

3D Single-object Tracking in Point Clouds with High Temporal Variation

Qiao Wu¹, Kun Sun², Pei An³, Mathieu Salzmann⁴, Yanning Zhang¹, and Jiaqi Yang^{1*}

¹ Northwestern Polytechnical University

² China University of Geosciences, Wuhan

³ HuaZhong University of Science and Technology

⁴ École Polytechnique Fédérale de Lausanne

qiaowu@mail.nwpu.edu.cn, jqyang@nwpu.edu.cn

Abstract. The high temporal variation of the point clouds is the key challenge of 3D single-object tracking (3D SOT). Existing approaches rely on the assumption that the shape variation of the point clouds and the motion of the objects across neighboring frames are smooth, failing to cope with high temporal variation data. In this paper, we present a novel framework for 3D SOT in point clouds with high temporal variation, called HVTrack. HVTrack proposes three novel components to tackle the challenges in the high temporal variation scenario: 1) A Relative-Pose-Aware Memory module to handle temporal point cloud shape variations; 2) a Base-Expansion Feature Cross-Attention module to deal with similar object distractions in expanded search areas; 3) a Contextual Point Guided Self-Attention module for suppressing heavy background noise. We construct a dataset with high temporal variation (KITTI-HV) by setting different frame intervals for sampling in the KITTI dataset. On the KITTI-HV with 5 frame intervals, our HVTrack surpasses the state-of-the-art tracker CXTracker by **11.3%/15.7%** in Success/Precision.

Keywords: 3D single-object tracking · High temporal variation · Point cloud

1 Introduction

3D single-object tracking (3D SOT) is pivotal for autonomous driving [2, 40] and robotics [16, 19, 25, 43]. Given the target point cloud and 3D bounding box as template, the goal of 3D SOT is to regress the target 3D poses in the tracking point cloud sequence. Existing approaches [5, 7, 9–12, 24, 28, 34, 39, 44–46] rely on the assumption that the point cloud variations and motion of the object across neighboring frames are relatively smooth. They crop out a small search area around the last proposal for tracking, thus dramatically reducing the complexity of the problem. The template and search area features are then typically correlated as shown in Fig. 1a, and used to regress the 3D bounding box.

* Corresponding author.




Fig. 1: Feature correlation in 3D SOT. (a) Feature correlation in the smooth case (1 frame interval). Correlating the features is relatively trivial as the target undergoes only small shape variations, and the observation angles are consistent in the three frames. (b-c) Feature correlation in high temporal variation cases (10 frames interval). The pose relative to the camera changes rapidly. Correlating the features using historical information is highly challenging (b). We encode the historical observation angles α into the features to guide the variation of relative pose to the camera (c).

In practice, these approaches are challenged by the presence of large point cloud variations due to the limited sensor temporal resolution and the moving speed of objects as shown in Fig. 1b. We refer to this significant variation in point cloud and object position between two frames as the high temporal variation (HV). The high temporal variation challenge is non-negligible in existing benchmarks, and exists in other scenarios not yet covered by them, such as:

- Skipped-tracking, which can greatly reduce computational consumption in tracking and serve a wide range of other tasks such as detection [20] and segmentation [41].
- Tracking in edge devices, which is essential for deploying trackers on common devices with limited frame rate, resolution, computation, and power *etc.*
- Tracking in highly dynamic scenarios [15], which is common in life. For example, tracking in sports events, highway, and UAV scenarios.

There are three challenges for 3D SOT in HV point clouds, and existing approaches are not sufficient to address these challenges. 1) *Strong shape variations of the point clouds*: Point cloud shape variations are usually caused by the occlusion and relative pose transformation between the object and the sensor. As illustrated in Fig. 1b, feature correlation in existing approaches fails because of the dramatic change in the density and distribution of points. 2) *Distractions due to similar objects*: When objects suffer from a significant motion, the search area needs to be enlarged to incorporate the target, thus introducing more distractions from similar objects. Most of the existing trackers focus on local scale features, which discards environmental spatial contextual information to handle distractions. 3) *Heavy background noise*: The expansion of the search area further reduces the proportion of target information in the scene. While aiming to find the high template-response features in the feature correlation stage, existing methods then neglect to suppress the noise interference and reduce the impact of noise features. We evaluate state-of-the-art (SOTA) trackers [24, 39, 44, 45] in the high temporal variation scenario as shown in Fig. 2. Their performance drops dramatically as the temporal variation of scene point clouds enlarges.




Fig. 2: Comparison of HVTrack with the SOTAs [24, 39, 44, 45] on ‘Car’ from KITTI-HV (KITTI [8] with different frame intervals, see Sec. 4).

To address the above challenges, we propose a novel framework for 3D SOT in point clouds with **H**igh temporal **V**ariation, which we call HVTrack. Specifically, we propose three novel modules to address each of the three above-mentioned challenges. 1) A Relative-Pose-Aware Memory (RPM) module to handle the strong shape variations of the point clouds. Different from [17], we integrate the foreground masks and observation angles into the memory bank. Therefore, the model can implicitly learn the distribution variation of point clouds from the relative pose in time. The information arising from observation angles has been overlooked by all existing trackers. 2) A Base-Expansion Feature Cross-Attention (BEA) module to deal with the problem of similar object distractions occurring in large scenes. We synchronize the correlation of the hybrid scales features (base and expansion scales, Sec. 3.4) in the cross-attention, and efficiently utilize spatial contextual information. 3) A Contextual Point Guided Self-Attention (CPA) module to suppress the background noise introduced by the expanded search area. It aggregates the features of points into contextual points according to their importance. Less important points share fewer contextual points and vice versa, thus suppressing most of the background noise. BEA and CPA are inspired by the SGFormer [26], which utilizes hybrid scale significance maps to assign more tokens to salient regions of 2D images. Our experiments clearly demonstrate the remarkable performance of HVTrack in high temporal variation scenarios, as illustrated in Fig. 2. Our contributions can be summarized as follows:

- For the first time, to the best of our knowledge, we explore the new 3D SOT task for high temporal variation scenarios, and propose a novel framework called HVTrack for the task.
- We propose three novel modules, RPM, BEA, and CPA, to address three challenges for 3D SOT in HV point clouds: strong point cloud variations, similar object distractions, and heavy background noise.
- HVTrack yields state-of-the-art results on KITTI-HV and Waymo, and ranks second on KITTI. Our experimental results demonstrate the robustness of HVTrack in both smooth and high temporal variation cases.

2 Related Work

2.1 3D Single-object Tracking

Most of the 3D SOT approaches are based on a Siamese framework, because the appearance variations of the target between neighboring frames are not significant. The work of Giancola *et al.* [9] constitutes the pioneering method in 3D SOT. However, it only solved the discriminative feature learning problem, and used a time-consuming and inaccurate heuristic matching to locate the target. Zarzar *et al.* [42] utilized a 2D RPN in bird’s eyes view to build an end-to-end tracker. The P2B network [24] employs VoteNet [22] as RPN and constructs the first point-based tracker. The following works [7, 11, 12, 28, 34, 44] develop different architectures of trackers based on P2B [24]. V2B [11] leverages the target completion model to generate the dense and complete targets and proposes a simple yet effective voxel-to-BEV target localization network. BAT [44] utilizes the relationship between points and the bounding box, integrating the box information into the point clouds. With the development of transformer networks, a number of works [5, 10, 12, 28, 39, 46] have proposed to exploit various attention mechanisms. STNet [12] forms an iterative coarse-to-fine cross-and self-attention to correlate the target and search area. CXTrack [39] employs a target-centric transformer to integrate targetness information and contextual information. TAT [17] leverages the temporal information to integrate target cues by applying an RNN-based [4] correlation module. Zheng *et al.* [45] presented a motion-centric method M2-Track, which is appearance matching-free and has made great progress in dealing with the sparse point cloud tracking problem. Wu *et al.* [37] proposed the first semi-supervised framework in 3D SOT.

While effective in their context, the above methods are designed based on the assumption that the point cloud variation and motion of the objects across neighboring frames are not significant. In high temporal variation scenarios, this assumption will lead to performance degradation because of the point cloud variations and interference naturally occurring in large scenes. Here, we introduce HVTrack to tackle the challenges of 3D SOT in high temporal variation scenarios.

2.2 3D Multi-object Tracking

3D multi-object tracking (MOT) in point clouds follows two main streams: Tracking-by-detection, and learning-based methods. Tracking-by-detection [1, 3, 32, 35] usually exploits methods such as Kalman filtering to correlate the detection results and track the targets. CenterTrack [47], CenterPoint [40], and SimTrack [18] replace the filter by leveraging deep networks to predict the velocity and motion of the objects. The learning-based methods [6, 27, 36] typically apply a Graph Neural Network to tackle the association challenge in MOT. GNN3DMOT [36] leverages both 2D images and 3D point clouds to obtain a robust association. 3DMOTFormer [6] constructs a graph transformer framework and achieves a good performance using only 3D point clouds.

3D MOT and 3D SOT have different purposes and their own challenges [13]. 3D MOT is object-level and focuses on correlating detected objects, whereas 3D SOT is intra-object-level [14] and aims to track a single object given a template. 3D SOT methods usually come with much lower computational consumption and higher throughput [46]. Also, 3D MOT is free from the challenges posed by the dynamic change in the search area size, as MOT is not required to adopt the search area cropping strategy in SOT.

3 Method

3.1 Problem Definition

Given the template of the target, the goal of 3D SOT is to continually locate the poses of the target in the search area point cloud sequence $\mathbf{P}^s = \{P_0^s, \dots, P_t^s, \dots, P_n^s | P_t^s \in \mathbb{R}^{N_s \times 3}\}$. Usually, the target point cloud with labels in the first frame is regarded as the template. Former trackers [5, 7, 9–12, 24, 28, 34, 39, 44–46] leverage a 3D bounding box label $B_0 = (x, y, z, w, l, h, \theta) \in \mathbb{R}^7$ to generate the template in the input. Here, (x, y, z) , (w, l, h) and θ are the center location, bounding box size (width, length, and height), and rotation angle of the target, respectively. As objects can be assumed to be rigid, the trackers only need to regress the center and rotation angle of the target.

3.2 Overview

We propose HVTrack to exploit both temporal and spatial information and achieve robust tracking in high temporal variation scenarios. As shown in Fig. 3, we take the point cloud P_t^s at time t as the search area, and leverage memory banks as the template. We first employ a backbone to extract the local spatial features $\mathcal{X}_0 \in \mathbb{R}^{N \times C}$ of P_t^s , with N and C the point number and feature channel, respectively. Then, L transformer layers are employed to extract spatio-temporal information. For each layer l , (i) we capture the template information $Mem_l \in \mathbb{R}^{K \times N \times C}$ from the Relative-Pose-Aware Memory module, with K the memory bank size (Sec. 3.3); (ii) the memory features and search area features \mathcal{X}_{l-1} are correlated in the Base-Expansion Features Cross-Attention (Sec. 3.4); (iii) the Contextual Point Guided Self-Attention (Sec. 3.5) leverages the attention map in the Base-Expansion Features Cross-Attention to suppress the noise features; (iv) we update the Layer Features memory bank using \mathcal{X}_{l-1} . After the transformer layers, an RPN is applied to regress the location $(x_t, y_t, z_t, \theta_t)$, the mask $\mathcal{M}_t \in \mathbb{R}^{N \times 1}$, and the observation angle $\alpha \in \mathbb{R}^2$. Finally, the mask and observation angle memory banks are updated using the predicted results.

3.3 Relative-Pose-Aware Memory Module

As shown in Fig. 1(b), rapid changes in relative pose lead to large variations in the shape of the object point cloud across the frames. Correlating the object

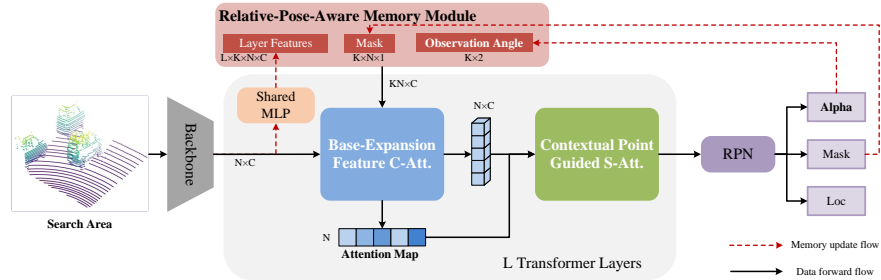


Fig. 3: HVTrack framework. We first utilize a backbone to extract the local embedding features of the search area. Then, we construct L transformer layers to fuse spatio-temporal information. For each transformer layer, (i) we apply three memory bank features in the Relative-Pose-Aware Memory module to generate temporal template information; (ii) we employ the Base-Expansion Feature Cross-Attention to correlate the template and search area by leveraging hybrid scale spatial context-aware features; (iii) we introduce a Contextual Point Guided Self-Attention to suppress unimportant noise. After each layer, we update the layer features memory bank using the layer input. Finally, we apply an RPN to regress the 3D bounding box, and update the mask and observation angle memory banks.

features in $(t-2, t-1, t)$ then becomes difficult, as they have a low overlap with each other. To address this, we introduce the observation angle into the memory bank. The observation angle gives us knowledge of the coarse distribution of an object’s point cloud. Thus, the model can learn the variations in point cloud distribution from the historical changes of observation angle.

To exploit the temporal information as the template, we propose a Relative-Pose-Aware Memory (RPM) module. RPM contains 3 memory banks. 1) A layer features memory bank (LM) $\in \mathbb{R}^{L \times K \times N \times C}$: We leverage the historical transformer layer features as the template features to reduce the template inference time in former trackers [5, 9–12, 24, 28, 34, 44, 46]. 2) A mask memory bank (MM) $\in \mathbb{R}^{K \times N \times 1}$: Inspired by the mask-based trackers [39, 45], we utilize the mask as the foreground representation. 3) An observation angle memory bank (OM) $\in \mathbb{R}^{K \times 2}$. For each transformer layer l , we process the memory features as

$$T_l = \text{Linear}([\text{LM}_l, \text{MM}, \text{Repeat}(\text{OM})]), \quad (1)$$

where $T_l \in \mathbb{R}^{KN \times C}$ denotes the template features, $\text{Linear}(\cdot)$ is a linear layer that projects the features from $\mathbb{R}^{KN \times (C+3)}$ to $\mathbb{R}^{KN \times C}$, $[\cdot]$ is the concatenation operation, and $\text{Repeat}(\cdot)$ stacks the OM to $\mathbb{R}^{K \times N \times 2}$. Then, we project T_l into Query (Q), Key (K), and Value (V) using the learnable parameter matrices as

$$\begin{aligned} Q_l^T &= \text{LN}(\text{LN}(T_l)W_l^{TQ} + \text{PE}^T), \\ K_l^T &= \text{LN}(T_l)W_l^{TK}, \\ V_l^T &= \text{LN}(T_l)W_l^{TV}, \end{aligned} \quad (2)$$

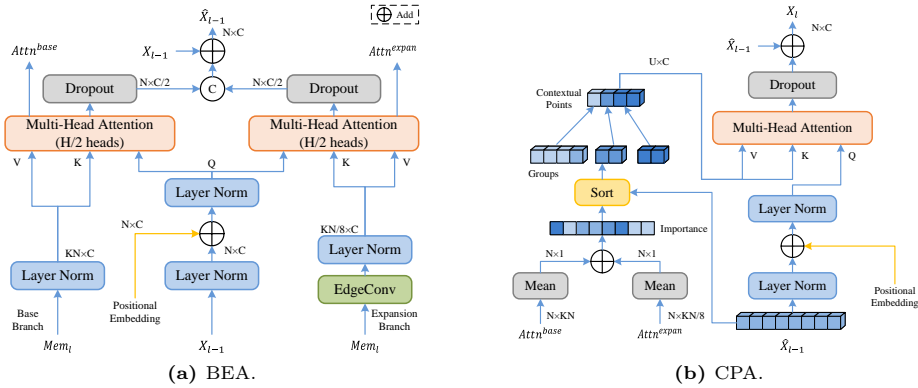


Fig. 4: (a) Base-Expansion Feature Cross-Attention (BEA). The H heads in the multi-head attention (MHA) are split to process hybrid scale features. For the base scale branch, we directly put the local features into the MHA. For the expansion scale branch, we apply an EdgeConv [33] to expand the receptive field of each point and extract more abstract features before MHA. BEA captures the spatial context-aware information with a humble extra computational cost. **(b) Contextual Point Guided Self-Attention (CPA).** We determine the importance of each point by both base and expansion scale attention maps. Then, we aggregate all the points into U clusters (contextual points) according to their importance and project the clusters to K and V . We assign fewer contextual points for low-importance points, and vice versa. CPA not only suppresses the noise but also reduces the computational cost of the attention.

where $\text{LN}(\cdot)$ is the layer norm, and $\text{PE}^T \in \mathbb{R}^{KN \times C}$ is the positional embedding of the historical point cloud coordinates. We utilize a linear layer to project the point cloud coordinates to their positional embedding. Finally, a self-attention is applied for internal interactions between temporal information as

$$Mem_t^* = T_t + \text{Dropout}(\text{MHA}(Q_t^T, K_t^T, V_t^T)), \quad (3)$$

where MHA is the multi-head attention in [31], and Dropout is the random dropping operation in [29]. Following CXTrack [39], we apply dropout and feed-forward network (FFN) after self-attention, i.e.,

$$Mem_t = Mem_t^* + \text{Dropout}(\text{FFN}(\text{LN}(Mem_t^*))), \quad (4)$$

$$\text{FFN}(x) = \max(0, xW_1 + b_1)W_2 + b_2. \quad (5)$$

3.4 Base-Expansion Feature Cross-Attention

Most of the existing trackers [11, 24, 28, 34, 39, 44, 46] employ a point based backbone [23, 33] and focus on local region features, which we call base scale features. Using only base scale features in the whole pipeline is quite efficient and effective

in small scenes. However, the base scale features are limited in representing the neighboring environment features around the object in large search areas. To tackle the challenge of similar object distractions, spatial context information across consecutive frames is crucial for effective object tracking [39]. Expanding the receptive field of features can help capture spatial contextual information, and such features are called expansion scale features. Inspired by [26], we propose Base-Expansion Feature Cross-Attention (BEA) to capture both local and more abstract features, and exploit spatial context-aware information.

As shown in Fig. 4a, the input features X_{l-1} are projected into Q. Usually, the memory features Mem_l would be projected into K and V. Then, multi-head cross-attention adopts H independent heads, and processes them using the same base scale features. By contrast, we split the H heads into 2 groups. $H/2$ heads exploit local spatial context information. We directly process the base scale features with normal cross-attention, and output base scale features $\hat{X}_{l-1}^{base} \in \mathbb{R}^{N \times C/2}$ and attention map $Attn^{base} \in \mathbb{R}^{N \times KN}$. The other $H/2$ heads capture environment context features. We first apply an EdgeConv [33] to extract more abstract features $Mem_l^{expan} \in \mathbb{R}^{KN/8 \times C}$, which are expansion scale features, i.e.,

$$Mem_l^{expan} = \text{EdgeConv}(Mem_l). \quad (6)$$

Then, we project the expansion features into K and V, and perform multi-head cross-attention with Q. Specifically, for the i -th head belonging to the expansion scale branch, we generate Q, K, and V as

$$\begin{aligned} Q_i &= LN(LN(X_{l-1})W_i^Q + PE_i^S), \\ K_i &= LN(Mem_l^{expan})W_i^K, \\ V_i &= LN(Mem_l^{expan})W_i^V, \end{aligned} \quad (7)$$

where PE_i^S is the positional embedding of search area point cloud coordinates. Then, cross-attention is performed as

$$Attn_i^{expan} = \text{Softmax}\left(\frac{Q_i K_i}{\sqrt{d_h}}\right), \quad (8)$$

$$h_i^{expan} = Attn_i^{expan} V_i, \quad (9)$$

where d_h is the feature dimension of the heads, and h_i^{expan} is the output features of the i -th head. After that, we concatenate the output features and attention map of each head as

$$\begin{aligned} \hat{X}_{l-1}^{expan} &= [h_1, \dots, h_{H/2}], \\ Attn^{expan} &= [Attn_1^{expan}, \dots, Attn_{H/2}^{expan}], \end{aligned} \quad (10)$$

where $\hat{X}_{l-1}^{expan} \in \mathbb{R}^{N \times C/2}$, and $Attn^{expan} \in \mathbb{R}^{N \times KN/8}$. Finally, we concatenate the base scale and expansion scale outputs as the resulting correlation feature $\hat{X}_{l-1} \in \mathbb{R}^{N \times C}$. Thus, BEA provides rich hybrid scale spatial contextual information for each point, with a very humble extra computational cost.

3.5 Contextual Point Guided Self-Attention

Most of the information in the search area will be regarded as noise, because we are only interested in one single object to be tracked. Existing trackers [11, 24, 28, 34, 44] aim to find the features with high template-response in the search area, but neglect the suppress to the noise. Zhou *et al.* [46] proposed a Relation-Aware Sampling for preserving more template-relevant points in the search area before inputting it to the backbone. By contrast, we focus on suppressing the noise after feature correlation via a Contextual Point Guided Self-Attention (CPA).

As shown in Fig. 4b, we leverage the base and expansion scale attention maps to generate the importance map $I \in \mathbb{R}^{N \times 1}$ as

$$I = \text{Mean}(\text{Attn}^{base}) + \text{Mean}(\text{Attn}^{expn}). \quad (11)$$

The higher the importance of the point, the more spatial context-aware information related to the target it contains. We sort the points according to the magnitude of their importance values. Then, all the points will be separated into G groups according to their importance. For each group with points $P_i^G \in \mathbb{R}^{G_i \times C}$, we aggregate the points into U_i clusters, which we call contextual points. Specifically, we first reshape the points as $P_i^G \in \mathbb{R}^{U_i \times C \times G_i / U_i}$. Second, a linear layer is employed to project the group to the contextual points $P_i^U \in \mathbb{R}^{U_i \times C}$. We assign fewer contextual points for the groups with lower importance, and suppress the noise feature expression. Finally, all the contextual points are concatenated and projected into Key $K^U \in \mathbb{R}^{U \times C}$ and Value $V^U \in \mathbb{R}^{U \times C}$. We project \hat{X}_{l-1} to Q and perform a multi-head attention with K^U and V^U , and an FFN is applied after attention. CPA shrinks the length of K and V, and leads to a computational cost decrease in self-attention.

3.6 Implementation Details

Backbone & Loss Functions. Following CXTrack [39], we adopt DGCNN [33] as our backbone, and apply X-RPN [39] as the RPN of our framework. We add two Shared MLP layers to X-RPN for predicting the observation angles (α) and the masks. Therefore, the overall loss is expressed as

$$\mathcal{L} = \gamma_1 \mathcal{L}_{cc} + \gamma_2 \mathcal{L}_{mask} + \gamma_3 \mathcal{L}_{alpha} + \gamma_4 \mathcal{L}_{rm} + \gamma_5 \mathcal{L}_{box}, \quad (12)$$

where \mathcal{L}_{cc} , \mathcal{L}_{mask} , \mathcal{L}_{alpha} , \mathcal{L}_{box} , and \mathcal{L}_{rm} are the loss for the coarse center, foreground mask, observation angle, targetness mask, and bounding box, respectively. We apply the L_2 loss for \mathcal{L}_{cc} , the standard cross entropy loss for \mathcal{L}_{mask} and \mathcal{L}_{rm} , and the Huber loss for \mathcal{L}_{alpha} and \mathcal{L}_{box} . γ_1 , γ_2 , γ_3 , γ_4 , and γ_5 are empirically set as 10.0, 0.2, 1.0, 1.0, and 1.0.

Training & Testing. We train our model on NVIDIA RTX-3090 GPUs with the Adam optimizer and an initial learning rate of 0.001. Due to GPU memory limitation, we construct point cloud sequences with 8 frames for training, and set $K = 2$ in training, and $K = 6$ in testing. Following existing methods [39, 45], we set N and C to 128. We stack $L = 2$ transformer layers and apply $H = 4$ heads in BEA and CPA. We adopt $G = 3$ groups in CPA, and assign [32, 64, 32] points and $U = [4, 32, 16]$ contextual points for the groups, respectively.

4 Experiments

We leverage two famous 3D tracking benchmarks of KITTI [8] and Waymo [30] to evaluate the general performance of our approach in regular 3D SOT. In addition, we establish a new KITTI-HV dataset to test our performance in high temporal variation scenarios.

Regular Datasets. The KITTI tracking dataset comprises 21 training sequences and 29 test sequences, encompassing eight object types. Following prior studies [9, 24, 34, 39, 44, 45], we use the sequences 0-16 as training data, 17-18 for validation, and 19-20 for testing. The Waymo dataset is large-scale. We adopt the approach outlined in LiDAR-SOT [21] to utilize 1121 tracklets, which are subsequently categorized into easy, medium, and hard subsets based on the number of points in the first frame of each tracklet.

HV Dataset. We build a dataset with high temporal variation for 3D SOT based on KITTI, called KITTI-HV. Although high temporal variation scenarios are present in the existing benchmarks, there is no exact threshold to determine whether the scenario is a high temporal variation scenario or not. Large point cloud variations and significant object motions are two major challenges in high temporal variation scenarios. Sampling at frame intervals is a good way to simulate these two challenges. Also, the constructed KITTI-HV can provide a preliminary platform for exploring tracking in scenarios such as skipped-tracking, edge devices, and high dynamics. For a fairer comparison with existing methods, we set the frame interval to 2, 3, 5, and 10. We set up more dense testings at low frame intervals to exploit the performance of the existing methods in point cloud variations close to smooth scenarios. We train and test all methods from scratch individually on each frame interval.

Evaluation Metrics. We employ One Pass Evaluation [38] to evaluate the different methods in terms of Success and Precision. Success is determined by measuring the Intersection Over Union between the proposed bounding box and the ground-truth (GT) bounding box. Precision is evaluated by computing the Area Under the Curve of the distance error between the centers of the two bounding boxes, ranging from 0 to 2 meters.

4.1 Comparison with the State of the Art

Results on HV tracking. We evaluate our HVTrack in 4 categories (‘Car’, ‘Pedestrian’, ‘Van’, and ‘Cyclist’) following existing methods [24, 39, 44, 45] in the KITTI-HV dataset. *The methods we choose to compare with HVTrack are the most representative SOT methods from 2020 to 2023 (Most cited methods published in each year according to Google Scholar).* As illustrated in Tab. 1, our approach consistently outperforms the state-of-the-art methods [24, 39, 44, 45] across all frame intervals, confirming the effectiveness of the proposed tracking framework for high temporal variation scenarios. Notably, the performance gap between our HVTrack and existing trackers widens as variations are exacerbated. In the particularly challenging scenario of 10 frame intervals, we achieve a substantial **9.1%**↑ improvement in success and a remarkable **10.4%**↑ enhancement

Table 1: Comparison of HVTrack with the state-of-the-art methods on each category of the KITTI-HV dataset. We construct the HV dataset KITTI-HV for training and testing by setting different frame intervals for sampling in the KITTI dataset. **Bold** and underline denote the best and second-best performance, respectively. Success/Precision are used for evaluation. Improvement and deterioration are shown in **green** and **red**, respectively.

Frame Intervals	2 Intervals					3 Intervals				
Category	Car	Pestrian	Van	Cyclist	Mean	Car	Pestrian	Van	Cyclist	Mean
Frame Number	6424	6088	1248	308	14068	6424	6088	1248	308	14068
P2B [24]	56.3/71.0	30.8/53.0	33.4/38.4	41.8/61.4	42.9/60.1	43.4/51.8	27.9/46.8	27.9/31.8	44.8/64.4	35.4/48.1
BAT [44]	61.8/74.2	36.5/61.1	26.8/30.4	54.1/78.7	47.6/64.7	51.7/61.9	31.8/53.5	24.0/28.2	50.5/72.6	40.6/55.5
M2-Track [45]	<u>63.0/76.6</u>	54.6/81.7	<u>52.8/66.5</u>	<u>68.3/89.3</u>	58.6/ <u>78.2</u>	<u>62.1/72.7</u>	<u>51.8/74.3</u>	33.6/41.6	<u>64.7/82.0</u>	<u>55.1/70.8</u>
CXTrack [39]	61.4/70.9	62.6/86.3	56.0/69.1	59.2/76.9	<u>61.4/77.5</u>	47.4/53.1	57.9/79.3	48.5/58.8	40.7/58.4	51.9/65.1
HVTrack	67.1/77.5	<u>60.0/84.0</u>	50.6/61.7	73.9/93.6	62.7/79.3	66.8/76.5	51.1/71.9	<u>38.7/46.9</u>	66.5/89.7	57.5/72.2
Improvement	4.1%↑/0.9%↑	2.6%↓/2.3%↓	6.0%↓/7.4%↓	5.6%↑/4.3%↑	1.3%↑/1.1%↑	4.7%↑/3.8%↑	6.8%↓/7.4%↓	9.8%↓/11.9%↓	1.8%↑/7.7%↑	2.4%↑/1.4%↑
Frame Intervals	5 Intervals					10 Intervals				
Category	Car	Pestrian	Van	Cyclist	Mean	Car	Pestrian	Van	Cyclist	Mean
Frame Number	6424	6088	1248	308	14068	6424	6088	1248	308	14068
P2B [24]	39.3/46.1	27.4/43.5	27.2/30.4	35.0/44.4	33.0/43.5	28.6/29.2	23.1/31.1	25.9/27.3	29.1/28.3	<u>26.0/29.8</u>
BAT [44]	44.1/51.1	21.1/32.8	26.1/29.5	35.7/46.3	32.4/41.1	30.6/33.1	21.7/ <u>29.2</u>	20.8/20.7	<u>29.3/29.1</u>	25.9/ <u>30.2</u>
M2-Track [45]	<u>50.9/58.6</u>	31.6/45.4	30.0/36.5	<u>47.4/61.0</u>	<u>40.6/51.0</u>	<u>33.0/35.1</u>	17.5/24.1	20.7/20.8	27.7/26.6	25.0/28.9
CXTrack [39]	38.6/42.2	<u>35.0/47.8</u>	21.6/24.3	25.7/33.3	35.3/42.8	30.2/32.4	18.2/21.4	17.5/17.9	27.7/26.5	23.8/26.2
HVTrack	60.3/68.9	35.1/52.1	<u>28.7/32.4</u>	58.2/71.7	46.6/58.5	49.4/54.7	<u>22.5/29.1</u>	<u>22.2/23.4</u>	39.5/45.4	35.1/40.6
Improvement	9.4%↑/10.3%↑	0.1%↑/4.3%↑	1.3%↓/4.1%↓	10.8%↑/10.7%↑	6.0%↑/7.5%↑	16.4%↑/19.6%↑	0.6%↓/0.1%↓	3.7%↓/3.9%↓	10.2%↑/16.3%↑	9.1%↑/10.4%↑

in precision. This showcases the robustness of our method in accommodating various levels of point cloud variation. Our method delivers outstanding performance on ‘Car’ and ‘Cyclist’, in which we gain a great improvement in 5 frame intervals (**9.4%↑/10.3%↑** for ‘Car’ and **10.8%↑/10.7%↑** for ‘Cyclist’) and 10 frame intervals (**16.4%↑/19.6%↑** for ‘Car’ and **10.2%↑/16.3%↑** for ‘Cyclist’). However, the challenge of tracking large objects persists in high temporal variation cases for our method. Note that the performance of CXTrack drops dramatically after 3 frame intervals. In particular, in the medium variation case of 5 frame intervals, we achieve **11.3%↑/15.7%↑** improvement in overall success/precision compared to CXTrack, despite the fact that *our HVTrack shares the same backbone and RPN with CXTrack* [39]. Furthermore, HVTrack surpasses CXTrack on ‘Car’ and ‘Cyclist’ by a very large margin (**21.7%↑/26.7%↑** for ‘Car’ and **32.5%↑/38.4%↑** for ‘Cyclist’). The distinct performance gap between HVTrack and CXTrack in HV tracking showcases the effectiveness of our feature correlation module design.

Results on regular tracking. For the KITTI dataset, we compare HVTrack with 12 top-performing trackers [7, 9, 11, 12, 17, 24, 28, 34, 39, 44–46]. As shown in Tab. 2, our overall performance is close to the SOTA tracker CXTrack [39], and achieves the second best result on the average in success (2.0%↓ w.r.t. CXTrack). Note that HVTrack outperforms TAT [17] on average (0.8%↑/0.3%↑), which utilizes temporal information by concatenating historical template features. This demonstrates our better design for leveraging the spatio-temporal context information. However, the performance of HVTrack drops when dealing with large objects (‘Van’). We conjecture this performance drop to be caused by CPA,

Table 2: Comparison of HVTrack with the SOTA methods on each category of the KITTI dataset.

Category Frame Number	Car 6424	Pedestrian 6088	Van 1248	Cyclist 308	Mean 14068
SC3D [9]	41.3/57.9	18.2/37.8	40.4/47.0	41.5/70.4	31.2/48.5
P2B [24]	56.2/72.8	28.7/49.6	40.8/48.4	32.1/44.7	42.4/60.0
3DSiamRPN [7]	58.2/76.2	35.2/56.2	45.7/52.9	36.2/49.0	46.7/64.9
MLVNet [34]	56.0/74.0	34.1/61.1	52.0/61.4	34.3/44.5	45.7/66.7
BAT [44]	60.5/77.7	42.1/70.1	52.4/67.0	33.7/45.4	51.2/72.8
PTT [28]	67.8/81.8	44.9/72.0	43.6/52.5	37.2/47.3	55.1/74.2
V2B [11]	70.5/81.3	48.3/73.5	50.1/58.0	40.8/49.7	58.4/75.2
PTTR [46]	65.2/77.4	50.9/81.6	52.5/61.8	65.1/90.5	57.9/78.1
STNet [12]	<u>72.1</u> / 84.0	49.9/77.2	58.0/70.6	73.5/93.7	61.3/80.1
TAT [17]	72.2 / <u>83.3</u>	57.4/84.4	<u>58.9</u> / 69.2	74.2 / <u>93.9</u>	64.7/82.8
M2-Track [45]	65.5/80.8	61.5/88.2	53.8/ <u>70.7</u>	73.2/93.5	62.9/ <u>83.4</u>
CXTrack [39]	69.1/81.6	67.0 / 91.5	60.0 / 71.8	74.2 / 94.3	67.5 / 85.3
HVTrack	68.2/79.2	<u>64.6</u> / <u>90.6</u>	54.8/63.8	72.4/93.7	<u>65.5</u> / 83.1

Table 3: Comparison of HVTrack with the SOTA methods on the Waymo dataset.

Method	Vehicle (185632)				Pedestrian (241168)				Mean (426800)
	Easy	Medium	Hard	Mean	Easy	Medium	Hard	Mean	
P2B [24]	57.1/65.4	52.0/60.7	47.9/58.5	52.6/61.7	18.1/30.8	17.8/30.0	17.7/29.3	17.9/30.1	33.0/43.8
BAT [44]	61.0/68.3	53.3/60.9	48.9/57.8	54.7/62.7	19.3/32.6	17.8/29.8	17.2/28.3	18.2/30.3	34.1/44.4
V2B [11]	64.5/71.5	55.1/63.2	52.0/62.0	57.6/65.9	27.9/43.9	22.5/36.2	20.1/33.1	23.7/37.9	38.4/50.1
STNet [12]	<u>65.9</u> / <u>72.7</u>	57.5 / 66.0	<u>54.6</u> / <u>64.7</u>	<u>59.7</u> / <u>68.0</u>	29.2/45.3	24.7/38.2	22.2/35.8	25.5/39.9	40.4/52.1
CXTrack [39]	63.9/71.1	54.2/62.7	52.1/63.7	57.1/66.1	35.4 / 55.3	29.7 / 47.9	<u>26.3</u> / <u>44.4</u>	30.7 / 49.4	<u>42.2</u> / <u>56.7</u>
HVTrack(Ours)	66.2 / 75.2	<u>57.0</u> / 66.0	55.3 / 67.1	59.8 / 69.7	<u>34.2</u> / <u>53.5</u>	28.7 / 47.9	26.7 / 45.2	<u>30.0</u> / <u>49.1</u>	43.0 / 58.1

which will be further explored in Sec. 4.2. For the Waymo dataset, following the benchmark setting in LiDAR-SOT [21] and STNet [12], we test our HVTrack in 2 categories (‘Vehicle’, ‘Pedestrian’) with 3 difficulty levels. All the methods are pre-trained on KITTI. The results of P2B [24], BAT [44], and V2B [11] on Waymo are provided by STNet [12]. As shown in Tab. 3, our method achieves the best performance in success (0.8% \uparrow) and precision (1.4% \uparrow). Notably, HVTrack does not surpass CXTrack and reach SOTA on the KITTI benchmark, while the opposite situation occurs in the larger dataset of Waymo. The improvement on Waymo clearly demonstrates the robustness of our method in the large-scale dataset. Also, HVTrack surpasses other SOTA methods on all categories of ‘Hard’ difficulty, revealing our excellent ability to handle sparse cases. *The experimental results show that our method can generally solve the problem of 3D SOT under various levels of point cloud variations, and achieve outstanding performance.*

4.2 Analysis Experiments

In this section, we extensively analyze HVTrack via a series of experiments. All the experiments are conducted on KITTI-HV with 5 frame intervals unless otherwise stated.

Ablation Study. We conduct experiments to analyze the effectiveness of different modules in HVTrack. As shown in Tab. 4, we respectively ablate OM, BEA, and CPA from HVTrack. We only ablate OM in RPM because LM and MM

Table 4: Ablation analysis of HVTrack.

OM	BEA	CPA	Car	Pedestrian	Van	Cyclist	Mean
	✓	✓	60.0/ <u>69.0</u>	33.9/50.0	<u>28.4</u> /32.2	<u>54.2</u> / <u>67.1</u>	45.8/ <u>57.5</u>
✓		✓	60.3 / 69.4	<u>35.0</u> / <u>50.2</u>	26.7/30.7	43.9/61.5	<u>46.0</u> / <u>57.5</u>
✓	✓		58.2/66.9	34.7/49.8	28.1/ 33.5	47.7/63.9	45.1/56.5
✓	✓	✓	60.3 /68.9	35.1 / 52.1	28.7 / <u>32.4</u>	58.2 / 71.7	46.6 / 58.5

Table 5: Ablation experiment of BEA. ‘Base’/‘Expansion’ denotes only using the base/expansion branch in BEA.

Category	Car	Pedestrian	Van	Cyclist	Mean
Base	60.3/69.4	35.0/50.2	26.7/30.7	43.9/61.5	46.0/57.5
Expansion	60.0/68.6	34.7/50.5	31.4/36.8	54.5/67.5	46.4/57.9

serve as the template and are the indivisible parts of HVTrack. BEA and CPA are replaced by vanilla cross-attention and self-attention. In general, all components have been proven to be effective; removing an arbitrary module degrades the ‘mean’ performance.

Analysis Experiment of BEA. The performance slightly drops on the ‘Car’ when we apply BEA on HVTrack as shown in Tab. 4. We conjecture this to be caused by the side effect of aggregating larger scale features in BEA, which will involve more background noise at each point. Further, ‘Car’ has a medium size and does not have the distraction of crowded similar objects like small objects (‘Pedestrian’ and ‘Cyclist’), nor does it require a larger receptive field like large objects (‘Van’). To verify this issue, we further analyze each branch of BEA as shown in Tab. 5. ‘Pedestrian’, ‘Van’, and ‘Cyclist’ benefit from the expansion branch and achieve a better performance compared to using only the base branch in BEA. On the other hand, the performance in the ‘Car’ category has the opposite behavior to the other categories. The experimental results validate our hypothesis that BEA is beneficial to small and large objects, while negatively affecting medium-sized objects.

Analysis Experiment of CPA. Our method yields better results on ‘Van’ after we remove CPA as shown in Tab. 4, which reveals the relation between CPA and the large object tracking challenge. We believe that this is caused by the suppressing strategy in CPA. Large objects usually have more points, and under the same probability of misclassification of importance, they will have more foreground points assigned as low importance in the attention map, resulting in a part of useful information being suppressed in CPA. As shown in Fig. 5b, the importance conflict in the object leads to tracking failure. That part of the information will be further suppressed when stacking multiple transformer layers. However, the performance drops in other categories, without CPA to suppress the background noise for medium and small objects. As shown in Fig. 5a, most of the background points are assigned with low importance and suppressed in the success case, which proves our idea of CPA.




Fig. 5: The attention maps of ‘Van’ in CPA.

Table 6: Results of HVTrack when using different memory sizes. We train HVTrack with a memory size of 2, and evaluate it with memory sizes ranging from 1 to 8 .

Memory Size	Car	Pedestrian	Van	Cyclist	Mean
1	58.3/66.5	30.9/46.2	26.8/29.8	57.1/70.5	43.6/54.6
2	58.6/67.0	31.7/47.9	27.1/30.6	57.6/70.9	44.1/55.6
3	59.2/67.6	33.8/49.9	27.7/31	55.8/67.7	45.3/56.7
4	<u>60.0/68.5</u>	33.7/50.6	29.5/33.6	57.9/71.3	45.9/57.7
5	<u>60.0/68.5</u>	33.8/51.2	<u>28.7/32.6</u>	57.8/70.8	45.8/57.9
6	60.3/68.9	<u>35.1/52.1</u>	<u>28.7/32.4</u>	58.2/71.7	46.6/58.5
7	59.7/68.2	35.6/52.9	28.0/31.5	<u>58.1/71.4</u>	<u>46.4/58.4</u>
8	59.8/68.3	<u>35.1/52.4</u>	28.2/32.0	<u>58.1/71.4</u>	46.3/58.3

Memory Size. Intuitively, trackers will achieve better performance when leveraging more temporal information. However, the performance of the trackers cannot continuously improve with the accumulation of historical information, due to inaccuracies in the historical tracklets. As shown in Tab. 6, we train HVTrack with a memory size of 2 due to the GPU memory limitation, and evaluate it with memory sizes from 1 to 8. The performance peaks for a memory size of 6, which is consistent with our assumption. Thus, we set 6 as our memory size and achieve a tracking speed of 31 FPS.

5 Conclusion

In this paper, we have explored a new task in 3D SOT, and presented the first 3D SOT framework for high temporal variation scenarios, HVTrack. Its three main components, RPM, BEA, and CPA, allow HVTrack to achieve robustness to point cloud variations, similar object distractions, and background noise. Our experiments have demonstrated that HVTrack significantly outperforms the state of the art in high temporal variation scenarios, and achieves remarkable performance in regular tracking.

Limitation. Our CPA relies on fixed manual hyperparameters to suppress noise. This makes it difficult to balance the performance in different object and search area sizes, leading to a performance drop in tracking large objects. In the future, we will therefore explore the use of a learnable function to replace the manual hyperparameters in CPA and overcome the large object tracking challenge.

Acknowledgements

This work is supported in part by the National Natural Science Foundation of China (NFSC) under Grants 62372377 and 62176242.

References

1. Chen, X., Shi, S., Zhang, C., Zhu, B., Wang, Q., Cheung, K.C., See, S., Li, H.: Trajectoryformer: 3d object tracking transformer with predictive trajectory hypotheses. arXiv preprint arXiv:2306.05888 (2023)
2. Cheng, R., Wang, X., Sohel, F., Lei, H.: Topology-aware universal adversarial attack on 3d object tracking. *Visual Intelligence* **1**(1), 31 (2023)
3. Chiu, H.k., Prioletti, A., Li, J., Bohg, J.: Probabilistic 3d multi-object tracking for autonomous driving. arXiv preprint arXiv:2001.05673 (2020)
4. Chung, J., Gulcehre, C., Cho, K., Bengio, Y.: Empirical evaluation of gated recurrent neural networks on sequence modeling. arXiv preprint arXiv:1412.3555 (2014)
5. Cui, Y., Fang, Z., Shan, J., Gu, Z., Zhou, S.: 3d object tracking with transformer. arXiv preprint arXiv:2110.14921 (2021)
6. Ding, S., Rehder, E., Schneider, L., Cordts, M., Gall, J.: 3dmotformer: Graph transformer for online 3d multi-object tracking. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 9784–9794 (2023)
7. Fang, Z., Zhou, S., Cui, Y., Scherer, S.: 3d-siamrpn: An end-to-end learning method for real-time 3d single object tracking using raw point cloud. *IEEE Sensors Journal* **21**(4), 4995–5011 (2020)
8. Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? the KITTI vision benchmark suite. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 3354–3361 (2012)
9. Giancola, S., Zarzar, J., Ghanem, B.: Leveraging shape completion for 3d siamese tracking. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 1359–1368 (2019)
10. Guo, Z., Mao, Y., Zhou, W., Wang, M., Li, H.: Cmt: Context-matching-guided transformer for 3d tracking in point clouds. In: *European Conference on Computer Vision*. pp. 95–111. Springer (2022)
11. Hui, L., Wang, L., Cheng, M., Xie, J., Yang, J.: 3d siamese voxel-to-bev tracker for sparse point clouds. *Advances in Neural Information Processing Systems* **34**, 28714–28727 (2021)
12. Hui, L., Wang, L., Tang, L., Lan, K., Xie, J., Yang, J.: 3d siamese transformer network for single object tracking on point clouds. arXiv preprint arXiv:2207.11995 (2022)
13. Jiao, L., Wang, D., Bai, Y., Chen, P., Liu, F.: Deep learning in visual tracking: A review. *IEEE transactions on neural networks and learning systems* (2021)
14. Jiayao, S., Zhou, S., Cui, Y., Fang, Z.: Real-time 3d single object tracking with transformer. *IEEE Transactions on Multimedia* (2022)
15. Kapania, S., Saini, D., Goyal, S., Thakur, N., Jain, R., Nagrath, P.: Multi object tracking with uavs using deep sort and yolov3 retinanet detection framework. In: *Proceedings of the 1st ACM Workshop on Autonomous and Intelligent Mobile Systems*. pp. 1–6 (2020)

16. Kart, U., Lukezic, A., Kristan, M., Kamarainen, J.K., Matas, J.: Object tracking by reconstruction with view-specific discriminative correlation filters. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1339–1348 (2019)
17. Lan, K., Jiang, H., Xie, J.: Temporal-aware siamese tracker: Integrate temporal context for 3d object tracking. In: Proceedings of the Asian Conference on Computer Vision. pp. 399–414 (2022)
18. Luo, C., Yang, X., Yuille, A.: Exploring simple 3d multi-object tracking for autonomous driving. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 10488–10497 (2021)
19. Machida, E., Cao, M., Murao, T., Hashimoto, H.: Human motion tracking of mobile robot with kinect 3d sensor. In: Proceedings of SICE Annual Conference (SICE). pp. 2207–2211. IEEE (2012)
20. Nishimura, H., Komorita, S., Kawanishi, Y., Murase, H.: Sdof-tracker: Fast and accurate multiple human tracking by skipped-detection and optical-flow. *IEICE TRANSACTIONS on Information and Systems* **105**(11), 1938–1946 (2022)
21. Pang, Z., Li, Z., Wang, N.: Model-free vehicle tracking and state estimation in point cloud sequences. In: 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 8075–8082. IEEE (2021)
22. Qi, C.R., Litany, O., He, K., Guibas, L.J.: Deep hough voting for 3d object detection in point clouds. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9277–9286 (2019)
23. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems* **30** (2017)
24. Qi, H., Feng, C., Cao, Z., Zhao, F., Xiao, Y.: P2b: Point-to-box network for 3d object tracking in point clouds. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6329–6338 (2020)
25. Ren, C., Xu, Q., Zhang, S., Yang, J.: Hierarchical prior mining for non-local multi-view stereo. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3611–3620 (2023)
26. Ren, S., Yang, X., Liu, S., Wang, X.: Sg-former: Self-guided transformer with evolving token reallocation. *arXiv preprint arXiv:2308.12216* (2023)
27. Sadjadpour, T., Li, J., Ambrus, R., Bohg, J.: Shasta: Modeling shape and spatio-temporal affinities for 3d multi-object tracking. *IEEE Robotics and Automation Letters* (2023)
28. Shan, J., Zhou, S., Fang, Z., Cui, Y.: Ptt: Point-track-transformer module for 3d single object tracking in point clouds. In: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 1310–1316 (2021)
29. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* **15**(1), 1929–1958 (2014)
30. 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: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2446–2454 (2020)
31. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
32. Wang, Q., Chen, Y., Pang, Z., Wang, N., Zhang, Z.: Immortal tracker: Tracklet never dies. *arXiv preprint arXiv:2111.13672* (2021)

33. Wang, Y., Sun, Y., Liu, Z., Sarma, S.E., Bronstein, M.M., Solomon, J.M.: Dynamic graph cnn for learning on point clouds. *ACM Transactions on Graphics (tog)* **38**(5), 1–12 (2019)
34. Wang, Z., Xie, Q., Lai, Y.K., Wu, J., Long, K., Wang, J.: Mlvsnet: Multi-level voting siamese network for 3d visual tracking. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 3101–3110 (2021)
35. Weng, X., Wang, J., Held, D., Kitani, K.: 3d multi-object tracking: A baseline and new evaluation metrics. In: *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. pp. 10359–10366. IEEE (2020)
36. Weng, X., Wang, Y., Man, Y., Kitani, K.M.: Gnn3dmot: Graph neural network for 3d multi-object tracking with 2d-3d multi-feature learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 6499–6508 (2020)
37. Wu, Q., Yang, J., Sun, K., Zhang, C., Zhang, Y., Salzmann, M.: Mixcycle: Mixup assisted semi-supervised 3d single object tracking with cycle consistency. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 13956–13966 (2023)
38. Wu, Y., Lim, J., Yang, M.H.: Online object tracking: A benchmark. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 2411–2418 (2013)
39. Xu, T.X., Guo, Y.C., Lai, Y.K., Zhang, S.H.: Cxtrack: Improving 3d point cloud tracking with contextual information. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1084–1093 (2023)
40. Yin, T., Zhou, X., Krahenbuhl, P.: Center-based 3d object detection and tracking. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 11784–11793 (2021)
41. Yoo, J.S., Lee, H., Jung, S.W.: Video object segmentation-aware video frame interpolation. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 12322–12333 (2023)
42. Zarzar, J., Giancola, S., Ghanem, B.: Efficient bird eye view proposals for 3d siamese tracking. *arXiv preprint arXiv:1903.10168* (2019)
43. Zhang, X., Yang, J., Zhang, S., Zhang, Y.: 3d registration with maximal cliques. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 17745–17754 (2023)
44. Zheng, C., Yan, X., Gao, J., Zhao, W., Zhang, W., Li, Z., Cui, S.: Box-aware feature enhancement for single object tracking on point clouds. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 13199–13208 (2021)
45. Zheng, C., Yan, X., Zhang, H., Wang, B., Cheng, S., Cui, S., Li, Z.: Beyond 3d siamese tracking: A motion-centric paradigm for 3d single object tracking in point clouds. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 8111–8120 (2022)
46. Zhou, C., Luo, Z., Luo, Y., Liu, T., Pan, L., Cai, Z., Zhao, H., Lu, S.: Pptr: Relational 3d point cloud object tracking with transformer. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 8531–8540 (2022)
47. Zhou, X., Koltun, V., Krähenbühl, P.: Tracking objects as points. In: *European conference on computer vision*. pp. 474–490. Springer (2020)