

This WACV 2023 paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

TransPillars: Coarse-to-Fine Aggregation for Multi-Frame 3D Object Detection

Zhipeng Luo^{1,3} Gongjie Zhang¹ Changqing Zhou^{1,3} Tianrui Liu^{1,3} Shijian Lu^{1*} Liang Pan²

¹ Nanyang Technological University ² S-Lab, Nanyang Technological University ³ Sensetime Research

Abstract

3D object detection using point clouds has attracted increasing attention due to its wide applications in autonomous driving and robotics. However, most existing studies focus on single point cloud frames without harnessing the temporal information in point cloud sequences. In this paper, we design TransPillars, a novel transformerbased feature aggregation technique that exploits temporal features of consecutive point cloud frames for multiframe 3D object detection. TransPillars aggregates spatialtemporal point cloud features from two perspectives. First, it fuses voxel-level features directly from multi-frame feature maps instead of pooled instance features to preserve instance details with contextual information that are essential to accurate object localization. Second, it introduces a hierarchical coarse-to-fine strategy to fuse multi-scale features progressively to effectively capture the motion of moving objects and guide the aggregation of fine features. Besides, a variant of deformable transformer is introduced to improve the effectiveness of cross-frame feature matching. Extensive experiments show that our proposed TransPillars achieves state-of-art performance as compared to existing multi-frame detection approaches.

1. Introduction

3D object detection using point clouds has been actively studied in recent years due to its wide applications in the fields of autonomous driving and robotics. With the advent of large-scale datasets [8, 30, 1], popular deep learning-based single-frame 3D object detectors [47, 27, 26, 40, 20, 37, 36] are proposed. However, there are challenges difficult to resolve with a single sweep of point cloud. *First*, point clouds are sparse, especially at a long distance away from the LiDAR sensor. *Second*, incomplete point clouds due to partial observation, occlusion, and view truncation lead to ambiguities in the object geometry (as shown in Fig. 1 upper).



Figure 1: Upper: Illustration of challenges in point cloud object detection. Occluded objects are indicated by red boxes and LiDAR is highlighted by the axes. Lower: Sequential frames contain complementary information and aggregating the complementary information leads to a more complete view.

Leveraging multiple frames of point clouds, on the other hand, provides critical temporal cues that mitigate the above-mentioned challenges [38, 24, 41]. The lower section of Fig. 1 illustrates that points accumulated from consecutive frames progressively form a holistic depiction of the object. Nevertheless, leveraging temporal information is not trivial. A straightforward approach is simply concatenating points from multiple frames [1]. However, such a method does not explicitly model cross-frame relations and the performance deteriorates for moving objects as the number of concatenated frames exceeds a certain threshold [38]. Some recent approaches resort to fusing multi-frame information via feature aggregation [19, 13]. In particular, inspired by relational networks [12], [38] proposes a 3D multi-frame attention network that performs feature alignment and aggregation on pooled instance-level features. Despite its remarkable performance gain as compared to the single frame baseline, the RoI pooling process inevitably leads to 1) loss of instance details due to the misalignment between RoI and the object as dimension and location estimation is imper-

^{*}Corresponding author

fect, 2) loss of contextual information due to separation of the objects from the scene, both undermining cross-frame correlation modeling. The proposed feature pooling method also introduces extra class-specific heuristic designs, such as the number of key points for each class, which make the method less adaptable. Moreover, such a method depends on high-quality region proposals, which are not guaranteed under the above challenges associated with point clouds.

Motivated by the above observations, in this paper, we explore aggregating multi-frame information directly from feature maps. Specifically, we propose TransPillars, which builds on top of PointPillars [17] and employs the attention mechanism of transformer to perform cross-frame feature aggregation at voxel-level. To avoid prohibitive computational complexity and memory consumption, instead of performing global attention computations between each pair of tokens from the feature maps as in regular transformers, we adopt deformable attention [48] for feature aggregation that each voxel adaptively attends to a small number of target voxels. To address the limitations of the original deformable attention in the cross-frame matching process, we develop a variant of deformable attention by incorporating the query-key matching operation in the attention module to better adapt to moving objects. In order to effectively capture the motion of fast-moving objects, we introduce a novel coarse-to-fine aggregation strategy to first identify the highlevel instance correspondences to guide the subsequent fusion of fine features for accurate localization.

In summary, the contributions of this work are threefold. *First*, We propose a transformer-based multi-frame point cloud detection model named TransPillars. TransPillars performs voxel-level feature aggregation instead of using pooled instance features to preserve instance details and exploit contextual information for accurate localization. We show that the performance of our proposed method surpasses state-of-the-art multi-frame approaches on standard benchmarks. *Second*, we design a novel hierarchical coarse-to-fine feature aggregation strategy to effectively capture the motion of fast-moving objects to guide subsequent fusion of fine features. *Third*, we develop a variant of deformable attention for effective cross-frame feature matching in the feature aggregation process.

2. Related Work

Single-frame Point Cloud Object Detection. Single-frame 3D detectors can generally be classified into two categories, namely voxel-based methods and point-based methods. Voxel-based methods [47, 16, 34, 28, 40] first project points to grids of fixed size to form voxel representations and process the input with convolutions. In particular, 3D sparse convolution [10] is widely used to speed up the computation by leveraging the sparsity of point cloud and only considering the non-empty voxels. On the other hand,

point-based approaches [21, 27, 20, 37, 36] first utilize point cloud feature extractors [22, 23] to perform point sampling and feature extraction for subsequent region proposal generation or bounding box prediction. Recent method [26] propose to combine voxel-based and point-based approaches to obtain refined predictions. In this work, we adopt Point-Pillars [17], a lightweight single stage voxel-based detector commonly used by multi-frame detectors [39, 38, 42], as our base model.

Multi-frame Point Cloud Object Detection. A basic approach to utilize multiple frames for point cloud detection is through point concatenation of consecutive frames [1]. Despite its effectiveness, the performance gain over the single-frame method is limited especially over longer time intervals since it does explicitly model the relations among frames [38]. FaF [19] proposes to concatenate feature maps extracted from point clouds instead, but faces the similar limitation of misaligned representations. Some recent approaches [13, 39] employ recurrent networks to aggregate multi-frame features but such methods often lead to high computational cost. [24] explores an offboard setting for auto-labeling by performing detection for individual frames and aggregating the results from the entire sequence. Inspired by relational networks [12] and its applications [3, 29, 33, 9, 11] on 2D video object detection, 3D-MAN [38] proposes to apply attention mechanisms on pooled RoI features for multi-frame alignment and aggregation. However, RoI pooling separates the objects from the context and leads to the loss of details. Our method also resorts to attention mechanisms but performs feature aggregation at voxel-level to maximally preserve the instance details and contextual information.

Vision Transformers. Transformers [31] were first proposed in natural language processing (NLP) as an attentionbased building block, which allows for information aggregation from the entire input sequence. In recent years, transformers have been adopted to tackle various computer vision problems and achieved remarkable success [6, 2, 6, 32, 45, 15, 43, 44, 46]. One main advantage of transformers is their global receptive field and capability in capturing long-range relationships. However, such characteristics also bring unfavorable high computational cost and memory usage, which makes it challenging for applications that involve large-scale inputs or feature representations. Extensive studies have been carried out to address this issue and a number of researches [18, 5, 14, 4, 7, 35] propose to reduce the complexity with restricted attention patterns such as local windows. Such approaches may not be suitable for our task as objects are moving at various speeds, which makes it difficult to determine a suitable window size. On the other hand, [48] proposes deformable attention, which learns query-dependent sparse sampling locations to adaptively gather features from the value input.



Figure 2: The framework of the proposed TransPillars: Given multiple consecutive point cloud frames as inputs, multi-scale features are first extracted by a feature extractor which are then aggregated with multiple Fusion Aggregation Modules in a coarse-to-fine manner. We first aggregate coarse features to extract high-level cross-frame correspondences and then employ them to guide the aggregation of fine features. Finally, the multi-scale aggregated features are fused to produce the final predictions. Best viewed in color.

However, deformable attention does not explicitly enforce query-key matching, which leads to limited capability in the cross-frame feature matching in our task. In this work, we propose a variant of deformable attention to mitigate this issue.

3. Method

In this section, we introduce our proposed TransPillars for multi-frame point cloud object detection in detail. Section 3.1 briefly introduces the PointPillars [17] model, which is used as the base model in our proposed method. Section 3.2 gives an overview of our proposed method, followed by Sections 3.3 and 3.4, which explain the Fusion Aggregation Module and our proposed attention mechanism. Lastly, we describe the losses for model optimization in Section 3.5.

3.1. Preliminaries

We use PointPillars as the base model in our proposed method. PointPillars differs from regular voxel-based detectors that it only discretizes the input point cloud with a grid of fixed size in the x-y plane, thus forming pillars instead of cubic voxels. The points inside each pillar are then augmented and used to generate a feature vector. Note that only non-empty pillars are processed by the network to speed up the feature extraction. The obtained pillar features are then scattered back to their corresponding locations in the scene to form a pseudo-image representation. Subsequently, the pseudo-image features are processed by a feature pyramid network, which uses convolution layers to extract multi-scale features \mathbf{F}^i , where $i \in \{1, 2, 3\}$ denotes the scale level that i = 1 refers to the feature map of the smallest scale. Finally, the downsampled features are upsampled with transpose convolutions, and the final prediction is generated using a detection head based on the concatenated features. We refer the readers to [17] for details.

PointPillars is known as a lightweight 3D detector that achieves a good balance between efficiency and accuracy, and it has been widely adopted by multi-frame 3D detection approaches [38, 39, 42] as the baseline model. We also base our proposed method on PointPillars without making modifications to the baseline model and focus on the multiframe feature aggregation.

3.2. TransPillars

Fig.2 gives an overview of our proposed TransPillars. In the feature extraction stage, given a sequence of point cloud input $\{I_{T-n}\}_{n=0}^{N-1}$, where N denotes the number of frames in the sequence and I_T refers the current (latest) frame, a shared feature extractor of the base model is used to extract multi-scale features $\{\mathbf{F}_{T-n}\}_{n=0}^{N-1}$. In the subsequent feature aggregation stage, the goal is to aggregate useful information from the past frames to enrich the feature representations of the current frame in order to enhance the detection prediction. Therefore, we regroup the extracted features into current and past representations and perform attention-based feature aggregation.

As introduced in the related works, a commonly adopted paradigm in the 2D video detection field is to extract instance-level features with RoI Pooling or RoIAlign operations and perform feature fusion with relation networks [12] or its variants. However, unlike image frames which are essentially projections from the 3D real-world space to the flattened 2D space, point clouds possess different characteristics. First, due to the loss of depth information in projection, images do not contain explicit spatial information such as object size and location, while point clouds offer accurate 3D coordinates. Moreover, as compared to typical video detection datasets such as ImageNet VID [25], object movements pose a bigger challenge in point cloud applications such as autonomous driving where rapid moving objects are often observed. Region pooling methods separate objects from the scene and lead to loss of instance details as well as contextual information, which restricts the effectiveness of feature aggregation. Motivated by the above analysis, we propose to directly perform voxel-level aggregation from feature maps to maximally exploit the context and preserve details for accurate localization.

To perform voxel-level aggregation, two challenges are faced. The first challenge is the high computational cost and memory consumption associated with the attention mechanism due to the global matching process, which can be prohibitive for large-scale point cloud scenes studied in this work. To alleviate this issue, we resort to the recently proposed deformable attention [48], where each element in the query only adaptively attends to a small number of value tokens. However, the original deformable attention does not explicitly enforce query-key matching, which makes it less effective for our multi-frame detection task where crossframe feature matching is required. Therefore, we introduce a variant of deformable attention that incorporates the said matching process, which will be detailed in Section 3.4. In practice, in order to further reduce the computation, we select a proportion of the voxels from the current frame with high classification scores predicted by the base model to form the query feature.

The second challenge is to capture the motion of fastmoving objects to establish the cross-frame correspondence. To achieve this goal directly on fine features is difficult due to the large searching space brought by the small voxel size. To address this issue, we design a novel coarse-to-fine aggregation strategy by leveraging the multiscale features generated by the base model. Specifically, we employ a transformer-based Fusion Aggregation Module (FAM) to perform feature fusion and feature aggregation starting from the coarse feature maps. The output aggregated features are then fused with the feature maps of the next scale level for subsequent aggregations. The main idea is to utilize coarse feature maps to perform a coarse matching for cross-frame feature tokens, and the aggregated features are used to guide the following matching process of fine features whose tokens are of smaller physical size. Finally, the outputs of all FAMs are combined as in Point-Pillars [17] to generate the final prediction. The details of FAM are elaborated in Section 3.3.

We highlight that our proposed method emphasizes feature aggregation and does not require modifications to the base model. The same prediction head is used to generate



Figure 3: Architecture of the proposed Fusion Aggregation Module (FAM): FAM fuses (optional) and aggregates features of each scale level. For feature fusion, the previously aggregated feature is first upsampled and then fused with the current frame feature. For feature aggregation, the fused feature is aggregated with features of past frames using a transformer.

the final prediction based on the aggregated features.

3.3. Fusion Aggregation Module

Fig. 3 illustrates the architecture of the Fusion Aggregation Module (FAM). The FAM consists of feature fusion and feature aggregation operations. The feature fusion is an operation that takes as input the aggregated feature $\hat{\mathbf{F}}_T^{i-1}$ from the previous FAM as well as the current frame feature \mathbf{F}_T^i , where *i* denotes the scale level. The previous aggregated feature $\hat{\mathbf{F}}_T^{i-1}$ is upsampled by a transpose convolution layer to match the size of the current scale level and concatenated with the current frame feature \mathbf{F}_T^i . A convolution layer is used to fuse the two concatenated inputs and the fused feature $\bar{\mathbf{F}}_T^i$ is obtained by:

$$\bar{\mathbf{F}}_T^i = \operatorname{Conv}([\operatorname{upsample}(\hat{\mathbf{F}}_T^{i-1}), \mathbf{F}_T^i])$$
(1)

where $[\cdot]$ denotes concatenation. By fusing the previous aggregated feature of smaller scale with the current frame feature, the fused feature integrates the coarse cross-frame matching information, which is used to guide the subsequent feature aggregation. Note that feature fusion is not required for the smallest scale level since it is the first level.

In the feature aggregation stage, the fused feature $\mathbf{\bar{F}}_T^i$ aggregates information from the features from past frames denoted by $\{\mathbf{F}_{T-n}^i\}_{n=1}^{N-1}$, where N is the number of frames. As illustrated in Fig. 3, the fused feature is first processed by a multi-head self-attention module to gather features within the frame, and the output feature serves as the query for the subsequent cross-attention module which performs crossframe feature aggregation, while the features of the past frames are concatenated and serve as the key and value. The feature aggregation process can be summarized as:

$$\hat{\mathbf{F}}_{T}^{i} = \operatorname{Attn}(\operatorname{Attn}(\bar{\mathbf{F}}_{T}^{i}, \bar{\mathbf{F}}_{T}^{i}), [\mathbf{F}_{T-1}^{i}, ..., \mathbf{F}_{T-N+1}^{i}]) \quad (2)$$

where the output $\hat{\mathbf{F}}_T^i$ is the aggregated feature of the current scale *i*. Note that the normalization layer, skip connection and feed-forward network in the transformer are omitted for simplicity, and each FAM consists of *L* Transformer layers described above for feature aggregation.

Positional Objectiveness Encoding. On top of the regular positional encoding generated with sinusoidal functions as introduced in [31], we incorporate an additional objectiveness encoding to facilitate the feature aggregation process. Specifically, we obtain the classification prediction from the base model for each frame and use a convolution layer to encode it into the same dimension as the feature maps. In the case of multi-class prediction, the highest score is selected. Formally, the objectiveness encoding E_{obj} is obtained by:

$$E_{obj} = \operatorname{Conv}(\sigma(\max_{c=1}^{C} S))$$
(3)

where *C* denotes the number of classes and *S* represents the single-frame classification prediction. $\sigma(\cdot)$ is the sigmoid function that converts the scores to the [0,1] range. The objectiveness encoding is then summed with the positional encoding to form the positional objectiveness encoding. The classification prediction reflects the objectiveness of each token in the feature map, and explicitly adding this information to the query and key features helps the transformer locate foreground objects during the matching process.

3.4. Attention Mechanism

One known issue of regular transformers is their high computational complexity and memory consumption since global matching is performed between query and key tokens. The complexity grows quadratically as the number of tokens increases. This issue is even amplified as we deal with large-scale point cloud scenes. Deformable attention [48] is proposed in Deformable DETR as an alternative attention mechanism to reduce the computation and speed up the convergence of the detection model. As shown in Fig.4(a), instead of computing the global attention matrix, deformable attention generates a small number of sampling locations based on the position of the query element through a linear projection from the query feature. At the same time, the query feature is also projected to a set of attention weights corresponding to the sampling locations. Value samples are obtained from the value feature based on the sampling locations, and the output is computed by multiplying the attention weights with the value samples.



Figure 4: Comparison of the original deformable attention [48] and our proposed variant. (a) The original deformable attention generates the attention weights directly via projection from the query features. (b) Our proposed attention samples key elements and performs query-key matching to obtain the attention matrix, which is applied on the value samples to yield the output.

Deformable attention greatly reduces the computation complexity from quad-ratic to linear w.r.t. the feature length. Moreover, it does not restrict the attention pattern to a fixed local range, which makes it favorable for our application where objects are of different moving speeds. However, the original deformable attention does not explicitly enforce query-key matching for attention matrix generation, which limits the capability in modeling the crossframe correlation of moving objects in our task. Specifically, when the current frame feature serves as the query in cross-frame matching, the value feature comes from the past frames and might be misaligned with the query due to object movements. However, the attention weights are directly generated by projecting the query feature, which does not contain the motion information. Therefore, it is difficult for the attention module to generate meaningful attention weights in order to focus on the moving objects (as shown in Fig.6). To address this issue, we introduce a variant of deformable attention by incorporating the query-key matching process. Concretely, the projected sampling locations are used to generate key samples and value samples, respectively. The key samples are multiplied with the query feature to obtain the attention matrix, which is further applied to the value samples to generate the final output. Formally, we use $\{q_i\}$ to denote a set of query tokens as well as $\{\mathbf{k}_i\}$ and $\{\mathbf{v}_i\}$ for sampled key and value tokens. The attention matrix is computed by:

$$A_{i,j}^{h} = \operatorname{softmax}\left(\frac{(W_{q}\mathbf{q}_{i})^{T}(W_{k}\mathbf{k}_{j})}{\sqrt{d}}\right)$$
(4)

where h indexes the attention heads, and W_q and W_k are the query and key projections. \sqrt{d} is a scaling factor [31] while d is the feature dimension. The output of the attention is

calculated by:

$$\operatorname{Attn}(\mathbf{q}_{i}, \mathbf{k}_{j}, \mathbf{v}_{j}) = \sum_{h=1}^{H} W_{o}(\sum_{j=1}^{K} A_{i,j}^{h} \cdot W_{v} \mathbf{v}_{j}) \qquad (5)$$

where W_o and W_v are the output projection and value projection, respectively. K denotes the number of sample locations for each query element and H is the number of heads.

3.5. Losses

In the feature aggregation stage of FAMs, apart from the final aggregated feature, the intermediate outputs from each transformer layer is also kept to provide additional supervision. We fuse the multi-scale features in the same manner as the base model [17] to generate the final prediction for loss computation. The losses are the same as the base model, which consist of the classification loss \mathcal{L}_{cls} , the localization loss \mathcal{L}_{loc} , and the direction loss \mathcal{L}_{dir} . We sum the losses for all the transformer layers to get the aggregation loss:

$$\mathcal{L}_{aggr} = \frac{1}{L} \sum_{l=1}^{L} (\beta_{cls} \mathcal{L}_{cls} + \beta_{loc} \mathcal{L}_{loc} + \beta_{dir} \mathcal{L}_{dir}) \quad (6)$$

where *L* is the number of transformer layers and β_{cls} , β_{loc} , β_{dir} are the loss coefficients. On the other hand, the base model loss is computed for all the input frames, and the aggregation loss and the base model loss are summed to get the total loss:

$$\mathcal{L} = \mathcal{L}_{base} + \mathcal{L}_{aggr} \tag{7}$$

The whole model is optimized in an end-to-end manner.

4. Experiments

4.1. Datasets

We evaluate our proposed method on two popular point cloud detection benchmarks that include sequential frames. nuScenes. The nuScenes [1] dataset contains 700 sequences for training and 150 for validation. Each point cloud sequence is around 20s in length with a frame interval of 0.05s. Annotations are provided for every 10 consecutive frames, which are known as key frames. The main evaluation metric for the detection task is mean average precision (mAP). The mAP calculation uses a series of center distance thresholds instead of the commonly used box IoU threshold. Waymo. The Waymo Open Dataset [30] is a large-scale autonomous driving dataset that contains 798 point cloud sequences for training and 202 sequences for validation. The point clouds are collected using 64-line LiDAR, which gives around 180k points for each frame. The frame interval between consecutive frames is 0.1s and the annotations are provided for each frame in the training data. The detection

performance is evaluated by mean average precision (mAP) and mAP weighted by heading accuracy (mAPH) while the objects are categories into 2 difficulty levels based on the number of points contained.

4.2. Implementation Details

We use an input range of [-51.2m, 51.2m] for nuScenes. the x-y plane, and [-5m, 3m] for the z-axis. The voxel size is set to 0.2m, which leads to 512×512 pillars in total. The initial feature maps are downsampled with factors of [2, 4, 8] to obtain the multi-scale feature maps as described in [17]. Ego-motion correction is performed to compensate the self-motion. We concatenate the points from the associated sweeps with the key frame to form one input frame. We follow the practice in [39, 42] to use 3 key frames and their associated sweeps as input. For the FAM, we use 6 transformer layers with 8 attention heads. The number of deformable points is set as 8. We select voxels with top 5% classification scores to form the query feature. The base model is trained for 40 epochs with a learning rate of 0.001 and cosine annealing scheduling, and the whole model is then trained for 30 epochs with a learning rate 0.0005.

Waymo. For the experiments on Waymo, we use input point cloud range of [-76.8m, 76.8m] for x- and y-axis and [-2m, 4m] for the z-axis, following the settings in [38]. The voxel size is set as 0.3m, which gives us the same number of pillars as the nuScenes dataset. We follow [38] to use 16 frames as the input by dividing them into 4 windows of 4 frames, where the points inside each window are concatenated. The base model is first trained for 20 epochs with a learning rate of 0.003, and we then train the whole model for another 10 epochs with a learning rate of 0.0016. The rest of the settings are the same as nuScenes.

4.3. Benchmarking Results

We compare with existing multi-frame detection methods that exploit temporal information to improve the detection performance. We do not include single-frame detectors as they build on distinct architectures and parameters, which are not the focus of this study. Table 1 reports the detection performance on the nuScenes dataset. Our proposed method outperforms the current start-of-the-art multi-frame detection method TCTR [42] by a margin of 1.8 in mAP while achieving the best result for the majority of the categories. Table 2 reports the detection performance on the Waymo dataset. TransPillars outperforms the state-of-theart multi-frame method 3D-MAN [38]: although 3D-MAN employs a stronger single-frame base model than ours, we still achieve better results with convincingly larger margins of improvements upon the base models. The difference in the base model performance comes from the modifications made to the backbone by 3D-MAN, while we use the original PointPillars [17]. From the distance breakdown, it is

Table 1: Performance comparison with multi-frame 3D detectors on the nuScenes dataset. T.C., Moto. and Cons. stand for traffic cone, motorcycle, and construction vehicle, respectively. Mean Average Precision (mAP) is used for evaluation.

Method	Car	Ped	Bus	Barrier	T.C.	Truck	Trailer	Moto.	Cons.	Bicycle	Mean
3DVID [39]	79.7	76.5	47.1	48.8	58.8	33.6	43.0	40.7	18.1	7.9	45.4
TCTR [42]	83.2	74.9	63.7	53.8	52.5	51.5	33.0	54.0	15.6	22.6	50.5
Ours	84.0	77.9	62.0	55.1	55.4	52.4	34.3	55.2	18.9	27.6	52.3

Table 2: Performance comparison with multi-frame 3D detectors on the Waymo validation dataset. Even with a less competitive baseline model, our method outperforms the state-of-the-art multi-frame detector 3D-MAN.

Difficulty	Method	mAP (IoU=0.7)				mAPH (IoU=0.7)			
Difficulty	Wiethou	Overall	0-30m	30-50m	50m-Inf	Overall	0-30m	30-50m	50m-Inf
	ConvLSTM [13]	63.6	-	-	-	-	-	-	-
Level 1	3D-MAN (baseline)	69.03	87.99	66.55	43.15	68.52	87.57	65.92	42.37
	3D-MAN [38]	74.53	92.19	72.77	51.66	74.03	91.76	72.15	51.02
	Ours (baseline)	67.40	86.91	62.93	41.35	66.74	86.34	62.19	40.47
	Ours	74.97	91.39	73.10	53.52	74.42	90.91	72.48	52.75
	3D-MAN (baseline)	60.16	87.10	59.27	32.69	59.71	86.68	58.71	32.08
Level 2	3D-MAN [38]	67.61	92.00	67.20	41.38	67.14	91.57	66.62	40.84
	Ours (baseline)	58.94	86.07	57.13	31.72	58.36	85.50	56.44	31.02
	Ours	67.89	90.69	67.37	42.57	67.35	90.20	66.78	41.91



Figure 5: Compared with 3D-MAN, our proposed method has a more significant performance gain for objects that are further away (>30m), which are more susceptible to point sparsity and occlusions, thus benefiting more from multi-frame information.

observed from Fig. 5 that our model achieves higher performance gain for objects at a longer distance (>30m) compared to 3D-MAN, e.g. an improvement of 10.29 in mAPH for level 1 objects in the 30-50m range as compared to the 6.23 increment for 3D-MAN. As pointed out in [38], further objects typically have sparser point distributions and are more susceptible to occlusions, which make the multi-frame information more useful in complementing the single-frame view. The significant performance gain demonstrates the effectiveness of our method in aggregating multi-frame features.

4.4. Ablation Studies

We conduct ablation studies to investigate the effectiveness of our proposed components on the Waymo dataset. Due to the overwhelming number of training samples and limited computational resources, we sample 10% of se-

Table 3: Ablation study on different aggregation methods.

Aggragation Mathod	Le	vel 1	Level 2		
Aggregation Method	mAP	mAPH	mAP	mAPH	
Single-Scale	59.78	58.86	52.05	51.26	
Separate Multi-Scale	61.21	60.44	53.38	52.70	
Hierarchical Multi-Scale (Ours)	62.09	61.23	55.45	54.59	

quences from the training set for the experiments. We use 4 point cloud frames with an interval of 0.4s between adjacent frames as the input to report the results unless specified otherwise. The evaluation is conducted with the full validation dataset.

Aggregation Methods. We compare our proposed hierarchical coarse-to-fine feature aggregation method with two different approaches, namely *Single-scale* and *Separate Multi-scale*. For the Single-scale setting, the final concatenated feature maps of the base model are used for feature aggregation. Separate Multi-scale refers to performing feature aggregated features. As shown in Table 3, we observe that the performance drops for both methods while Single-scale suffers a larger loss. It shows that multi-scale features and our proposed hierarchical coarse-to-fine aggregation strategy makes the feature aggregation more effective.

Attention Mechanisms. We compare our proposed attention mechanism with the original deformable attention [48] and report the results in Table 4. The model using the original deformable attention [48] experiences a clear per-

Table 4: Ablation study on attention mechanisms.

Attention Type	Le	vel 1	Level 2		
Attention Type	mAP	mAPH	mAP	mAPH	
DeformAttention [48]	60.86	60.04	53.03	52.32	
Ours	62.09	61.23	55.45	54.59	

Table 5: Ablation study on transformer encodings.

Encoding Type	Le	vel 1	Level 2		
	mAP	mAPH	mAP	mAPH	
w/o Pos. Encoding	61.52	60.68	54.87	54.09	
w/o Obj. Encoding	61.29	60.47	53.32	52.60	
Ours	62.09	61.23	55.45	54.59	

formance drop as compared to our proposed attention mechanism, which shows the importance of query-key matching for aggregating multi-frame features in the task this paper studies.

Transformer Encoding. To investigate the effectiveness of our proposed positional objectiveness encoding, we conduct experiments by removing the positional encoding and the objectiveness encoding. The results in Table 5 show that the removal of each type of encoding leads to a loss in model performance, while the objectiveness encoding has a more significant impact. It demonstrates that the objectiveness encoding has a positive effect on guiding the feature aggregation by adding an additional signal to indicate the objectiveness of each input token.

4.5. Further Discussions

Number of Input Frames. We study the performance of our proposed method under different numbers of input frames. For the experiments with 8 and 16 input frames, we divide them into 4 windows of 2 and 4 concatenated frames, respectively. Additionally, we also compare with the basic point concatenation [1] method using 4 frames (denoted as 4^*). As shown in Table 6, by using 4 consecutive frames as input, there is already a substantial improvement over the single-frame baseline (e.g. an increase of 6.86 in mAP for level 1 objects.). Our method outperforms the point concatenation method by a clear margin when using the same number of frames, which demonstrates the benefit of explicit cross-frame correlation modeling. As the number of frames is further increased, we notice a higher improvement for level 2 objects, which is in line with the observation that objects with sparse point distributions can benefit more from multi-frame information.

Attention Visualization. We visualize the learned attention patterns on the bird eye's view of object instances. As illustrated in Fig. 6, we use red circles to represent the attention sampling locations while the radius of each circle indicates the attention magnitude. It is observed that the orig-

 Table 6: Performance comparison on number of input frames. * denotes point concatenation.

Number of	Lev	vel 1	Level 2		
Frames	mAP	mAPH	mAP	mAPH	
1	54.43	53.52	47.04	46.24	
4*	57.01	56.21	50.39	49.66	
4	60.75	59.91	52.76	52.02	
8	61.26	60.43	53.36	52.64	
16	62.16	61.41	55.55	54.85	



Figure 6: Visualization of attention weights. Current frame points are shown in dark blue, and the intensity decreases as the time interval increases for past frame points. The arrows indicate the objects' moving directions. The attention sampling locations are shown in red, while the radius of each circle indicates the attention weight magnitude.

inal deformable attention [48] demonstrates highly sparse attention patterns haphazardly scattered around the object, while our proposed attention mechanism effectively focuses on the trajectories of moving objects. It demonstrates that the proposed query-key matching operation is essential to the cross-frame feature aggregation task as it is more aware of the object motion in the aggregation process.

5. Conclusion

In this paper, we present a multi-frame 3D object detection method named TransPillars, which performs voxellevel feature aggregation. We design a novel hierarchical aggregation strategy to perform coarse-to-fine aggregation based on multi-scale features: coarse features are used to capture object motions to guide the fusion of fine features for accurate localization. A variant of deformable attention is introduced for efficient and motion-aware feature matching. Experimental results show that TransPillars outperforms state-of-the-art multi-frame detection methods on standard benchmarks.

References

- [1] 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.
- [2] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-toend object detection with transformers. In *European Conference on Computer Vision*, pages 213–229. Springer, 2020.
- [3] Yihong Chen, Yue Cao, Han Hu, and Liwei Wang. Memory enhanced global-local aggregation for video object detection. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 10337–10346, 2020.
- [4] Xiangxiang Chu, Zhi Tian, Yuqing Wang, Bo Zhang, Haibing Ren, Xiaolin Wei, Huaxia Xia, and Chunhua Shen. Twins: Revisiting the design of spatial attention in vision transformers. *arXiv preprint arXiv:2104.13840*, 1(2):3, 2021.
- [5] Xiaoyi Dong, Jianmin Bao, Dongdong Chen, Weiming Zhang, Nenghai Yu, Lu Yuan, Dong Chen, and Baining Guo. Cswin transformer: A general vision transformer backbone with cross-shaped windows. *arXiv preprint arXiv:2107.00652*, 2021.
- [6] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929, 2020.
- [7] Jiemin Fang, Lingxi Xie, Xinggang Wang, Xiaopeng Zhang, Wenyu Liu, and Qi Tian. Msg-transformer: Exchanging local spatial information by manipulating messenger tokens. arXiv preprint arXiv:2105.15168, 2021.
- [8] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013.
- [9] Qichuan Geng, Hong Zhang, Na Jiang, Xiaojuan Qi, Liangjun Zhang, and Zhong Zhou. Object-aware feature aggregation for video object detection. arXiv preprint arXiv:2010.12573, 2020.
- [10] Benjamin Graham, Martin Engelcke, and Laurens Van Der Maaten. 3d semantic segmentation with submanifold sparse convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9224–9232, 2018.
- [11] Mingfei Han, Yali Wang, Xiaojun Chang, and Yu Qiao. Mining inter-video proposal relations for video object detection. In *European Conference on Computer Vision*, pages 431– 446. Springer, 2020.
- [12] Han Hu, Jiayuan Gu, Zheng Zhang, Jifeng Dai, and Yichen Wei. Relation networks for object detection. In *Proceed*ings of the IEEE conference on computer vision and pattern recognition, pages 3588–3597, 2018.

- [13] Rui Huang, Wanyue Zhang, Abhijit Kundu, Caroline Pantofaru, David A Ross, Thomas Funkhouser, and Alireza Fathi. An 1stm approach to temporal 3d object detection in lidar point clouds. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVIII 16*, pages 266–282. Springer, 2020.
- [14] Zilong Huang, Youcheng Ben, Guozhong Luo, Pei Cheng, Gang Yu, and Bin Fu. Shuffle transformer: Rethinking spatial shuffle for vision transformer. arXiv preprint arXiv:2106.03650, 2021.
- [15] Drew A Hudson and C Lawrence Zitnick. Generative adversarial transformers. arXiv preprint arXiv:2103.01209, 2021.
- [16] Hongwu Kuang, Bei Wang, Jianping An, Ming Zhang, and Zehan Zhang. Voxel-fpn: Multi-scale voxel feature aggregation for 3d object detection from lidar point clouds. *Sensors*, 20(3):704, 2020.
- [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.
- [18] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. arXiv preprint arXiv:2103.14030, 2021.
- [19] Wenjie Luo, Bin Yang, and Raquel Urtasun. Fast and furious: Real time end-to-end 3d detection, tracking and motion forecasting with a single convolutional net. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 3569–3577, 2018.
- [20] Charles R Qi, Or Litany, Kaiming He, and Leonidas J Guibas. Deep hough voting for 3d object detection in point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9277–9286, 2019.
- [21] Charles R Qi, Wei Liu, Chenxia Wu, Hao Su, and Leonidas J Guibas. Frustum pointnets for 3d object detection from rgbd data. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 918–927, 2018.
- [22] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference* on computer vision and pattern recognition, pages 652–660, 2017.
- [23] Charles R Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. arXiv preprint arXiv:1706.02413, 2017.
- [24] Charles R Qi, Yin Zhou, Mahyar Najibi, Pei Sun, Khoa Vo, Boyang Deng, and Dragomir Anguelov. Offboard 3d object detection from point cloud sequences. In *Proceedings of* the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 6134–6144, 2021.
- [25] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.

- [26] Shaoshuai Shi, Chaoxu Guo, Li Jiang, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. Pv-rcnn: Pointvoxel feature set abstraction for 3d object detection. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 10529–10538, 2020.
- [27] Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. Pointrcnn: 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.
- [28] 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*, 2020.
- [29] Mykhailo Shvets, Wei Liu, and Alexander C Berg. Leveraging long-range temporal relationships between proposals for video object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9756– 9764, 2019.
- [30] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 2446–2454, 2020.
- [31] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In Advances in neural information processing systems, pages 5998–6008, 2017.
- [32] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. arXiv preprint arXiv:2102.12122, 2021.
- [33] Haiping Wu, Yuntao Chen, Naiyan Wang, and Zhaoxiang Zhang. Sequence level semantics aggregation for video object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9217–9225, 2019.
- [34] Yan Yan, Yuxing Mao, and Bo Li. Second: Sparsely embedded convolutional detection. *Sensors*, 18(10):3337, 2018.
- [35] Jianwei Yang, Chunyuan Li, Pengchuan Zhang, Xiyang Dai, Bin Xiao, Lu Yuan, and Jianfeng Gao. Focal self-attention for local-global interactions in vision transformers. arXiv preprint arXiv:2107.00641, 2021.
- [36] 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.
- [37] Zetong Yang, Yanan Sun, Shu Liu, Xiaoyong Shen, and Jiaya Jia. Std: Sparse-to-dense 3d object detector for point cloud. In *Proceedings of the IEEE/CVF International Conference* on Computer Vision, pages 1951–1960, 2019.
- [38] 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.

- [39] 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.
- [40] Tianwei Yin, Xingyi Zhou, and Philipp Krahenbuhl. Centerbased 3d object detection and tracking. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 11784–11793, 2021.
- [41] Cunjun Yu, Zhongang Cai, Daxuan Ren, and Haiyu Zhao. Leveraging temporal information for 3d detection and domain adaptation. arXiv preprint arXiv:2006.16796, 2020.
- [42] 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 Trans*actions on Circuits and Systems for Video Technology, 2021.
- [43] Gongjie Zhang, Zhipeng Luo, Kaiwen Cui, and Shijian Lu. Meta-detr: Few-shot object detection via unified image-level meta-learning. arXiv preprint arXiv:2103.11731, 2(6), 2021.
- [44] Gongjie Zhang, Zhipeng Luo, Yingchen Yu, Kaiwen Cui, and Shijian Lu. Accelerating detr convergence via semanticaligned matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 949–958, 2022.
- [45] Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip HS Torr, et al. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6881– 6890, 2021.
- [46] Changqing Zhou, Zhipeng Luo, Yueru Luo, Tianrui Liu, Liang Pan, Zhongang Cai, Haiyu Zhao, and Shijian Lu. Pttr: Relational 3d point cloud object tracking with transformer. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 8531–8540, 2022.
- [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.
- [48] 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.