.8
CenterPoint [54]	CVPR 21		74.2/73.6	66.2/65.7	76.6/70.5	68.8/63.2	72.3/71.1	69.7/68.5	68.2/65.8
PV-RCNN‡ [35]	CVPR 20		78.0/77.5	69.4/69.0	79.2/73.0	70.4/64.7	71.5/70.3	69.0/67.8	69.6/67.2
SST_TS‡ [10]	CVPR 22		76.2/75.8	68.0/67.6	81.4/74.0	72.8/65.9	-/-	-/-	-/-
AFDetV2 [17]	AAAI 22		77.6/77.1	69.7/69.2	80.2/74.6	72.2/67.0	73.7/72.7	71.0/70.1	71.0/68.8
SWFormer [40]	ECCV 22		77.8/77.3	69.2/68.8	80.9/72.7	72.5/64.9	-/-	-/-	-/-
PillarNet-34 [34]	ECCV 22	×	79.1/78.6	70.9/70.5	80.6/74.0	72.3/66.2	72.3/71.2	69.7/68.7	77.3/74.6
CenterFormer [60]	ECCV 22		75.0/74.4	69.9/69.4	78.6/73.0	73.6/68.3	72.3/71.3	69.8/68.8	71.1/68.9
PV-RCNN++‡ [36]	IJCV 22		79.3/78.8	70.6/70.2	81.3/76.3	73.2/68.0	73.7/72.7	71.2/70.2	71.7/69.5
FSD‡ [11]	NeurIPS 22		79.2/78.8	70.5/70.1	82.6/77.3	73.9/69.1	77.1/76.0	74.4/73.3	72.9/70.8
OcTr [58]	CVPR 23		78.1/77.6	69.8/69.3	80.8/74.4	72.5/66.5	72.6/71.5	69.9/68.9	70.7/68.2
PillarNeXt [20]	CVPR 23		78.4/77.9	70.3/69.8	82.5/77.1	74.9/69.8	73.2/72.2	70.6/69.6	71.9/69.7
VoxelNext [8]	CVPR 23		78.2/77.7	69.9/69.4	81.5/76.3	73.5/68.6	76.1/74.9	73.3/72.2	72.2/70.1
DSVT-Pillar [43]	CVPR 23		79.3/78.8	70.9/70.5	82.8/77.0	75.2/69.8	76.4/75.4	73.6/72.7	73.2/71.0
DSVT-Voxel [43]	CVPR 23		79.7/79.3	71.4/71.0	83.7/78.9	76.1/71.5	77.5/76.5	74.6/73.7	74.0/72.1
BoxeR-3D [30]	CVPR 22		70.4/70.0	63.9/63.7	64.7/53.5	61.5/53.7	50.2/48.9	-/-	-/-
TransFusion [1]	CVPR 22		-/-	-/65.1	-/-	-/63.7	-/-	-/65.9	-/64.9
ConQueR [61]	CVPR 23		76.1/75.6	68.7/68.2	79.0/72.3	70.9/64.7	73.9/72.5	71.4/70.1	70.3/67.7
FocalFormer3D [7]	ICCV 23	✓	-/-	68.1/67.6	-/-	72.7/66.8	-/-	73.7/72.6	71.5/69.0
SEED-S (Ours)	-		78.2/77.7	70.2/69.7	81.3/75.8	73.3/68.1	78.4/77.2	75.7/74.5	73.1/70.8
SEED-B (Ours)	-		79.7/79.2	71.8/71.4	83.1/78.3	75.5/70.8	80.0/78.8	77.3/76.1	74.9/72.8
SEED-L (Ours)	-		79.8/79.3	71.9/71.5	83.6/79.1	76.2/71.8	81.2/80.0	78.4/77.3	75.5/73.5

our SEED-S outperforms the advanced ConQueR [61] and FocalFormer3D [7] with 3.1 and 1.8 mAPH/L2, respectively. Moreover, SEED-S possesses a satisfactory running speed with about 13.5 FPS on an NVIDIA GeForce RTX 3090 (see Figure 1). These benefits on both the detection performance and the running speed effectively illustrate the superiority of SEED. Besides, we compare our SEED with other advanced methods that do not belong to the DETR paradigm. SEED-B exceeds the representative two-stage methods PV-RCNN++ [36] with 3.3 mAPH/L2. Considering that the high-resolution feature map is necessary to detect small 3D objects, SEED-L further boosts the detection performance of SEED-B with 1.0 (71.8 *vs.* 70.8) APH/L2 on *Pedestrian*. It is noteworthy that SEED-L even outperforms the previous state-of-the-art (SOTA) method DSVT-Voxel [43] with 1.4 (73.5 *vs.* 72.1) mAPH/L2, leading to a new SOTA. Note that the BEV resolution of SEED-L is still much smaller than DSVT-Voxel [43] since DSVT uses one-stride 3D backbone like SST [10] to improve the detection performance. Furthermore, DSVT [43] focuses on enhancing the representation ability of the 3D backbone, which is orthogonal to our SEED.

In Table 2a, we also provide the results of our SEED with three frames as inputs on the WOD *validation* split. We observe that SEED-L produces a leading performance with 76.1 mAPH/L2, even surpassing the temporal 3D object detection method MPPNet [6] with an obvious margin. To further verify the effectiveness of our SEED, we evaluate the performance of our SEED-L with one frame and three frames as inputs on the WOD *test* benchmark, as shown

Table 2: Effectiveness of our SEED with multiple frames as inputs on the Waymo Open Dataset *validation* and *test* split.

Methods	Frames	mAP/mAPH (L1)	mAP/mAPH (L2)	Methods	Frames	mAP/mAPH (L1)	mAP/mAPH (L2)
CenterPoint [54]	4	76.4/74.9	70.8/69.4	PV-RCNN++ [36]	1	78.0/75.7	72.4/70.2
CenterFormer [60]	4	78.5/77.0	74.7/73.2	AFDetV2 [17]	1	77.6/75.2	72.2/70.3
MPPNet [6]	4	81.1/79.9	75.4/74.2	PillarNet [34]	1	77.5/74.7	72.2/69.6
MSF [15]	4	81.1/80.2	76.0/74.6	FSD [11]	1	80.4/78.2	74.4/72.4
PillarNeXt [34]	3	81.5/80.0	75.9/74.5	ConQueR [61]	1	—/—	—/72.0
DSVT-Voxel [43]	3	82.1/80.8	76.3/75.0	SEED-L (Ours)	1	81.7/79.7	76.5/74.5
SEED-S (Ours)	3	81.6/80.1	75.8/74.3	CenterPoint++ [54]	3	79.4/77.9	74.2/72.8
SEED-B (Ours)	3	82.9/81.4	77.2/75.8	PillarNeXt [20]	3	80.5/79.0	75.5/74.1
SEED-L (Ours)	3	83.1/81.6	77.5/76.1	SEED-L (Ours)	3	83.5/82.1	78.7/77.3

(a) Effectiveness of SEED on *validation* split.**(b)** Effectiveness of SEED on *test* benchmark.**Table 3:** Comparison with state-of-the-art methods on the nuScenes *validation* split. All results are presented without any test-time augmentation or model ensembling. Here, our SEED adopts the same 3D backbone with CenterPoint [54] for a fair comparison. * denotes the reproduced result from official code.

Method	Present at	mATE	mASE	mAOE	mAVE	mAAE	NDS	mAP
PointPillar [18]	CVPR 19	0.424	0.284	0.529	0.377	0.194	49.1	34.3
CenterPoint [54]	CVPR 21	0.291	0.252	0.324	0.284	0.189	64.9	56.6
TransFusion-L [1]	CVPR 22	—	—	—	—	—	70.1	65.1
PillarNet [34]	ECCV 22	0.277	0.252	0.289	0.247	0.191	67.4	59.8
UVTR-L [21]	NeurIPS 22	0.334	0.257	0.300	0.204	0.182	67.7	60.9
VoxelNeXt* [8]	CVPR 23	0.301	0.252	0.406	0.217	0.186	66.7	60.5
Uni3DETR [45]	NeurIPS 23	0.288	0.249	0.303	0.216	0.181	68.5	61.7
SEED (Ours)	—	0.279	0.257	0.284	0.208	0.187	71.2	66.6

in Table 2b. SEED-L with single frame achieves 74.5 mAPH/L2, which exceeds DETR-based method ConQueR [61] with 2.5 (74.5 *vs.* 72.0) mAPH/L2. For three frames, SEED-L outperforms the representative method PillarNeXt [20] with 3.2 (77.3 *vs.* 74.1) mAPH/L2. All the experimental results clearly demonstrate the superiority of our SEED.

Results on nuScenes. We also evaluate our SEED on the *validation* split of nuScenes dataset [2] to further verify the effectiveness of our SEED. As shown in Table 3, SEED achieves 71.2 NDS and 66.6 mAP, which exceeds the popular DETR-based detector TransFusion-L [1] with 1.1 NDS and 1.5 mAP under the same 3D backbone. This demonstrates the generalization of our method.

4.4 Ablation Study

In this section, we conduct extensive ablation studies to investigate the effectiveness of SEED on the Waymo *validation* set with 20% training data by default if not specified. And we adopt SEED-S as our default SEED model in the following ablation studies. For more ablation studies about our SEED, please refer to our supplemental materials.

Effectiveness of the Proposed Components. As shown in Table 4, we conduct ablation studies on our proposed two components in SEED. First, we set the baseline by replacing our DQS with heatmap-based query selection [1] and

Table 4: Ablation study for each component in SEED. We use mAP/mAPH (L2) to evaluate the overall detection performance.

DQS	DGA	3D AP/APH (L2)			mAP/mAPH (L2)
		<i>Vehicle</i>	<i>Pedestrian</i>	<i>Cyclist</i>	
–	–	65.4/64.9	68.8/63.3	66.7/65.5	67.0/64.6
✓	–	68.0/67.5	70.9/65.2	70.7/69.5	69.9/67.4
–	✓	65.7/65.2	69.9/64.4	71.0/69.6	68.8/66.4
✓	✓	68.5/68.1	72.1/66.5	71.2/70.0	70.6/68.2

Table 5: Ablation study for DQS in SEED. We adopt mAP/mAPH (L2) to evaluate the detection performance.

Methods	3D AP/APH (L2)			mAP/mAPH (L2)
	<i>Vehicle</i>	<i>Pedestrian</i>	<i>Cyclist</i>	
Learnable [46]	66.1/65.6	70.1/64.5	71.0/69.8	69.1/66.6
Heatmap-based [1]	64.3/63.8	69.7/64.1	68.4/67.1	67.5/65.0
Top-N [61]	66.1/65.7	70.6/65.3	70.7/69.5	69.1/66.8
DQS (Ours)	68.5/68.1	72.1/66.5	71.2/70.0	70.6/68.2

replacing our DGA with box attention [30] in our SEED. Compared with the baseline, the proposed DQS module brings 2.8 mAPH/L2 (67.4 *vs.* 64.6) performance improvement in average for *Vehicle*, *Pedestrian*, and *Cyclist* on Waymo dataset, which demonstrates the superiority of DQS for selecting out high-quality queries. Besides, the proposed DGA boosts the performance of the baseline with 1.8 mAPH/L2, demonstrating the effectiveness of DGA in achieving feature interaction. Finally, thanks to the benefits of the proposed DQS and DGA, SEED has an obvious gain of 3.6 mAPH/L2 over the baseline model.

Superiority of the DQS. To further demonstrate the superiority of the proposed DQS module, we provide three representative query selection strategies for comparison with DQS in Table 5, namely *Learnable*, *Heatmap-based* and *Top-N* manners. Specifically, *Learnable* means that we obtain queries by adopting a learnable manner like CMT [46]. *Heatmap-based* (*e.g.*, TransFusion-L [1]) is a classic query selection strategy in LiDAR-based 3D object detection, which collects the local maximum elements in BEV features as queries. The advanced query selection manner *Top-N* (*e.g.*, ConQueR [61]) uses a class-agnostic feed-forward network (FFN) head to obtain Top-N scored box proposals, which are selected as object queries. As shown in Table 5, *Heatmap-based* produces the worst performance among these methods. We think the main reason is that the queries (as Query) of *Heatmap-based* method is obtained directly from BEV features (as Key, Value), rather than learnable queries or geometric-aware queries. This often leads to sub-optimal feature interaction, which makes it difficult for the subsequent decoder to stack more layers due to the potential risk of overfitting. However, our DQS respectively outperforms the *Learnable* and *Top-N* manners with 1.6 and 1.4 mAPH/L2, which illustrates the superiority of dual query selection for picking out high-quality geometric-aware queries.




Fig. 5: Visualization of SEED without DQS (the first row) and with DQS (the second row). We highlight the challenging queries with red circles. The colormap indicates the values of the confidence scores for selected queries on the BEV map. Green boxes are the ground truth boxes.

Table 6: Ablation study for DGA in SEED. We adopt mAP/mAPH (L2) to evaluate the detection performance.

Methods	3D AP/APH (L2)			mAP/mAPH (L2)
	<i>Vehicle</i>	<i>Pedestrian</i>	<i>Cyclist</i>	
Global Attention [41]	– (OOM)	– (OOM)	– (OOM)	– (OOM)
Deformable Attention [62]	67.5/67.0	71.3/65.9	70.8/69.7	69.9/67.5
Box Attention [30]	67.9/67.4	71.1/65.5	70.9/69.7	70.0/67.5
DGA (Ours)	68.5/68.1	72.1/66.5	71.2/70.0	70.6/68.2

Effectiveness of the DGA. To illustrate the effectiveness of the proposed DGA, we compare it with three representative query interaction operations, including global attention [41], deformable attention [62] and box attention (or named non-deformable grid attention) [30], whose results are summarized in Table 6. Since most LiDAR-based 3D detection methods [1, 54] have larger feature maps (*e.g.*, 180×180), performing global attention operation on mainstream GPUs (*e.g.*, NVIDIA Tesla V100 or NVIDIA GeForce RTX 3090) is unable to bear such a large computational cost, causing the GPU to run out of memory (OOM). Therefore, the local attention operation, including the deformable attention and the box attention, is still reasonable to promise performance and efficiency for query interaction. For deformable attention operation, it produces 67.0 APH/L2 on *Vehicle*, which is inferior to the box attention manner with 67.4 APH/L2. This indicates that effectively leveraging the box geometric information is important in feature interaction. However, the box attention operation depends on the precision of box regression, and its receptive field is not as flexible as deformable attention. Therefore, on hard objects such as *Pedestrian*, the deformable attention outperforms the detection performance of the box attention. In contrast, our DGA has the advantages of both the flexible receptive field of deformable attention and the rich geometric information of box attention. Not surprisingly, our DGA surpasses these local attention methods, demonstrating its effectiveness.

Table 7: Ablation study for quality-aware Hungarian Matching (QHM) in SEED. THM is short for traditional Hungarian Matching. Besides, we adopt mAP/mAPH (L2) to evaluate the detection performance.

Methods	3D AP/APH (L2)			mAP/mAPH (L2)
	<i>Vehicle</i>	<i>Pedestrian</i>	<i>Cyclist</i>	
THM [3]	67.3/66.8	71.7/66.4	70.9/69.7	70.0/67.6
QHM (Ours)	68.5/68.1	72.1/66.5	71.2/70.0	70.6/68.2

Influence of Quality-aware Hungarian Matching. Hungarian Matching is an indispensable component in the existing DETR-based approaches. In SEED, we adopt a new matching manner, named quality-aware Hungarian Matching (QHM). To verify the influence of QHM on 3D detection performance, we compare it with the traditional Hungarian Matching [3] (THM), whose result is presented in Table 7. It can be observed that QHM produces a consistent performance improvement over THM in terms of mAPH/L2, which benefits from taking the quality scores of 3D objects into account when computing classification cost in Hungarian Matching. Moreover, we carefully find there is a better gain on *Vehicle* than *Pedestrian* and *Cyclist*. The main reason is that the localization score of *Vehicle* is more easily estimated than that of *Pedestrian* and *Cyclist* due to some factors such as a large size and a rigid object.

4.5 Visualization Analysis

As shown in Figure 5, we present the visualization of our SEED with and without DQS (*i.e.*, directly select Top N_f queries in one step). Specifically, we visualize the positions of final selected queries from query selection on the BEV map, whose corresponding colormap denotes the confidence score of the selected queries. It can be observed that after utilizing DQS, some hard queries are successfully captured, which indicates that DQS can enhance the confidence score of some potential hard objects. This phenomenon further demonstrates the superiority of our DQS. Besides, we provide more visualization of our SEED in the supplemental materials.

5 Conclusion

In this paper, we have presented a simple and effective 3D DETR framework named SEED to detect 3D objects from point clouds. Specifically, SEED involves two key components: a dual query selection (DQS) module to retain high-quality queries in a coarse-to-fine manner and a deformable grid attention (DGA) module to capture informative features by performing sufficient query interaction. Extensive ablation studies have demonstrated the effectiveness of the proposed DQS and DGA. Thanks to the superiority of the proposed DQS and DGA, our SEED has achieved state-of-the-art 3D detection performance on the large-scale Waymo and nuScenes dataset. Finally, we hope SEED could become a new strong baseline for the community of DETR-based 3D object detection.

Acknowledgements

This work was supported by the National Science Fund for Distinguished Young Scholars of China (Grant No. 62225603) and the Hubei Key R&D Program (Grant No. 2022BAA078).

References

1. Bai, X., Hu, Z., Zhu, X., Huang, Q., Chen, Y., Fu, H., Tai, C.L.: Transfusion: Robust lidar-camera fusion for 3d object detection with transformers. In: CVPR. pp. 1090–1099 (2022)
2. Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O.: nuscenes: A multimodal dataset for autonomous driving. In: CVPR. pp. 11621–11631 (2020)
3. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: ECCV. pp. 213–229. Springer (2020)
4. Chen, C., Chen, Z., Zhang, J., Tao, D.: Sasa: Semantics-augmented set abstraction for point-based 3d object detection. In: AAAI. vol. 36, pp. 221–229 (2022)
5. Chen, F., Zhang, H., Hu, K., Huang, Y.K., Zhu, C., Savvides, M.: Enhanced training of query-based object detection via selective query recollection. In: CVPR. pp. 23756–23765 (2023)
6. Chen, X., Shi, S., Zhu, B., Cheung, K.C., Xu, H., Li, H.: Mppnet: Multi-frame feature intertwining with proxy points for 3d temporal object detection. In: ECCV. pp. 680–697. Springer (2022)
7. Chen, Y., Yu, Z., Chen, Y., Lan, S., Anandkumar, A., Jia, J., Alvarez, J.M.: Focalformer3d: Focusing on hard instance for 3d object detection. In: ICCV. pp. 8394–8405 (2023)
8. Chen, Y., Liu, J., Zhang, X., Qi, X., Jia, J.: Voxelnext: Fully sparse voxelnet for 3d object detection and tracking. In: CVPR. pp. 21674–21683 (2023)
9. Deng, J., Shi, S., Li, P., Zhou, W., Zhang, Y., Li, H.: Voxel r-cnn: Towards high performance voxel-based 3d object detection. In: AAAI. vol. 35, pp. 1201–1209 (2021)
10. Fan, L., Pang, Z., Zhang, T., Wang, Y.X., Zhao, H., Wang, F., Wang, N., Zhang, Z.: Embracing single stride 3d object detector with sparse transformer. In: CVPR (2022)
11. Fan, L., Wang, F., Wang, N., Zhang, Z.: Fully sparse 3d object detection. In: NeurIPS (2022)
12. Gao, P., Zheng, M., Wang, X., Dai, J., Li, H.: Fast convergence of detr with spatially modulated co-attention. In: ICCV. pp. 3621–3630 (2021)
13. Guan, T., Wang, J., Lan, S., Chandra, R., Wu, Z., Davis, L., Manocha, D.: M3detr: Multi-representation, multi-scale, mutual-relation 3d object detection with transformers. In: WACV. pp. 772–782 (2022)
14. He, C., Li, R., Li, S., Zhang, L.: Voxel set transformer: A set-to-set approach to 3d object detection from point clouds. In: CVPR. pp. 8417–8427 (2022)
15. He, C., Li, R., Zhang, Y., Li, S., Zhang, L.: Msf: Motion-guided sequential fusion for efficient 3d object detection from point cloud sequences. In: CVPR. pp. 5196–5205 (2023)
16. He, C., Zeng, H., Huang, J., Hua, X.S., Zhang, L.: Structure aware single-stage 3d object detection from point cloud. In: CVPR. pp. 11873–11882 (2020)

17. Hu, Y., Ding, Z., Ge, R., Shao, W., Huang, L., Li, K., Liu, Q.: Afdetv2: Rethinking the necessity of the second stage for object detection from point clouds. In: AAAI (2022)
18. Lang, A.H., Vora, S., Caesar, H., Zhou, L., Yang, J., Beijbom, O.: Pointpillars: Fast encoders for object detection from point clouds. In: CVPR (2019)
19. Li, F., Zhang, H., Liu, S., Guo, J., Ni, L.M., Zhang, L.: Dn-detr: Accelerate detr training by introducing query denoising. In: CVPR (2022)
20. Li, J., Luo, C., Yang, X.: Pillarnext: Rethinking network designs for 3d object detection in lidar point clouds. In: CVPR. pp. 17567–17576 (2023)
21. Li, Y., Chen, Y., Qi, X., Li, Z., Sun, J., Jia, J.: Unifying voxel-based representation with transformer for 3d object detection. In: NeurIPS. vol. 35, pp. 18442–18455 (2022)
22. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: ECCV. pp. 740–755 (2014)
23. Liu, S., Li, F., Zhang, H., Yang, X., Qi, X., Su, H., Zhu, J., Zhang, L.: Dab-detr: Dynamic anchor boxes are better queries for detr. In: ICLR (2022)
24. Liu, Z., Huang, T., Li, B., Chen, X., Wang, X., Bai, X.: Epnet++: Cascade bi-directional fusion for multi-modal 3d object detection. In: IEEE TPAMI. vol. 45, pp. 8324–8341. IEEE (2022)
25. Liu, Z., Zhao, X., Huang, T., Hu, R., Zhou, Y., Bai, X.: Tanet: Robust 3d object detection from point clouds with triple attention. In: AAAI. vol. 34, pp. 11677–11684 (2020)
26. Liu, Z., Yang, X., Tang, H., Yang, S., Han, S.: Flatformer: Flattened window attention for efficient point cloud transformer. In: CVPR. pp. 1200–1211 (2023)
27. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: arXiv preprint arXiv:1711.05101 (2017)
28. Meng, D., Chen, X., Fan, Z., Zeng, G., Li, H., Yuan, Y., Sun, L., Wang, J.: Conditional detr for fast training convergence. In: ICCV (2021)
29. Misra, I., Girdhar, R., Joulin, A.: An end-to-end transformer model for 3d object detection. In: ICCV. pp. 2906–2917 (2021)
30. Nguyen, D.K., Ju, J., Booi, O., Oswald, M.R., Snoek, C.G.: Boxer: Box-attention for 2d and 3d transformers. In: CVPR. pp. 4773–4782 (2022)
31. Qi, C.R., Litany, O., He, K., Guibas, L.J.: Deep hough voting for 3d object detection in point clouds. In: ICCV (2019)
32. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In: CVPR. pp. 652–660 (2017)
33. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In: NeurIPS. pp. 5099–5108 (2017)
34. Shi, G., Li, R., Ma, C.: Pillarnet: High-performance pillar-based 3d object detection. In: ECCV (2022)
35. Shi, S., Guo, C., Jiang, L., Wang, Z., Shi, J., Wang, X., Li, H.: Pv-rcnn: Point-voxel feature set abstraction for 3d object detection. In: CVPR (2020)
36. Shi, S., Jiang, L., Deng, J., Wang, Z., Guo, C., Shi, J., Wang, X., Li, H.: Pv-rcnn++: Point-voxel feature set abstraction with local vector representation for 3d object detection. IJCV (2021)
37. Shi, S., Wang, X., Li, H.: Pointrcnn: 3d object proposal generation and detection from point cloud. In: CVPR (2019)
38. Shi, S., Wang, Z., Shi, J., Wang, X., Li, H.: From points to parts: 3d object detection from point cloud with part-aware and part-aggregation network. IEEE TPAMI **43**(8), 2647–2664 (2020)

39. Sun, P., Kretzschmar, H., Dotiwalla, X., Chouard, A., Patnaik, V., Tsui, P., Guo, J., Zhou, Y., Chai, Y., Caine, B., et al.: Scalability in perception for autonomous driving: Waymo open dataset. In: CVPR. pp. 2446–2454 (2020)
40. Sun, P., Tan, M., Wang, W., Liu, C., Xia, F., Leng, Z., Anguelov, D.: Swformer: Sparse window transformer for 3d object detection in point clouds. In: ECCV (2022)
41. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. In: NeurIPS. vol. 30 (2017)
42. Wang, C., Ma, C., Zhu, M., Yang, X.: Pointaugmenting: Cross-modal augmentation for 3d object detection. In: CVPR. pp. 11794–11803 (2021)
43. Wang, H., Shi, C., Shi, S., Lei, M., Wang, S., He, D., Schiele, B., Wang, L.: Dsvt: Dynamic sparse voxel transformer with rotated sets. In: CVPR. pp. 13520–13529 (2023)
44. Wang, Y., Zhang, X., Yang, T., Sun, J.: Anchor detr: Query design for transformer-based detector. In: AAAI (2022)
45. Wang, Z., Li, Y.L., Chen, X., Zhao, H., Wang, S.: Uni3detr: Unified 3d detection transformer. In: NeurIPS (2023)
46. Yan, J., Liu, Y., Sun, J., Jia, F., Li, S., Wang, T., Zhang, X.: Cross modal transformer: Towards fast and robust 3d object detection. In: CVPR. pp. 18268–18278 (2023)
47. Yan, Y., Mao, Y., Li, B.: Second: Sparsely embedded convolutional detection. Sensors **18**(10), 3337 (2018)
48. Yang, H., Liu, Z., Wu, X., Wang, W., Qian, W., He, X., Cai, D.: Graph r-cnn: Towards accurate 3d object detection with semantic-decorated local graph. In: ECCV. pp. 662–679. Springer (2022)
49. Yang, H., Wang, W., Chen, M., Lin, B., He, T., Chen, H., He, X., Ouyang, W.: Pvt-ssd: Single-stage 3d object detector with point-voxel transformer. In: CVPR. pp. 13476–13487 (2023)
50. Yang, J., Song, L., Liu, S., Mao, W., Li, Z., Li, X., Sun, H., Sun, J., Zheng, N.: Dbq-ssd: Dynamic ball query for efficient 3d object detection. In: ICLR (2022)
51. Yang, Z., Sun, Y., Liu, S., Jia, J.: 3dssd: Point-based 3d single stage object detector. In: CVPR. pp. 11040–11048 (2020)
52. Yang, Z., Sun, Y., Liu, S., Shen, X., Jia, J.: Std: Sparse-to-dense 3d object detector for point cloud. In: CVPR. pp. 1951–1960 (2019)
53. Yang, Z., Zhou, Y., Chen, Z., Ngiam, J.: 3d-man: 3d multi-frame attention network for object detection. In: CVPR. pp. 1863–1872 (2021)
54. Yin, T., Zhou, X., Krahenbuhl, P.: Center-based 3d object detection and tracking. In: CVPR (2021)
55. Zhang, G., Junnan, C., Gao, G., Li, J., Hu, X.: Hednet: A hierarchical encoder-decoder network for 3d object detection in point clouds. In: NeurIPS. vol. 36 (2024)
56. Zhang, H., Li, F., Liu, S., Zhang, L., Su, H., Zhu, J., Ni, L.M., Shum, H.Y.: Dino: Detr with improved denoising anchor boxes for end-to-end object detection. In: ICLR (2023)
57. Zhang, Y., Hu, Q., Xu, G., Ma, Y., Wan, J., Guo, Y.: Not all points are equal: Learning highly efficient point-based detectors for 3d lidar point clouds. In: CVPR. pp. 18953–18962 (2022)
58. Zhou, C., Zhang, Y., Chen, J., Huang, D.: Octr: Octree-based transformer for 3d object detection. In: CVPR. pp. 5166–5175 (2023)
59. Zhou, Y., Tuzel, O.: Voxelnet: End-to-end learning for point cloud based 3d object detection. In: CVPR (2018)

60. Zhou, Z., Zhao, X., Wang, Y., Wang, P., Foroosh, H.: Centerformer: Center-based transformer for 3d object detection. In: ECCV (2022)
61. Zhu, B., Wang, Z., Shi, S., Xu, H., Hong, L., Li, H.: Conquer: Query contrast voxel-detr for 3d object detection. In: CVPR. pp. 9296–9305 (2023)
62. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. In: ICLR (2021)