

D-Align: Dual Query Co-attention Network for 3D Object Detection Based on Multi-frame Point Cloud Sequence

Junhyung Lee¹, Junho Koh², Youngwoo Lee² and Jun Won Choi^{2,*}

Abstract—LiDAR sensors are widely used for 3D object detection in various mobile robotics applications. LiDAR sensors continuously generate point cloud data in real-time. Conventional 3D object detectors detect objects using a set of points acquired over a fixed duration. However, recent studies have shown that the performance of object detection can be further enhanced by utilizing spatio-temporal information obtained from point cloud sequences. In this paper, we propose a new 3D object detector, named *D-Align*, which can effectively produce strong bird’s-eye-view (BEV) features by aligning and aggregating the features obtained from a sequence of point sets. The proposed method includes a novel dual-query co-attention network that uses two types of queries, including target query set (T-QS) and support query set (S-QS), to update the features of target and support frames, respectively. D-Align aligns S-QS to T-QS based on the temporal context features extracted from the adjacent feature maps and then aggregates S-QS with T-QS using a gated fusion mechanism. The dual queries are updated through multiple attention layers to progressively enhance the target frame features used to produce the detection results. Our experiments on the nuScenes dataset show that the proposed D-Align method greatly improved the performance of a single frame-based baseline method and significantly outperformed the latest 3D object detectors. Code is available at <https://github.com/junhyung-SPALab/D-Align>.

I. INTRODUCTION


LiDAR is a widely used sensor modality for perception tasks in various mobile robotics applications. LiDAR sensors generate point cloud data corresponding to observations of laser reflections from the surfaces of objects. Point cloud data are particularly useful for 3D object detection tasks, which estimate objects’ locations in 3D coordinate systems. Recently, LiDAR-based 3D object detection has advanced rapidly with the adoption of deep neural networks to extract features from point cloud data.

Conventional 3D object detectors operate on a set of LiDAR points, called *point clouds*, which are acquired by a fixed number of consecutive laser scans. The geometrical distribution of point clouds is used to detect objects in 3D space. In practical applications, LiDAR sensors continuously scan their surroundings to generate point cloud sequences in real time. Therefore, the quality of point cloud data can be improved by using more than one set of point clouds for 3D object detection. One naïve approach is to merge multiple consecutive sets of point clouds acquired in the fixed duration

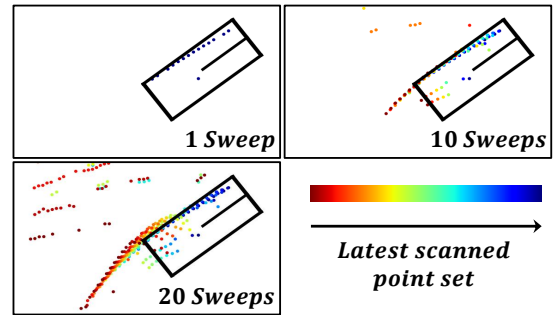
¹Department of Future Mobility, Hanyang University, Seoul, 04763, Korea. junhyunglee@spa.hanyang.ac.kr

²Department of Electrical Engineering, Hanyang University, Seoul, 04763, Korea. {jkhkoh, youngwoolee}@spa.hanyang.ac.kr, junwchoi@hanyang.ac.kr

* Corresponding author



(a)



(b)

Fig. 1. Effects of using larger LiDAR point sets for 3D object detection: (a) Evaluation results of the PointPillars [11] method on the nuScenes [4] validation set. NDS denotes the nuScenes detection score. The NDS performance improvement of the 3D object detector quickly diminished as the size of LiDAR point set increased. (b) Object motion leads to dynamic changes in the distribution of point clouds as more LiDAR points are merged.

and use the larger set of points as input to a 3D object detection model [4]. This effectively improves the density of the point clouds and thus improves performance, as shown in Fig. 1 (a). However, this performance improvement quickly diminishes with increasing numbers of point clouds merged because the distribution of the set of points changes across multiple scans. Fig. 1 (b) shows that a moving object exhibits a dynamically changing point distribution even after compensating for the ego vehicle’s motion that occurred in the duration. Therefore, more sophisticated 3D object detection models that capture the dynamic temporal structure of LiDAR points are needed to improve accuracy.

Several recent studies have explored methods to utilize temporal information existing in long sequences of point clouds for 3D object detection. In this study, we refer to these detectors as *3D multi-frame object detectors* (3D-MOD). Various architectures have been proposed for 3D-MOD [9], [18], [21], [22], [24]. These methods extracted geometric features from each set of points (called a *frame*) within which

the point distribution did not change much and modeled their temporal variations to improve performance. In LSTM-TOD [9] and 3DVID [22], the multiple bird’s-eye-view (BEV) feature maps extracted from the point cloud sequence were processed by the variant of the recurrent module, i.e., ConvLSTM [17] and ConvGRU [2]. TCTR [24] exploited temporal-channel relations over multiple feature maps using an encoder-decoder structure [16]. 3D-MAN [21] aggregated box-level features with a memory bank containing sequential temporal view information.

In this paper, we present a new 3D-MOD method called *Dual-Query Align (D-Align)*, which can produce robust spatio-temporal BEV representations using a multi-frame point cloud sequence. We propose a novel *dual query co-attention network* that employs two types of queries: *target query set* (T-QS) and *support query set* (S-QS) to facilitate the co-attention to both the target and support frame features. T-QS and S-QS serve to carry the target frame features and the support frame features to enhance the feature quality through multiple stages of co-attention mechanism.

In each co-attention layer, the dual queries are updated in two steps. First, the *inter-frame deformable alignment network* (IDANet) aligns S-QS to T-QS using deformable attention. The deformable attention mechanism is applied to S-QS with mask offsets and weights determined by multi-scale temporal context features generated from two adjacent BEV feature maps. This step yields S-QS that is aligned with T-QS. Next, the *inter-frame gated aggregation network* (IGANet) aggregates S-QS and T-QS using a gated fusion network [10]. The aggregated query features finally update T-QS. After going through multiple attention layers, D-Align produces the improved BEV features of the target frame, which are consumed by the detection head for 3D object detection.

We evaluated the proposed D-Align on the widely used public nuScenes dataset [4]. Our experimental results show that the proposed method improves the 3D detection baseline by significant margins and outperforms the latest LiDAR-based state-of-the-art (SOTA) 3D object detectors.

The contributions of our paper are summarized as follows.

- We propose an enhanced 3D object detection architecture, D-Align, which exploits temporal structure in point cloud sequences. We devise a novel dual query co-attention network that transforms the dual queries S-QS and T-QS through successive operations of feature alignment and feature aggregation. This co-attention mechanism allows attending to both the support and target frame features to gather useful spatio-temporal information from multiple frames of the point data.
- We design the temporal context-guided deformable attention to achieve inter-frame feature alignment. Our deformable attention mechanism differs from the original model proposed in [29] in that the attention mask is adjusted by the motion context obtained from two adjacent BEV feature maps. Our analysis shows that the use of such motion features contributes significantly to the overall detection performance.

II. RELATED WORK

3D object detection techniques based on a single point set [7], [11], [19], [23], [26] have advanced rapidly since deep neural networks have been adopted to encode irregular and unordered point clouds. However, the performance of these 3D object detectors remains limited because they do not exploit temporal information in sequence data.

To date, several 3D-MOD methods have been proposed, which used point cloud sequences to perform 3D object detection [9], [21], [22], [24]. These methods explored ways to represent time-varying features obtained from long sequence of point clouds. In [9], [22], [24], LiDAR features obtained from multiple point cloud frames were combined to exploit temporal information in sequence data. LSTM-TOD [9] produced a spatio-temporal representation of point cloud sequences using a 3D sparse ConvLSTM [17] modified to encode the point cloud sequence data. 3DVID [22] improved the conventional ConvGRU by adopting a transformer attention mechanism to exploit the spatio-temporal coherence of point cloud sequences. TCTR [24] explored channel-wise temporal correlations among consecutive frames and decoded spatial information using Transformer. 3D-MAN [21] stored 3D proposals and 3D spatial features obtained from a single-frame 3D detector in a memory bank to align and aggregate the local features of objects extracted from each frame.

In our paper, we introduce a novel dual query co-attention architecture for utilizing spatio-temporal information obtained from point cloud sequences. The proposed method effectively aligns and aggregates multi-frame features through successive refinement of the dual query sets.

III. PROPOSED METHOD

Fig. 2 shows the overall structure of D-Align. It consists of three main blocks, including 1) BEV feature extractor, 2) *dual-query co-attention network* (DUCANet), and 3) 3D object detection head. The set of points acquired by the LiDAR sensor over the duration of T seconds is called a frame. Ego-motion compensation [4] is applied to the points within each frame. D-Align takes the sequence of point clouds $\{P_n\}_{n=t-N+1}^t$ in N successive frames as an input, where P_n denotes the point set obtained in the n th frame. The frame t is called a target frame because we aim to detect objects for the frame t . The remaining frames are called support frames.

A. Overview

The BEV feature extractor extracts N BEV feature maps $\{\mathbf{F}_{t-k}\}_{k=0}^{N-1}$ from the point sets in N successive frames using a grid-based backbone network [11], [19]. This backbone network produces the feature maps of S scales, i.e., $\mathbf{F}_{t-k} = \{\mathbf{F}_{t-k}^s\}_{s=1}^S$. Next, DUCANet produces the enhanced target frame BEV features by applying the dual-query co-attention mechanism to the multi-frame features $\{\mathbf{F}_{t-k}\}_{k=0}^{N-1}$. DUCANet maintains two query sets, T-QS and S-QS. T-QS serves to store the target frame features, whereas S-QS stores the support frame features. We denote T-QS and S-QS as $\mathbf{B}_t^{(l)}$ and $\{\mathbf{U}_{t-k}^{(l)}\}_{k=1}^{N-1}$, respectively, where l is the index for

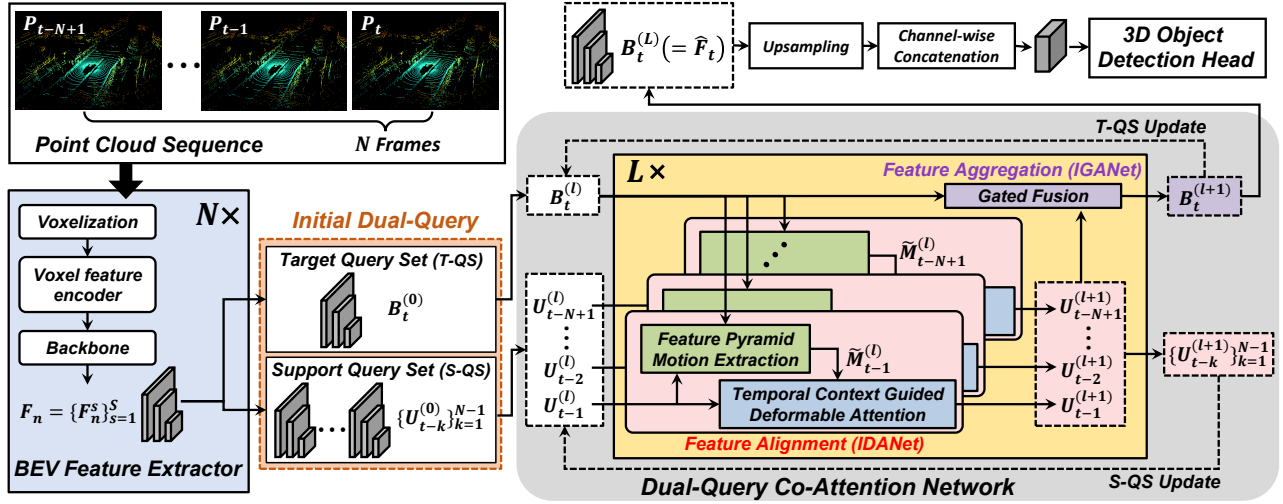


Fig. 2. **Overall architecture of the proposed D-Align:** D-Align employs a dual query co-attention network to align and aggregate multi-frame BEV features. T-QS and S-QS are dual queries used to refine the features of the target and support frames, respectively. For each attention layer, S-QS is refined by applying the temporal context-guided deformable attention to S-QS for inter-frame feature alignment. T-QS is also updated by aggregating the updated S-QS with T-QS via the gated fusion network. These updated dual queries are used as input queries for the next attention layer.

the attention layers. In the first attention layer, T-QS $B_t^{(0)}$ is initialized by F_t . Similarly, S-QS $\{U_{t-k}^{(0)}\}_{k=1}^{N-1}$ is initialized by $\{F_{t-k}\}_{k=1}^{N-1}$. For each attention layer, T-QS and S-QS are refined by aligning and aggregating the multi-frame features through IDANet and IGANet.

First, IDANet aligns S-QS to T-QS by applying the temporal context-guided deformable attention to S-QS. Deformable attention [29] has been widely used to spatially transform input features. The *feature pyramid motion extraction network* (FPMNet) constructs the temporal context features, $\tilde{M}_{t-k}^{(l)}$ from two query features $B_t^{(l)}$ and $U_{t-k}^{(l)}$. Then, IDANet determines the deformable mask based on the temporal context features $\{\tilde{M}_{t-k}^{(l)}\}_{k=1}^{N-1}$. S-QS is finally updated with the outputs of the deformable attention $\{U_{t-k}^{(l+1)}\}_{k=1}^{N-1}$.

Next, IGANet aggregates S-QS, which is refined by IDANet, with T-QS using the gated fusion network proposed in [10]. This step updates T-QS with $B_t^{(l+1)}$. Finally, these refined dual query sets $\{U_{t-k}^{(l+1)}\}_{k=1}^{N-1}$ and $B_t^{(l+1)}$ are used as input queries for the next attention layer.

After going through L attention layers, DUCANet comes up with the final T-QS, $B_t^{(L)}$. This T-QS is used as the multi-scale BEV feature maps $\hat{F}_t (= \{\hat{F}_t^s\}_{s=1}^S)$ to detect objects at frame t . After re-sizing these multi-scale features to have the same scale and concatenating them channel-wise, a 3D object detection head is applied to produce the final 3D detection results.

B. Inter-frame deformable alignment network (IDANet)

Fig. 3 depicts the structure of IDANet. In IDANet, FPMNet first extracts temporal context features from the adjacent multi-scale feature maps. Then, the proposed deformable attention is applied to perform an inter-frame feature alignment.

Feature pyramid motion extraction: The FPMNet operates on the dual query sets $B_t^{(l)}$ and $U_{t-k}^{(l)}$ to produce the

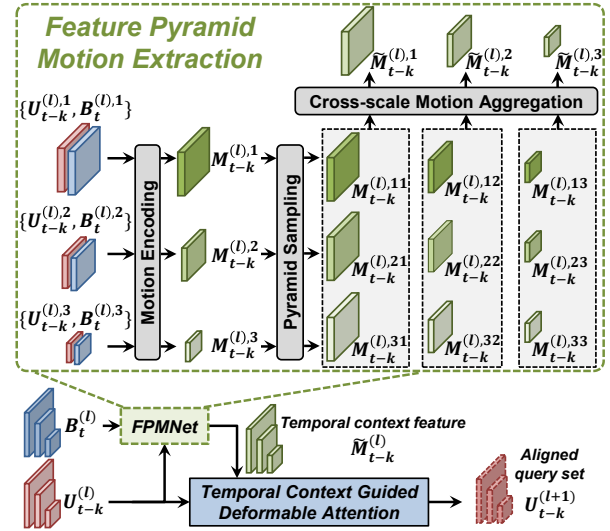


Fig. 3. **Structure of IDANet:** IDANet aligns S-QS using deformable attention. The temporal context information extracted from FPMNet is used to predict the sampling offsets and attention weights required for deformable attention.

temporal context features $\tilde{M}_{t-k}^{(l)}$. Recall that $B_t^{(l)}$ and $U_{t-k}^{(l)}$ represent the target frame features and the $(t-k)$ th support frame features updated up to the l th attention layer. First, FPMNet applies the motion encoding to $B_t^{(l)}$ and $U_{t-k}^{(l)}$ for each scale, i.e.,

$$M_{t-k}^{(l,s)} = \text{conv}_{3 \times 3}(B_t^{(l,s)} - U_{t-k}^{(l,s)}), \quad (1)$$

where $\text{conv}_{3 \times 3}$ denotes the convolutional layer with 3×3 kernels and $M_{t-k}^{(l,s)}$, $B_t^{(l,s)}$, and $U_{t-k}^{(l,s)}$ are the scale- s features of $M_{t-k}^{(l)}$, $B_t^{(l)}$, and $U_{t-k}^{(l)}$, respectively. Then, we aggregate the features across all scales as shown in Fig. 3. First, the pyramid sampling [13] transforms the features $M_{t-k}^{(l,s)}$ of the scale s into the features $M_{t-k}^{(l,ss')}$ of the scale s' . Then, the

final temporal context features $\tilde{\mathbf{M}}_{t-k}^{(l),s'}$ are obtained by fusing the set of the s' -scale features $\{\mathbf{M}_{t-k}^{(l),s,s'}\}_{s=1}^S$ as

$$\tilde{\mathbf{M}}_{t-k}^{(l),s'} = \text{conv}_{1 \times 1}([\mathbf{M}_{t-k}^{(l),1s'}, \mathbf{M}_{t-k}^{(l),2s'}, \dots, \mathbf{M}_{t-k}^{(l),Ss'}]), \quad (2)$$

where $[\cdot, \cdot, \dots, \cdot]$ denotes the channel-wise concatenation. The detailed process of extracting the temporal context features is described in Fig. 3. Finally, the output of FPMNet is the multi-scale temporal context features $\tilde{\mathbf{M}}_{t-k}^{(l)} = \{\tilde{\mathbf{M}}_{t-k}^{(l),s'}\}_{s'=1}^S$. Note that these feature maps are used to predict both the mask offsets and weights of the deformable attention for inter-frame feature alignment.

Temporal context guided deformable attention: For each attention layer, IDANet aligns S-QS to T-QS by applying the deformable attention to S-QS, $\mathbf{U}_{t-k}^{(l+1)}$ and generating the refined S-QS, $\mathbf{U}_{t-k}^{(l+1)}$. We employ the multi-head deformable attention, which supports H multi-head attentions. Let p_s be a two-dimensional reference position on the feature map of the scale s . Note that the reference position p_s is swept over the entire feature map to compute the attention values. The BEV embedding feature \mathbf{V}_{t-k}^{sh} for the h th head and the s th scale is obtained as

$$\mathbf{V}_{t-k}^{sh}(p_s) = W_h \mathbf{U}_{t-k}^{(l),s}(p_s), \quad (3)$$

where $W_h \in \mathbb{R}^{D' \times D}$ is the projection matrix and D' is (D/H) .

Next, the mask offsets $\Delta_{t-k}^{sh}(p_s)$ and the attention weights $A_{t-k}^{sh}(p_s)$ at the reference point p_s are given by

$$\Delta_{t-k}^{sh}(p_s) = W'_h \tilde{\mathbf{M}}_{t-k}^{(l),s}(p_s) \quad (4)$$

$$A_{t-k}^{sh}(p_s) = \text{softmax}(W''_h \tilde{\mathbf{M}}_{t-k}^{(l),s}(p_s)), \quad (5)$$

where $W'_h \in \mathbb{R}^{(2SJ) \times D}$ and $W''_h \in \mathbb{R}^{(SJ) \times D}$ are the linear projection matrices, J is the size of the mask, and $\text{softmax}(\cdot)$ is the softmax function. Note that $\Delta_{t-k}^{sh}(p_s)$ is the set $\{\Delta_{t-k}^{shij}(p_s) | 1 \leq i \leq S, 1 \leq j \leq J\}$, where $\Delta_{t-k}^{shij}(p_s)$ is the (x, y) mask offset of the j th sampling point at the i th scale. $A_{t-k}^{sh}(p_s)$ is defined similarly as $\{A_{t-k}^{shij}(p_s) | 1 \leq i \leq S, 1 \leq j \leq J\}$, except that each element has a scalar weight.

Given the mask offsets and weights, the attention value $z_{t-k}^s(p_s)$ can be obtained from

$$z_{t-k}^s(p_s) = \sum_{h=1}^H W_h''' \left(\sum_{i=1}^S \sum_{j=1}^J A_{t-k}^{shij}(p_s) \cdot \mathbf{V}_{t-k}^{ih}(\psi_i(p_s) + \Delta_{t-k}^{shij}(p_s)) \right), \quad (6)$$

where $W_h''' \in \mathbb{R}^{D \times D'}$ is the projection matrix and $\psi_i(\cdot)$ is a re-scaling function that adjusts the reference position p_s for the i th scale.

Finally, the refined S-QS $\mathbf{U}_{t-k}^{(l+1),s}$ is obtained as

$$\mathbf{U}_{t-k}^{(l+1),s}(p_s) = \text{FFN}(\text{LN}(\text{dropout}(z_{t-k}^s(p_s)) + \mathbf{U}_{t-k}^{(l),s}(p_s))), \quad (7)$$

where $\text{FFN}(\cdot)$ is the feed-forward network [29], $\text{dropout}(\cdot)$ is the dropout [15], and $\text{LN}(\cdot)$ is the layer normalization [1].

C. Inter-frame gated aggregation network (IGANet)

IGANet aggregates S-QS $\{\mathbf{U}_{t-k}^{(l+1)}\}_{k=1}^{N-1}$ and T-QS $\mathbf{B}_t^{(l)}$, producing the refined T-QS $\mathbf{B}_t^{(l+1)}$. IGANet applies the gated fusion network [10] to combine $\mathbf{U}_{t-k}^{(l+1)}$ and $\mathbf{B}_t^{(l)}$ for each scale

$$\tilde{\mathbf{B}}_{t-k}^{(l),s} = \mathbf{G}_{t-k}^{(l),s} \otimes \mathbf{B}_t^{(l),s} + (1 - \mathbf{G}_{t-k}^{(l),s}) \otimes \mathbf{U}_{t-k}^{(l+1),s}, \quad (8)$$

where the operation \otimes denotes element-wise multiplication and $\mathbf{G}_{t-k}^{(l),s}$ denotes the gating matrix that weights the contribution from S-QS and T-QS. The gating weights $\mathbf{G}_{t-k}^{(l),s}$ are obtained by applying the 3×3 convolutional layer to the concatenation of S-QS and T-QS,

$$\mathbf{G}_{t-k}^{(l),s} = \sigma(\text{conv}_{3 \times 3}([\mathbf{B}_t^{(l),s}, \mathbf{U}_{t-k}^{(l+1),s}])), \quad (9)$$

where $\sigma(\cdot)$ is the sigmoid function. Note that the gating matrix $\mathbf{G}_{t-k}^{(l),s}$ has a channel dimension of one and each element has a value between 0 and 1.

Finally, we aggregate the outputs $\tilde{\mathbf{B}}_{t-1}^{(l),s}, \dots, \tilde{\mathbf{B}}_{t-N+1}^{(l),s}$ of the gated fusion network and obtain the refined T-QS as

$$\mathbf{B}_t^{(l+1),s} = \text{conv}_{3 \times 3}([\tilde{\mathbf{B}}_{t-1}^{(l),s}, \dots, \tilde{\mathbf{B}}_{t-N+1}^{(l),s}]). \quad (10)$$

Putting $\mathbf{B}_t^{(l+1),s}$ together over all scales, IGANet produces the refined T-QS, $\mathbf{B}_t^{(l+1)}$.

IV. EXPERIMENTS

A. nuScenes Dataset

The nuScenes dataset [4] is a challenging large-scale autonomous driving dataset comprising 1000 scenes, each with a duration of about 20 seconds. The dataset provides 32-channel LiDAR scan samples acquired at 20Hz with a full 360-degree view. Moreover, 3D bounding boxes from 10 categories are annotated at 2Hz, i.e., annotated key samples are supplied in every 10 samples. We evaluated our method in terms of mean average precision (mAP) and nuScenes detection score (NDS), which are the official 3D object detection benchmark metrics of nuScenes. The mAP is computed based on the 2D center distance at the bird's-eye-view between the ground-truth data and the predictions. The NDS is a weighted sum of mAP and several true positive metrics, i.e., the average translation, scale, orientation, velocity, and attribute errors [4].

B. Implementation Details

Point Cloud Sequence Processing. Following the nuScenes dataset [4], a single frame comprised 10 consecutive LiDAR point cloud samples ($T = 0.5$) with ego-vehicle motion compensation applied. Specifically, each point in the single frame is represented with 5 dimensional vector as $(x, y, z, r, \Delta t)$, where r is intensity of the point and Δt indicates the time lag to the key sample, i.e., a scalar time stamp in the range of $[0, 0.5)$. Unless otherwise stated, $N = 3$ frames were used, in which the input sequence

TABLE I

QUANTITATIVE COMPARISON WITH STATE OF THE ART METHODS ON nuSCENES TEST SET. C.V AND T.C PRESENTS THE CONSTRUCTION VEHICLE AND THE TRAFFIC CONE, RESPECTIVELY. MOTOR. AND PED. ARE SHORT FOR THE MOTORCYCLE AND THE PEDESTRIAN, RESPECTIVELY. L INDICATES THE LiDAR. THE BEST PERFORMANCE IS BOLDED.

Method	Input	NDS	mAP	Car	Truck	Bus	Trailer	C.V	Ped.	Motor.	Bicycle	T.C	Barrier
PointPillars [11]	Single frame	45.3	30.5	68.4	23.0	28.2	23.4	4.1	59.7	27.4	1.1	30.8	38.9
3DSSD [20]		56.4	42.6	81.2	47.2	61.4	30.5	12.6	7.2	36.0	8.6	31.1	47.9
SA-Det3D [3]		59.2	47.0	81.2	43.8	57.2	47.8	11.3	73.3	32.1	7.9	60.6	55.3
SSN V2 [28]		61.6	50.6	82.4	41.8	46.1	48.0	17.5	75.6	48.9	24.6	60.1	61.2
CBGS [27]		63.3	52.8	81.1	48.5	54.9	42.9	10.5	80.1	51.5	22.3	70.9	65.7
CVCNet [5]		64.2	55.8	82.7	46.1	45.8	46.7	20.7	81.0	61.3	34.3	69.7	69.9
CyliNet [14]		66.1	58.5	85.0	50.2	56.9	52.6	19.1	84.3	58.6	29.8	79.1	69.0
HotSpotNet [6]		66.6	59.3	83.1	50.9	56.4	53.3	23.0	81.3	63.5	36.6	73.0	71.6
CenterPoint [23]		67.3	60.3	85.2	53.5	63.6	56.0	20.0	84.6	59.5	30.7	78.4	71.1
AFDet V2 [8]		68.5	62.4	86.3	54.2	62.5	58.9	26.7	85.8	63.8	34.3	80.1	71.0
S2M2-SSD [25]		69.3	62.9	86.3	56.0	65.4	59.8	26.2	84.5	61.6	36.4	77.7	75.1
UVTR-L [12]		69.7	63.9	86.3	52.2	62.8	59.7	33.7	84.5	68.8	41.1	74.7	74.9
3DVID [22]	Multi frame	53.1	45.4	79.7	33.6	47.1	43.1	18.1	76.5	40.7	7.9	58.8	48.8
TCTR [24]		-	50.5	83.2	51.5	63.7	33.0	15.6	74.9	54.0	22.6	52.5	53.8
STGNN [18]		66.7	59.0	86.2	57.2	67.2	35.5	14.6	85.5	58.1	37.0	71.3	66.2
Our D-Align-PP		63.1	53	82.2	44.0	53.1	53.5	19.8	75.8	51.3	17.6	67.4	65.5
Our D-Align-CP		70.2	64.0	85.0	50.5	67.0	60.2	27.5	84.7	73.6	43.3	81.8	66.0

consists of a single target frame and two previous support frames.

Architecture. We constructed the proposed D-Align based on two types of 3D object detector baselines, including anchor-based and anchor-free detectors. First, we chose PointPillars [11] as an anchor-based detector baseline. The point cloud range was set to $[-51.2, 51.2] \times [-51.2, 51.2] \times [-5.0, 3.0]$ meters on the (x, y, z) axes, and the grid size was $0.2 \times 0.2 \times 8.0m$. The PointPillars baseline generated the BEV feature maps of $S = 3$ scales, the sizes of which were $(D, H/2, W/2)$, $(D, H/4, W/4)$, and $(D, H/8, W/8)$. (H, W) denotes the size of the voxel grid structure in the x - y axis. Next, we chose CenterPoint [23] with a VoxelNet [26] backbone as an anchor-free detector baseline. The point cloud range was $[-54.0, 54.0] \times [-54.0, 54.0] \times [-5.0, 3.0]$ meters on the (x, y, z) axes, and the grid size was $0.075 \times 0.075 \times 0.2m$. It generated the BEV feature maps of $S = 2$ scales and their sizes were set to $(D, H/8, W/8)$ and $(D, H/16, W/16)$. In DUCANet, the number of attention layers was set to $L = 3$. In the deformable attention of IDANet, the number of attention heads H and sampling locations J were set to 8 and 4, respectively.

Training. We adopted the loss functions used in PointPillars [11] and CenterPoint [23] to train D-Aligns and used the Adam optimizer. D-Align was trained in two steps: 1) pre-training and 2) fine-tuning. In pre-training step, we trained a single frame detector including only the BEV feature extraction backbone and the 3D detection head. We used the one-cycle *learning rate* (LR) scheduling for 20 epochs with a maximum LR of 0.001. Then, we fine-tuned the entire network with the pre-trained weights initialized. In fine-tuning, we used the same LR scheduler for 40 epochs with a maximum LR of 0.0002. The batch size was set to 16 for pre-

training and 8 for fine-tuning. We applied data augmentation methods in both steps, including random flipping, rotation, scaling, and ground truth box sampling [19].

C. Detection Performance on nuScenes test set

We evaluated the performance of our proposed D-Align in comparison with other LiDAR-based 3D object detectors. We trained our model on the entire nuScenes training set and evaluated its 3D object detection performance on the provided testing set. For convenience, we refer to D-Align with a PointPillars [11] baseline as *D-Align-PP* and D-Align with a CenterPoint [23] baseline as *D-Align-CP*. In Table I, we report the test detection accuracy for both D-Align-PP and D-Align-CP. We did not use any ensemble strategies. Other 3D object detectors include both single-frame methods and multi-frame methods. D-Align-CP achieved state-of-the-art (SOTA) performance, outperforming the latest 3D object detectors in terms of both mAP and NDS. D-Align exhibited significant improvements in performance over the single-frame based baselines. D-Align-PP achieved 22.5% mAP gain over PointPillar baseline and D-Align-PP achieved 3.7% mAP gain over CenterPoint baseline, which shows that the proposed method effectively utilized the spatio-temporal information in the point cloud sequence. D-Align performed particularly well in motorcycle, bicycle, and traffic cone object classes, which tend to be challenging to detect due to the lack of LiDAR points. D-Align compensated for the performance degradation for these classes caused by a lack of points by aggregating information over multi-frame features.

D. Ablation Studies on nuScenes validation set

We conducted ablation studies on the nuScenes validation set to confirm the effectiveness of each component of D-Align. PointPillars [11] was used as a baseline detector and

TABLE II

ABLATION STUDIES CONDUCTED ON nuSCENES VALIDATION SET

		Methods			
		(a)	(b)	(c)	(d)
IGANet			✓	✓	✓
IDANet	w/o FPMNet			✓	
	w/ FPMNet				✓
mAP [%]		44.40	46.17	48.19	49.66
NDS [%]		58.15	59.42	60.60	61.33

TABLE III

PERFORMANCE VS. THE NUMBER OF LAYERS

# of attention layers	D-Align		
	1	2	3
mAP [%]	46.44	48.58	49.66
NDS [%]	59.55	60.72	61.33

was trained with the whole nuScenes training set. We fine-tuned the entire D-Align network only with 1/7 training set to reduce the training effort.

Contribution of D-Align components. Table II provides the contributions of IGANet, IDANet, and FPMNet to overall detection performance. Method (a) is the single-frame-based PointPillars baseline. Method (b) adds only IGANet to the baseline detector. In this setting, the multi-frame features are directly aggregated by IGANet without feature alignment. This improves the performance of Method (a) by 1.77% in mAP and 1.27% in NDS. Note that the performance improvements that can be achieved without feature alignment are limited due to non-negligible discrepancy between multi-frame features. Method (c) adds inter-frame feature alignment by IDANet to Method (b) without using the temporal context features by FPMNet. Without FPMNet, T-QS is directly used to determine the mask offsets and weights. By aligning the multi-frame features through IDANet, Method (c) improves performance by 2.02% in mAP and 1.18% in NDS compared to Method (b). Finally, Method (d) is the proposed D-Align architecture with FPMNet added. Method (d) yields additional improvements of 1.47% in mAP and 0.73% in NDS over Method (c). This shows that the temporal context produced by FPMNet helps generate better attention to multi-frame features.

Performance versus the number of attention layers. Table III presents the performance of D-Align as a function of the number of attention layers, L . Due to the memory limitation, we increased L only up to three. The detection accuracy of D-Align improves with L . This confirms that the dual query attention progressively improves the BEV representation through successive layers of attention.

Performance versus length of input frames. Fig. 4 provides performance change with respect to the length of




Fig. 4. Performance vs. input frame length

the input frames. The accuracy of D-Align improves up to an input frame length of 3, but this trend stops at a length of 4. Considering that 4 frames in the nuScenes dataset are equivalent to 2 seconds, it seems difficult for D-Align to cope with large dynamic changes in features that can occur in duration over 2 seconds.

V. CONCLUSIONS

In this paper, we have proposed a novel dual-query based co-attention method called D-Align for 3D object detection, which can leverage spatio-temporal information in point cloud sequences. The proposed dual-query framework uses two query sets, T-QS and S-QS, which serve to refine the features of both the target and support frames in an iterative fashion. We designed IDANet to align S-QS to T-QS based on the temporal context features generated by FPMNet. We also introduced IGANet, which can aggregate dual query sets to refine T-QS. IDANet and IGANet transform the input query sets and produce the updated queries in each attention layer. Finally, the proposed D-Align produces the enhanced target frame BEV feature maps for 3D object detection. Our experimental results have confirmed that the proposed D-Align achieved substantial performance improvements compared to the baseline 3D object detectors and achieved state-of-the-art performance among the latest LiDAR-based 3D object detectors on the nuScenes dataset.

VI. ACKNOWLEDGEMENT

This work was partly supported by 1) the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No.2020R1A2C2012146), 2) the Institute of Information communications Technology Planning Evaluation (IITP) grant funded by the Korea government (MSIT) (No.2020-0-01373, Artificial Intelligence Graduate School Program (Hanyang University)) and 3) the Institute for Information communications Technology Promotion (IITP) grant funded by the Korea government (MSIP) (No.2021-0-01314, Development of driving environment data stitching technology to provide data on shaded areas for autonomous vehicles).

REFERENCES

- [1] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- [2] Nicolas Ballas, Li Yao, Chris Pal, and Aaron Courville. Delving deeper into convolutional networks for learning video representations. *arXiv preprint arXiv:1511.06432*, 2015.

- [3] Prarthana Bhattacharyya, Chengjie Huang, and Krzysztof Czarnecki. Sa-det3d: Self-attention based context-aware 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3022–3031, 2021.
- [4] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020.
- [5] 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. 33:21224–21235, 2020.
- [6] Qi Chen, Lin Sun, Zhixin Wang, Kui Jia, and Alan Yuille. Object as hotspots: An anchor-free 3d object detection approach via firing of hotspots. In *European conference on computer vision*, pages 68–84. Springer, 2020.
- [7] Runzhou Ge, Zhuangzhuang Ding, Yihan Hu, Yu Wang, Sijia Chen, Li Huang, and Yuan Li. Afdet: Anchor free one stage 3d object detection. *arXiv preprint arXiv:2006.12671*, 2020.
- [8] 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.
- [9] Rui Huang, Wanyue Zhang, Abhijit Kundu, Caroline Pantofaru, David A Ross, Thomas Funkhouser, and Alireza Fathi. An lstm approach to temporal 3d object detection in lidar point clouds. In *European Conference on Computer Vision*, pages 266–282. Springer, 2020.
- [10] Jaekyum Kim, Junho Koh, Yecheol Kim, Jaehyung Choi, Youngbae Hwang, and Jun Won Choi. Robust deep multi-modal learning based on gated information fusion network. In *Asian Conference on Computer Vision*, pages 90–106. Springer, 2018.
- [11] 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.
- [12] Yanwei Li, Yilun Chen, Xiaojuan Qi, Zeming Li, Jian Sun, and Jiaya Jia. Unifying voxel-based representation with transformer for 3d object detection. *arXiv preprint arXiv:2206.00630*, 2022.
- [13] 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.
- [14] Meytal Rapoport-Lavie and Dan Raviv. It’s all around you: Range-guided cylindrical network for 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2992–3001, 2021.
- [15] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.
- [16] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. 30, 2017.
- [17] SHI Xingjian, Zhouong Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting. In *Advances in neural information processing systems*, pages 802–810, 2015.
- [18] Ziyu Xiong, Huimin Ma, Yilin Wang, Tianyu Hu, and Qingmin Liao. Lidar-based 3d video object detection with foreground context modeling and spatiotemporal graph reasoning. In *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, pages 2994–3001. IEEE, 2021.
- [19] Yan Yan, Yuxing Mao, and Bo Li. Second: Sparsely embedded convolutional detection. *Sensors*, 18(10):3337, 2018.
- [20] 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.
- [21] Zetong Yang, Yin Zhou, Zhifeng Chen, and Jiquan Ngiam. 3d-man: 3d multi-frame attention network for object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1863–1872, 2021.
- [22] Junbo Yin, Jianbing Shen, Chenye Guan, Dingfu Zhou, and Ruigang Yang. Lidar-based online 3d video object detection with graph-based message passing and spatiotemporal transformer attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11495–11504, 2020.
- [23] 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.
- [24] Zhenxun Yuan, Xiao Song, Lei Bai, Zhe Wang, and Wanli Ouyang. Temporal-channel transformer for 3d lidar-based video object detection for autonomous driving. *IEEE Transactions on Circuits and Systems for Video Technology*, 2021.
- [25] 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.
- [26] 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.
- [27] Benjin Zhu, Zhengkai Jiang, Xiangxin Zhou, Zeming Li, and Gang Yu. Class-balanced grouping and sampling for point cloud 3d object detection. *arXiv preprint arXiv:1908.09492*, 2019.
- [28] Xinge Zhu, Yuexin Ma, Tai Wang, Yan Xu, Jianping Shi, and Dahua Lin. Ssn: Shape signature networks for multi-class object detection from point clouds. In *European Conference on Computer Vision*, pages 581–597. Springer, 2020.
- [29] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020.