

Synchronize Feature Extracting and Matching: A Single Branch Framework for 3D Object Tracking

Teli Ma^{1,2*}, Mengmeng Wang^{2*}, Jimin Xiao³, Huifeng Wu⁴, Yong Liu^{2†}

¹The Hong Kong University of Science and Technology, Guangzhou ²Zhejiang University

³Xi'an Jiaotong-Liverpool University ⁴Hangzhou Dianzi University

tma184@connect.hkust-gz.edu.cn mengmengwang@zju.edu.cn

jimin.xiao@xjtlu.edu.cn whf@hdu.edu.cn yongliu@iipc.zju.edu.cn

Abstract

*Siamese network has been a de facto benchmark framework for 3D LiDAR object tracking with a shared-parametric encoder extracting features from template and search region, respectively. This paradigm relies heavily on an additional matching network to model the cross-correlation/similarity of the template and search region. In this paper, we forsake the conventional Siamese paradigm and propose a novel single-branch framework, **SyncTrack**, synchronizing the feature extracting and matching to avoid forwarding encoder twice for template and search region as well as introducing extra parameters of matching network. The synchronization mechanism is based on the dynamic affinity of the Transformer, and an in-depth analysis of the relevance is provided theoretically. Moreover, based on the synchronization, we introduce a novel Attentive Points-Sampling strategy into the Transformer layers (APST), replacing the random/Farthest Points Sampling (FPS) method with sampling under the supervision of attentive relations between the template and search region. It implies connecting point-wise sampling with the feature learning, beneficial to aggregating more distinctive and geometric features for tracking with sparse points. Extensive experiments on two benchmark datasets (KITTI and NuScenes) show that SyncTrack achieves state-of-the-art performance in real-time tracking.*

1. Introduction

With recent advances in autonomous driving, 3D vision tasks based on LiDAR are becoming increasingly popular in the visual community. Among these tasks, 3D LiDAR single object tracking (SOT) aims to track a specific target object in a 3D video with the knowledge of 3D bounding

*Equal Contribution

†Corresponding Author

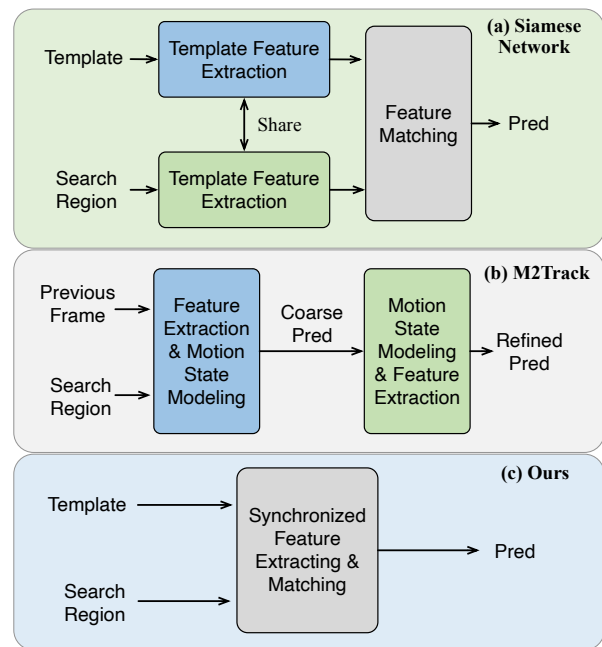


Figure 1. The comparison of (a) Siamese network based trackers, (b) previous single-branch, two-stage framework M2Track [51] with (c) our SyncTrack, which is a single-branch and single-stage framework.

box in the initial frame. This task meets numerous challenges, such as LiDAR point cloud sparsity, occlusions, and fast motions.

Most existing methods [32, 50, 19, 12, 33, 52, 20, 9, 39] of 3D SOT mainly adopt a Siamese-like backbone and incorporate an additional matching network to cope with the tracking challenges as shown in Fig. 1(a). Trackers based on the Siamese-like backbone separate feature extraction of template and search region, forwarding the two kinds of features with shared model parameters, respectively. Subsequently, an extra matching network is introduced to fuse

the extracted template and search region features to model the correlation or similarity between them. However, such a paradigm restricts the feature interaction to a post-matching network, correlating the template and search region insufficiently with merely the high-level extracted features. In other words, the matching process posterior to the encoder is incapable of modeling the relations of multi-scale features intra-backbone. Moreover, a standalone matching network results in extra model parameters and computational overheads, let alone the double forwarding process of the Siamese-backbone to extract the template and search region features. M2Track [51] proposed a motion-centric paradigm to replace the Siamese-like structure, constructing a spatial-temporal point cloud to predict the motion. However, they still rely heavily on an additional motion transformation network, which requires extra training input, to integrate the extracted template features into the search region, and another two-stage refinement network is leveraged to ensure the performance as illustrated in Fig. 1(b). Based on aforementioned problems, we ask the question: *Can feature extracting and matching be conducted simultaneously in a simple way?*

The answer is **Yes** and the solution is implied in the **dynamic global reasoning property** of Transformer [35, 7, 10]. Specifically, the affinity matrix of all tokens can be constructed dynamically via continuous computation of the *key* and *query* vectors in the attention mechanism. The spatial context is aggregated using affinity to attend features. Intuitively, this affinity matrix can intrinsically serve as the **matching matrix** for intermediate feature interactions between the template and search region if we merge them into one input of the Transformer layers. Therefore, we propose a single-branch and single-stage framework equipped with a Transformer-based backbone instead of the conventional Siamese-like PointNet++ [31] backbone, as shown in Fig 1(c). The framework is dubbed as SyncTrack, as the Transformer backbone synchronizes the feature extracting and matching process. The SyncTrack is composed of a simple backbone and prediction head, omitting the complex matching network design and motion state estimation, depending merely on point-wise features.

However, 3D point clouds have unique properties such as sparsity [54, 24], density variance, and implicit geometric features in data locality [29]. For example, 51% samples of KITTI [11] *Car* category have less than 100 points [19]. These problems are further aggravated when point clouds are grouped and down-sampled to formulate multi-scale point-wise feature maps. The phenomena stresses the vitality of point clouds sampling, aiming to improve the point-wise perception efficiency with limited points. PTTR [52] proposed to sample the input point clouds before the backbone, utilizing the L_2 distance as the similarity metric for sampling. However, as down-sampling layer by layer is

essential in tracking backbone for multi-scale feature fusion to strengthen representation, it is reasonable to consider the sampling strategy in the backbone. Therefore, we propose the attentive sampling strategy based on the attention map between the template and search region, and equip each Transformer layer with our sampling module as shown in Fig. 2(b). We name the Transformer containing Attentive Points-Sampling as APST. Specifically, the attentive response from template tokens to search region tokens is considered, as the positively respond tokens are more likely to be in the foreground and should be preserved for feature extracting. By contrast, as Fig. 3 shows, random sampling easily falls into the perceptive confusion as selected points hardly contain geometric features due to randomness.

The main contributions of our paper can be summarized as follows:

- We introduce a single-branch and single-stage framework for real-time 3D LiDAR SOT dubbed SyncTrack, without Siamese-like forward propagation and a standalone matching network. We ingeniously leverage the dynamic affinity characteristic of the self-attention mechanism to synchronize the feature extracting and matching. A detailed analysis is provided to explain the synchronizing mechanism.
- We propose a novel APST to build the backbone, replacing the random/FPS¹ down-sampling of point clouds with attentive sampling to preserve more target-relevant points, and thus improving the perceptive capability of feature extracting.
- Extensive results show that our method has achieved new state-of-the-art performance on the KITTI and NuScenes datasets in real-time tracking, up to 2.8% and 1.4% on *mean* results with a high-speed of around 45 fps. Besides, SyncTrack exhibits good scalability in both width and depth.

2. Related Works

2.1. 2D Visual Tracking

Advances in 2D video [53, 27, 26] and object tracking approaches [30, 16, 36, 23, 34, 48] stimulate the development of the 3D tracking, and the methods are evolving all the time. Early 2D SOT methods mainly focus on the classifier design like structured output SVM [36, 14], correlation filter [17, 25]. With the prevalence of deep learning, end-to-end trainable trackers emerged in the visual community. The Siamese-like structure [6, 18, 15, 23, 46, 42, 28] based methods have been popular in the tracking field and

¹In this paper, we use FPS to refer to Farthest Points Sampling and fps to denote frames per second.

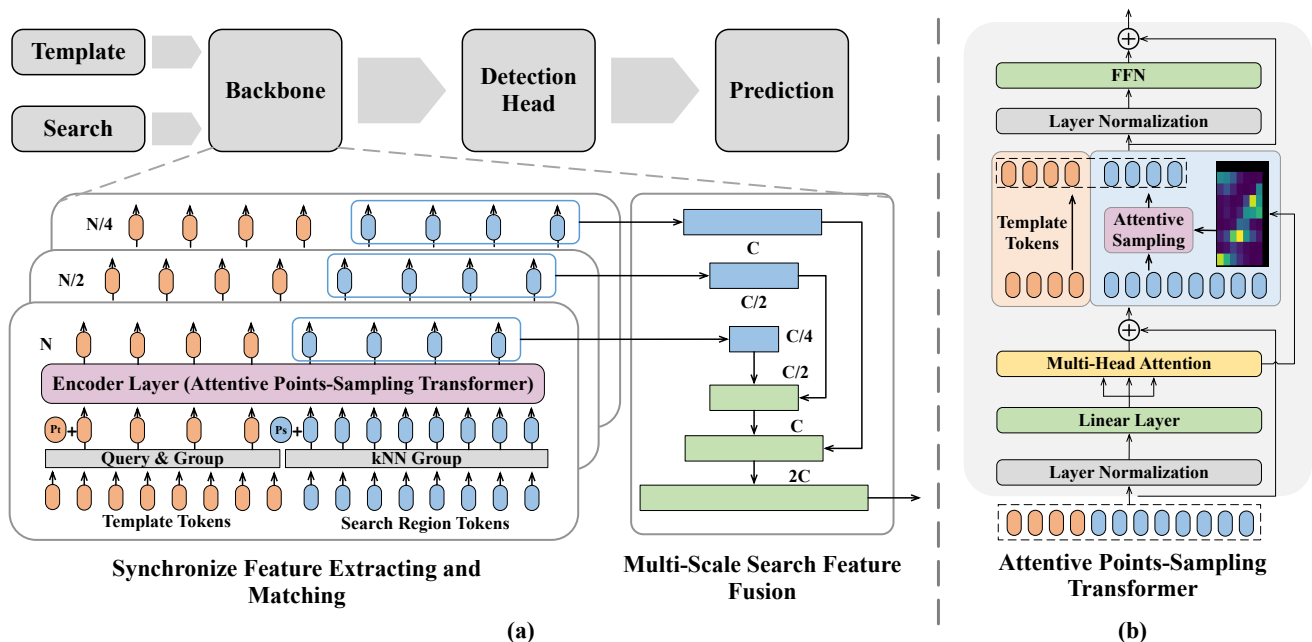


Figure 2. (a). The overall framework of SyncTrack and specific illustration of the single-branch backbone, the template and search region input are concatenated for feature extraction and matching synchronously. (b). The structure of the Attentive Points-Sampling Transformer (APST), which samples search region points after multi-head attention.

have numerous variants. SiamFC [1] is a pioneering work integrating feature correlation into a fully convolutional Siamese network for visual tracking. Subsequently, improvements like introducing detection components such as region proposal network [23, 22, 8], discriminating fore and back-ground [49], anchor-free detecting [13] *etc.*. Recently, vision transformers have been introduced into 2D SOT [4, 38, 43, 41] to exploit the long-range modeling of the attention mechanism to fuse the features effectively. Moreover, the one-stream network trackers based on transformers are proposed in 2D tracking. The OSTrack [44] proposes joint feature learning and relation modeling, masking a proportion of image patches to save computational overheads. SimTrack [3] concatenates the template and search input, improving the patch embedding method with a Foveal window strategy. Our work draws inspiration from the aforementioned trend in 2D vision. However, it is specifically tailored to address the challenges of the 3D Single Object Tracking (SOT) problem, taking into account the unique characteristics of point clouds. Furthermore, we analyze the success of the transformer-based single-branch backbone, attributing to the dynamic affinity of the attention mechanism, which enables the synchronization of feature extraction and matching processes.

2.2. 3D Visual Tracking

In this paper, we only discuss LiDAR-based 3D object tracking. Till to now, almost all the 3D tracking methods [12, 32, 9, 45, 5, 33, 52, 50, 19, 20, 37] are based on the Siamese structure. The pioneering work in this field is the SC3D [12], which initially defines the task. SC3D employs cosine similarity to measure the resemblance between template and search region features, and incorporates shape completion during training to enhance appearance refinement. Trackers following SC3D make advancements from two perspectives. Firstly, they enhance the matching network [32, 50, 19, 52, 5, 39, 51]. For instance, MLVSNet [39] uses the CBAM module [40] to enhance the vote cluster features with both channel and spatial attention. STNet [20] employs cross- and self-attention modules to enhance the interaction between the extracted template and search region features, boosting their feature-level integration. M2Track introduces a motion-centric paradigm and motion estimation module to correlate the template and search features instead of appearance matching. All these methods depend on a standalone module to match the features. Secondly, the trackers [32, 9, 19, 39, 33] have attempted to improve the prediction head part. P2B [32] employs Hough Voting to predict the target location, and [50, 33, 39] all follow the voting strategy to make predictions. 3D-SiamRPN [9] uses an RPN head to predict the final results. LTTR [9] and V2B [19] use center-based regres-

sions to predict several object properties. However, few efforts are made in exploring the encoder/backbone of trackers as PointNet++ [31] is the feature extractor by default. In this study, we focus on the backbone design and incorporate the matching process directly into the backbone, significantly streamlining the tracker network.

3. Method

In this section, we first define the 3D SOT task in Sec. 3.1. Then, a detailed introduction of the single-branch framework is included in the Sec. 3.2. Based on the single-branch framework, how the feature extracting and matching are synchronized is elaborated in (Sec. 3.3). Moreover, we propose the Attentive Points-Sampling Transformer to build the single-branch backbone and sampling search region tokens with a strategy of attentive sampling, as shown in Sec. 3.4. The decoder head and losses are described in Sec. 3.5.

3.1. Problem Definition

In the configuration of 3D LiDAR single object detection (SOT) task, the 3D bounding box (BBox) is defined as $(x, y, z, w, l, h, \theta) \in \mathbb{R}^7$, where the (x, y, z) represents the coordinate center of the BBox and $(w, l, h), \theta$ stand for the BBox size and heading angle (the rotation around the *up-axis*) respectively. Generally, the BBox size is assumed to be fixed by default even when the target object is non-rigid, thus minimizing the dimensions of BBox from \mathbb{R}^7 to \mathbb{R}^4 . Given a sequence of temporally-connected point clouds $\{\mathcal{P}_i\}_{i=1}^T$ (T is the number of points in each frame) and an initial BBox \mathcal{B}_1 of the target, the goal of SOT is to localize the target BBoxes $\{\mathcal{B}_i\}_{i=2}^T$ in all frames online. Following the previous manner, a template point cloud $\mathcal{P}^t = \{p_i^t\}_{i=1}^{N_t}$ and a search region $\mathcal{P}^s = \{p_i^s\}_{i=1}^{N_s}$ are generated, where N_t and N_s are number of template and search region points. The template P^t is generated by cropping and centering the target in the initial frame based on the initial BBox.

3.2. Single-Branch Structure

We propose to replace the conventional *Siamese-like backbone* paradigm with a sole backbone, eliminating the double forward process of Siamese structure. Therefore, template and search region seeds are concatenated to do forward propagation. The Transformers' property of long-range relation-modeling is the intrinsic merit to tackle the concatenated template and search region seeds. Based on that, we leverage the self-attention modules to build the single-branch backbone as shown in Fig. 2(a). In our approach, we utilize a Query & Group module to sample the template seeds, denoted as $\mathcal{P}^t \in \mathbb{R}^{N_t \times C}$, using the FPS method before jointly forwarding. This module also groups the k-nearest points to aggregate features. However, when

it comes to the search region point cloud, represented as $\mathcal{P}^s \in \mathbb{R}^{N_s \times C}$, we solely employ the Group module to aggregate neighbor information without reducing the number of search points. Subsequently, the template and search seeds are concatenated, incorporating a joint parametric positional embedding for the localization of tokens. This process is as follows:

$$\begin{aligned} \mathcal{T}^t &= \text{kNN}(\text{FPS}(\mathcal{P}^t)), \mathcal{P}^t \in \mathbb{R}^{N_t' \times C'}, N_t' < N_t, \\ \mathcal{T}^s &= \text{kNN}(\mathcal{P}^s), \mathcal{P}^s \in \mathbb{R}^{N_s \times C'}, \\ \mathcal{T}^{ts} &= [\mathcal{T}^t; \mathcal{T}^s] + \mathbf{P}_e, \mathcal{T}^{ts} \in \mathbb{R}^{(N_t + N_s) \times C'}. \end{aligned} \quad (1)$$

Afterward, linear layers are leveraged to project the input tokens into *query*, *key*, and *value* latent and the head-wise joint attention map is calculated to model the intra- & inter-relations of template and search tokens.

$$\begin{aligned} Q^{ts}, K^{ts}, V^{ts} &= W_q(\mathcal{T}^{ts}), W_k(\mathcal{T}^{ts}), W_v(\mathcal{T}^{ts}), \\ \mathcal{A}_m^{ts} &= \text{Softmax} \frac{Q_m^{ts} (K_m^{ts})^\top}{\sqrt{C'/M}}. \end{aligned} \quad (2)$$

Based on the joint head-wise attention map \mathcal{A}_m^{ts} of template and search tokens, the features extracted by multi-head attention are:

$$\mathcal{T}'^{ts} = [\mathcal{A}_1^{ts} V_1^{ts}, \mathcal{A}_2^{ts} V_2^{ts}, \dots, \mathcal{A}_M^{ts} V_M^{ts}] \mathcal{W} + \mathcal{T}^{ts}, \quad (3)$$

where M is the number of the attention heads and \mathcal{W} is the weight of a MLP.

3.3. Synchronize Feature Extracting and Matching

We illustrate how the single-branch structure can synchronize the process of feature extracting and matching with a simple backbone. The formula of joint attention \mathcal{A}_m^{ts} in Eq. 2 can be expanded as (σ represents *Softmax* operation):

$$\begin{aligned} [\mathcal{A}_m^t; \mathcal{A}_m^s] &= \sigma \frac{[Q_m^t; Q_m^s][K_m^t; K_m^s]^\top}{\sqrt{C'/M}} \\ &= \left[\sigma \frac{Q_m^t (K_m^t)^\top}{\sqrt{C'/M}}, \sigma \frac{Q_m^t (K_m^s)^\top}{\sqrt{C'/M}}; \sigma \frac{Q_m^s (K_m^t)^\top}{\sqrt{C'/M}}, \sigma \frac{Q_m^s (K_m^s)^\top}{\sqrt{C'/M}} \right]. \end{aligned} \quad (4)$$

The new extracted search region features can be obtained based on Eq 4 as:

$$\begin{aligned} \mathcal{T}'_m^s &= \sigma \frac{Q_m^s (K_m^t)^\top}{\sqrt{C'/M}} V_m^t + \sigma \frac{Q_m^s (K_m^s)^\top}{\sqrt{C'/M}} V_m^s, \\ [V_m^t; V_m^s] &= V_m^{ts}, \end{aligned} \quad (5)$$

depending on the projected features $[V_m^t, V_m^s]$ of both the template and search region from last layer. The attention queried from search region to template, $\sigma \frac{Q_m^s (K_m^t)^\top}{\sqrt{C'/M}}$, is the

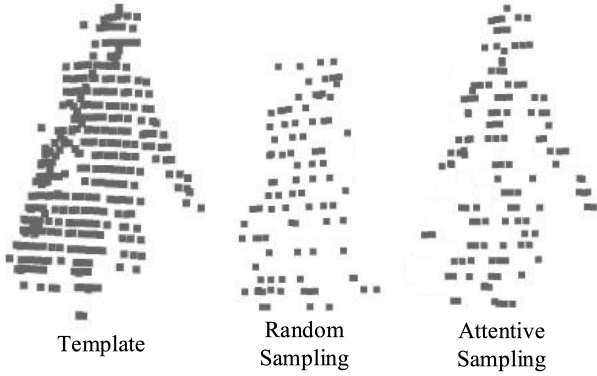


Figure 3. **Comparisons of random sampling with our attentive sampling in point-wise down-sampling.** We find that our attentive sampling approach excels in selecting more representative points, including those that contribute to the formation of the head, arms, and legs, which contain distinctive geometric features. In contrast, random sampling include more points that are far away from the target (we have removed such points for better visualization).

matching matrix that guides to aggregate highly-relevant template features. Moreover, the **matching matrix** is dynamic, which is determined by the changing *query* and *key* latent of search and template features as:

$$\sigma \frac{Q_m^s (K_m^t)^\top}{\sqrt{C'/M}} = \sigma \frac{W_{q,m}^s \mathcal{T}_m^s (\mathcal{T}_m^t)^\top (W_{k,m}^t)^\top}{\sqrt{C'/M}}, \quad (6)$$

i.e. $\sigma \frac{Q_m^s (K_m^t)^\top}{\sqrt{C'/M}} \propto [\mathcal{T}_m^s; \mathcal{T}_m^t].$

To conclude, the synchronization attributes to the dynamic mechanism of **matching matrix**, continuously adapting the matching relations according to the extracted features of template and search region.

Comparisons with Siamese Network. Previous Siamese-like trackers can be summarized as the paradigm of *Extracting then Matching*. If we name the feature extracting and matching as ϕ and δ , then the objective of this paradigm can be concluded as $\max(\delta(\phi(\mathcal{P}^t), \phi(\mathcal{P}^s)))$ for simplicity. It means the model training is to shorten the distances between correlated parts of the template and search region based on the extracted features from the backbone. However, compared to our approach, this matching mechanism is relatively *static* since it occurs only after feature extraction, resulting in inadequate modeling of inter-backbone relations. On the other hand, our single-branch framework facilitates dynamic interaction between the search region seeds and the template across all layers of the backbone. This allows for comprehensive learning of relations, encompassing both the local representation from early layers and the global representation from later ones.

3.4. Attentive Points-Sampling Transformer

Based on the synchronizing mechanism of feature extracting and matching, we propose to replace the Point-wise Transformer [47] with a novel Attentive Points-Sampling Transformer (APST). APST is proposed based on the observation that backbones utilized in previous 3D LiDAR trackers always adopt a non-parametric strategy of down-sampling the search region features/tokens with farthest point sampling (FPS) or randomly sampling to reduce the points as introduced in PointNet++ [31]. Nevertheless, this non-parametric sampling method lacks learnability and controllability since no parameters associated with the sampling process are updated during model training. Consequently, the final performance is compromised due to the significant influence of the points-center in LiDAR, which determines the effectiveness of extracting features around the foreground, as shown in Fig. 3.

Therefore, we introduce the APST, which involves the selection of points based on attentive relations between the tokens of the template and search region as illustrated in sub-figure (b) of Fig. 2. The attentive response of template tokens reacting to search region tokens is considered as positively responding search tokens are more likely to be in the foreground. In that case, we segment the attention map and separate the $\mathcal{A}_m^{t \rightarrow s} = \sigma \frac{Q_m^t (K_m^s)^\top}{\sqrt{C'/M}}$ from the Eq. 4, averaging along the dimension of template-tokens attention among all transformer heads to acquire the response scores. To ensure the maximum response of search tokens to template ones, a group of search region indexes, named Ω_s , is dynamically updated as:

$$\Omega_s^* = \arg \max_{\Omega_s} \frac{1}{N_t} \frac{1}{M} \sum_{t=0}^{N_t} \sum_{m=0}^M \frac{Q_m^t}{\sqrt{C'/M}} (K_m^s)^\top \Big|_{s \in \Omega_s}, \quad (7)$$

Based on the optimal solution Ω_s^* , the search region tokens are sampled to decrease the number of tokens after multi-head self-attention and concatenate with template ones. Note that the template tokens are sampled with the FPS method as introduced in Sec. 3.2, and only the search region tokens are sampled with the attentive sampling method.

The dynamic-affinity property of the Transformer suggests that the attention map is influenced by the latent tokens projected through learnable linear layers. Hence, it is reasonable to assert that sampling tokens guided by the attention map is linked to model learning, as the attention map is generated based on updated parameters. As a result, this type of token/point sampling proves advantageous in providing prior knowledge for centroid selection and enhancing the efficient aggregation of representational information, as shown in Fig. 3.

3.5. Decoder and Losses

With the encoded point-wise features of various scales, a multi-scale feature fusion module is adopted to fuse these features and output the features with the same origin input size, feeding into the decoder part for final predictions. Following the V2B [19], the features are voxelized as a volumetric representation and 3D convolutions are utilized on the encoded features. Afterward, the BEV feature maps are acquired by pooling on the z -axis for regression. Focal loss and L1 loss are leveraged for classification and BBox center offset and rotation regression, respectively. A detailed introduction is presented in the supplementary.

4. Experiments

4.1. Experiment Setups

Implementation Details. We set the number of input points as $N_t = 512$ and $N_s = 1024$ for template and search regions by randomly duplicating and discarding, respectively. The encoder backbone is merely consisted of three layers with Attentive Points-Sampling Transformers, and the number of both template and search region points output by each layer are 256, 128 and 64. Moreover, each layer’s feature dimensions in the encoder are 32, 64 and 128, respectively, whereas the final features for the prediction head are with 32 channels. Note that the heads for all APST are 2 by default. In the voxelization process, the region $[(x_{min}, x_{max}), (y_{min}, y_{max}), (z_{min}, z_{max})]$ is defined as $[(-5.6, 5.6), (-3.6, 3.6), (-2.4, 2.4)]$ to contain most target points. The voxel size (v_x, v_y, v_z) is set to $(0.3, 0.3, 0.3)$. For the detection head, four decomposed 3D (stride of 2, 1, 2, 1 along the z -axis) and 2D convolution blocks (stride of 2, 1, 1, 2) are leveraged to strengthen the feature aggregation.

Training and Testing. We train the model for 40 epochs with a batch size of 64. The Adam optimizer [21] is adopted with the initial learning rate of 0.001 and reduced by 5 every 10 epochs (every 2 epochs for nuScenes). The classification loss has a weight λ_{cls} of 1 and the regression loss has a weight λ_{reg} of 1.

Evaluation Metrics. Following the previous methods [32, 50], we measure the *Success* and *Precision* of the tracker. To be specific, *Success* is defined as the IoU between predicted boxes and the ground truth, and *Precision* measures the AUC (Area Under Curve) of the distance between predicted and ground truth boxes within the range of [0, 2] meters.

4.2. Comparison with State-of-the-Art Trackers

Results on KITTI. KITTI [11] is one of the most popular datasets used in mobile robotics and autonomous driving. The tracking benchmark of KITTI consists of 21 training sequences and 29 test sequences. Following the previous

methods [12, 51, 50], we split the training sequences into train/val/test splits due to the inaccessibility of the testing labels, scenes 0-16 for training, scenes 7-18 for validation and scenes 19-20 for testing, respectively.

We compare the SyncTrack with other state-of-the-art methods from the pioneering SC3D [12] to the most recent siamese network STNet [20] and single branch network M2Track [51], as shown in Table 1. We separate the trackers into *Siamese Network* and *Single Branch Network* categories to compare with our proposed SyncTrack. Compared with siamese structured trackers, the SyncTrack achieves the best results on rigid and non-rigid object tracking, outperforming current tracking methods based on siamese networks on most specific categories and the overall mean results. STNet [20] is the state-of-the-art siamese-based tracking method with self-attention and cross-attention modules to match the template and search-region features. Our SyncTrack outperforms the STNet by a relatively large margin on *Car*, *Van* and *Pedestrian* categories under the evaluation of *Success* metric. Also, SyncTrack surpasses the previous best *Mean* results by 2.8% and 1.8% on *Success* and *Precision* metrics, respectively.

We also compare the SyncTrack with the only current single-branch tracker M2Track [51]. Our SyncTrack outperforms M2Track by a large margin, up to 7.8% in *Success* in the category of *Car* whereas the M2Track is better in the category of *Pedestrian*. However, for the comprehensive evaluation, the *Mean* performance of all frames, our SyncTrack outperforms the M2Track by 1.2% on the *Success* metric.

Results on nuScenes. The nuScenes dataset [2] contains 1000 driving scenes collected from Boston and Singapore with a diverse and exciting set of driving maneuvers, traffic situations, and unexpected behaviors. In the configurations of LiDAR-based tracking methods, the train/val/test sets make up 700/150/150 of the whole 1000 scenes, respectively. Officially, the train set is evenly split into ‘train track’ and ‘train detect’ to remedy overfitting. Following [19], we train our model with “train track” split and test it on the val set.

Note that nuScenes dataset only annotates keyframes and provides official interpolated results for the remaining frames, so there are two configurations for this dataset. The first is from [50, 51] which trains and tests both only on the keyframes. The second one is from [19, 20], which trains and tests on all the frames. The results based on these two configurations are different. We believe that the motion in key frames is substantial, which does not conform to the practical applications. Therefore, we train and test all the frames in this paper. We train the previous methods on the nuScenes by ourselves using their official codes to compare with our SyncTrack as many results are missing or only reported by testing nuScenes test-split with a pre-

Table 1. *Success/Precision* comparisons among our SyncTrack and the state-of-the-art methods on the **KITTI** datasets. Mean shows frame-level averaging results. **Bold** and underline denote the best performance and the second-best performance, respectively. Improvements over previous state-of-the-arts are shown in *Italic* and color.

Methods	Car (6424)		Cyclist (308)		Van (1248)		Pedestrian (6088)		Mean (14068)	
	Success	Precision	Success	Precision	Success	Precision	Success	Precision	Success	Precision
<i>Siamese Network</i>										
SC3D [12]	41.3	57.9	41.5	70.4	40.4	47.0	18.2	37.8	31.2	48.5
SC3D-RPN[45]	36.3	51.0	43.0	81.4	-	-	17.9	47.8	-	-
P2B [32]	56.2	72.8	32.1	44.7	40.8	48.4	28.7	49.6	42.4	60.0
MLVSNet [39]	56.0	74.0	34.3	44.5	52.0	61.4	34.1	61.1	45.7	66.6
3DSiamRPN [9]	58.2	76.2	36.1	49.0	45.6	52.8	35.2	56.2	46.6	64.9
LTTR [5]	65.0	77.1	66.2	89.9	35.8	45.6	33.2	56.8	48.7	65.8
PTT [33]	67.8	81.8	37.2	47.3	43.6	52.5	44.9	72.0	55.1	74.2
BAT [50]	60.5	77.7	33.7	45.4	52.4	67.0	42.1	70.1	51.2	72.8
V2B [19]	70.5	81.3	40.8	49.7	50.1	58.0	48.3	73.5	58.4	75.2
PTTR [52]	65.2	77.4	65.1	90.5	52.5	61.8	<u>50.9</u>	81.6	58.4	77.8
STNet [20]	72.1	84.0	73.5	<u>93.7</u>	<u>58.0</u>	70.6	49.9	77.2	61.3	80.1
SyncTrack	73.3	85.0	<u>73.1</u>	93.8	60.3	<u>70.0</u>	54.7	<u>80.5</u>	64.1	81.9
improvement	<i>+1.2</i>	<i>+1.0</i>	<i>-0.4</i>	<i>+0.1</i>	<i>+2.3</i>	<i>-0.6</i>	<i>+3.8</i>	<i>-1.1</i>	<i>+2.8</i>	<i>+1.8</i>
<i>Single Branch Network</i>										
M2Track [51]	<u>65.5</u>	<u>80.8</u>	73.2	<u>93.5</u>	<u>53.8</u>	70.7	61.5	88.2	<u>62.9</u>	83.4
SyncTrack	73.3	85.0	<u>73.1</u>	93.8	60.3	<u>70.0</u>	<u>54.7</u>	<u>80.5</u>	64.1	<u>81.9</u>
improvement	<i>+7.8</i>	<i>+4.2</i>	<i>-0.1</i>	<i>+0.3</i>	<i>+2.3</i>	<i>-0.7</i>	<i>-6.8</i>	<i>-7.7</i>	<i>+1.2</i>	<i>-1.5</i>

Table 2. Comparison among SyncTrack and the state-of-the-art methods on the **nuScenes** datasets. Mean shows the average result weighed by frame numbers. **Bold** and underline denote the best performance and the second-best performance, respectively. Improvements over previous state-of-the-arts are shown in *Italic* and color.

Methods	Car (15578)		Bicycle (501)		Truck (3710)		Pedestrian (8019)		Mean (27808)	
	Success	Precision	Success	Precision	Success	Precision	Success	Precision	Success	Precision
SC3D [12]	24.5	25.9	16.6	18.8	32.5	30.6	13.8	14.7	22.3	23.2
P2B [32]	32.8	35.2	19.7	26.6	16.2	11.1	19.2	26.6	26.4	29.3
BAT [50]	26.5	28.8	17.8	22.8	16.5	10.6	19.4	28.2	23.0	27.9
V2B [19]	32.9	34.5	20.3	27.5	28.7	23.8	<u>20.1</u>	27.4	28.4	30.9
STNet [20]	<u>35.7</u>	<u>37.2</u>	22.3	29.3	<u>33.5</u>	<u>32.4</u>	<u>20.1</u>	27.8	<u>30.7</u>	<u>33.7</u>
M2Track [51]	31.4	33.9	<u>22.6</u>	<u>29.8</u>	30.1	28.8	20.7	<u>28.0</u>	28.0	31.4
SyncTrack	36.7	38.1	23.8	30.4	39.4	38.6	19.1	27.8	31.8	35.1
improvement	<i>+1.0</i>	<i>+0.9</i>	<i>+1.2</i>	<i>+0.6</i>	<i>+5.9</i>	<i>+6.2</i>	<i>-1.6</i>	<i>-0.4</i>	<i>+1.1</i>	<i>+1.4</i>

trained KITTI model.

As shown in Table 2, SyncTrack performs significantly better than other trackers on the mean results of four categories. Specifically, the SyncTrack yields the best results on both metrics and on most categories except the *Pedestrian*, which is lower than M2Track [51] and BAT [50] by a minor margin (1.6% and 0.4%). However, in the *Truck* category, our SyncTrack outperforms state-of-the-art by a large margin, up to 5.9% and 6.2% on *Success* and *Precision* respectively.

Computational Cost Comparison. We analyze the computational overheads and inference speed of SyncTrack and compare it with other trackers in Table 3. The reported results are tested by ourselves with official codebases with a single TITAN RTX GPU on *Car* category of KITTI. It can be seen that SyncTrack achieves the best *Success* performance with the lowest computational complexity (2.51

G). Compared with the most current Siamese-based tracker STNet [20] and single-branch tracker M2Track, our SyncTrack has the fewest number of parameters and fastest inference speed, satisfying the demand of real-time tracking.

Table 3. The computational cost of different trackers.

Methods	Parameters	FLOPs	FPS	Success
SC3D [12]	6.45 M	20.07 G	6	41.3
P2B [32]	1.34 M	4.28 G	48	56.2
BAT [50]	1.47 M	5.53 G	54	60.5
V2B [19]	1.36 M	5.57 G	39	70.5
STNet [20]	1.66 M	3.14 G	36	72.1
M2Track [51]	2.24 M	2.54 G	37	65.5
SyncTrack	1.47 M	2.51 G	45	73.3

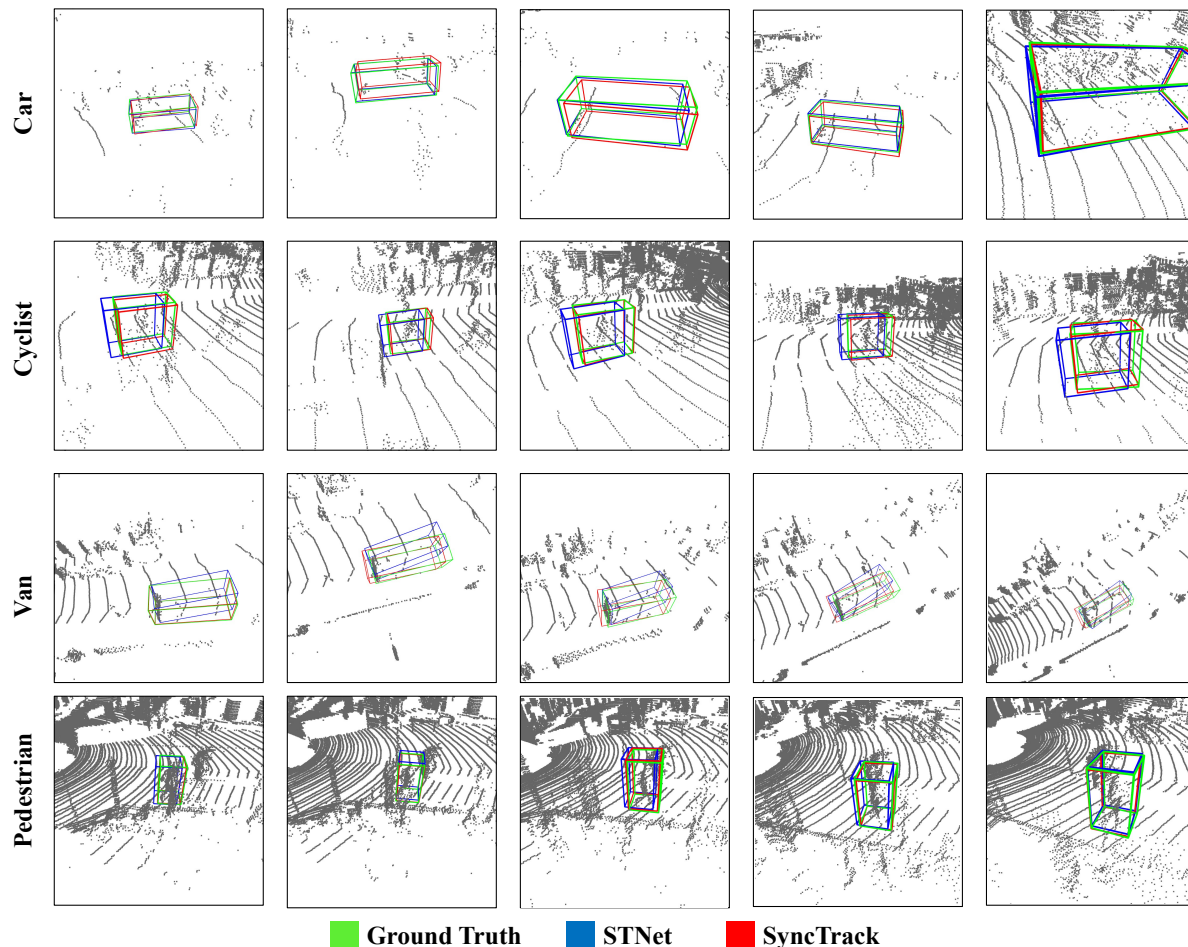


Figure 4. **Qualitative comparison with STNet.** We compare SyncTrack with STNet over four tracking sequences of different categories in KITTI. The predictions of SyncTrack fits the ground-truth boxes better.

4.3. Generalization Ability

To evaluate the generalization ability of SyncTrack, we pre-train the model on the KITTI dataset and test it directly on the nuScenes dataset without fine-tuning. The results are shown in Table 4. It can be observed that SyncTrack outperforms other methods on the mean results of four categories by a large margin. SyncTrack not only achieves a good balance between inference speed and tracking accuracy, but also generalizes very well to new domains.

Table 4. Testing results on the nuScenes for generalization ability.

Methods	Car (15578)	Bicycle (501)	Truck (3710)	Pedestrian (8019)	Mean (27808)
SC3D [12]	25.0/27.1	17.0/18.2	25.7/21.9	14.2/16.2	21.8/23.1
P2B [32]	27.0/29.2	20.0/26.4	21.5/16.2	15.9/22.0	22.9/25.3
BAT [50]	22.5/24.1	17.0/18.8	19.3/15.8	17.3/24.5	20.5/23.0
V2B [19]	31.3/35.1	22.2/19.1	21.7/19.1	17.3/23.4	25.8/29.0
STNet [20]	<u>32.2/36.1</u>	21.2/ 29.2	22.3/16.8	<u>19.1/27.2</u>	<u>26.9/30.8</u>
SyncTrack	32.8/36.3	<u>21.7/28.3</u>	<u>23.9/19.2</u>	19.3/27.1	27.5/31.2

4.4. Scalability of Backbone

We prove that our SyncTrack has good scalability in large-scale datasets like nuScenes. The basic model of SyncTrack has merely one Transformer layer in each stage to ensure real-time performance for tracking. However, the backbone of SyncTrack is scalable in both depth and width. We name the basic model as SyncTrack-Small. Furthermore, the SyncTrack-Mid has three Transformer layers at every stage, with nine layers total. As for the SyncTrack-Large, it doubles the number of feature channels of the SyncTrack-Mid to [256, 128, 64] for every stage. Table 5 reveals that performance improves when scalability increases not only in depth (Small v.s. Mid), but also in width (Mid v.s. Large).

4.5. Ablation Studies

We conduct comprehensive ablations to evaluate the components of SyncTrack.

Compare with Siamese Structure. The synchronized fea-

Table 5. The scalability of SyncTrack on nuScenes, (S, M, L represent small, middle, large model size of SyncTrack).

	Scale	#Param	FLOPs	Success	Precision
Bicycle	S	1.47 M	2.51 G	23.8	30.4
	M	1.82 M	2.63 G	25.0	33.6
	L	3.98 M	5.37 G	25.6	34.1
Truck	S	1.47 M	2.51 G	39.4	38.6
	M	1.82 M	2.63 G	40.1	38.8
	L	3.98 M	5.37 G	40.5	38.9

ture extracting and matching mechanism effectively aggregates features and models the relation. To make comparisons, we split the single branch into a Siamese structure based on SyncTrack. A matching network is added to the Siamese backbone for correlating the features. The shape-aware feature learning network in V2B [19] and iterative coarse-to-fine correlation network in STNet [20] are chosen as matching networks to compare with our single-branch framework. From Table 6, the single-branch structure of SyncTrack is quite significant, and when we split the branch and add a matcher to correlate, the performance drops heavily.

Table 6. Ablations of proposed Single-Branch and Siamese Structure on *Car* category of KITTI and nuScenes, (matcher1 is from the V2B [19] and matcher2 is from the STNet [20]).

Structure	KITTI		nuScenes	
	Success	Precision	Success	Precision
Siamese+matcher1	70.4	82.4	35.3	36.8
Siamese+matcher2	71.8	83.7	35.4	37.1
Single-Branch	73.3	85.0	36.7	38.1

Attentive Sampling. In this paper, we integrate attentive sampling into the multi-head Transformers for selecting search region points to aggregate neighborhood features. We ablate such configuration by performing attentive sampling on template tokens and both template and search region tokens, as well as comparing with standard Transformer without attentive sampling (using random sampling and random/FPS sampling) as shown in Table 7. It can be concluded that the pattern of performing attentive sampling on search region tokens in Transformers is the best. We hypothesize that sampling template tokens attentively is meaningless as template points are target-centric, and search regions’ responses include much noise from the background. Therefore, it is inefficient to down-sample template tokens based on attentive responses.

4.6. Visualization

In Figure 4, we present visualization results obtained from LiDAR video sequences taken from the KITTI dataset. These visualizations show the motion pattern of objects belonging to four categories within the KITTI dataset. Obviously, our SyncTrack excels in accurately tracking the in-

Table 7. Ablations of attentive sampling in APST on *Car* category of KITTI.

	Template	Search Region	Success	Precision
	✗	✗	70.9	82.8
✓APST	✓	✗	67.6	78.8
✗random	✗	✓	73.3	85.0
	✓	✓	69.6	81.1
	✗	✗	71.1	82.8
✓APST	✓	✗	68.0	78.4
✗FPS	✗	✓	73.2	85.0
	✓	✓	70.2	82.7

tended target and predicting bounding boxes when compared to STNet [20]. This achievement is primarily attributed to the dynamic and abundant feature interactions that occur between the template and search region seeds in SyncTrack. These interactions enable our algorithm to effectively distinguish the foreground from the background, leading to superior performance.

5. Conclusion

In this paper, we propose SyncTrack, a novel single-branch and single-stage framework for 3D LiDAR single object tracking. SyncTrack replaces the conventional Siamese-like backbones with a single-branch one, synchronizing the feature extracting and matching without an additional matching network. Moreover, the Attentive Points-Sampling Transformer is proposed for building the backbone, and sampling search region points attentively rather than randomly. Our SyncTrack achieves good tracking performance in accuracy, efficiency, and scalability. We hope it can help motivate further research on more simple yet efficient 3D trackers.

Limitations discussion. Compared with motion-centric tracking framework like M2Track [51], we find our SyncTrack achieves limited performance on tiny-sized and slow-moving objects like pedestrian. We attribute it to the global reasoning mechanism of self-attention. Specifically, the semantic density of small-sized objects’ token is much lower, which hinders effective informative interactions between tokens when performing self-attention. The fact that transformer-based method STNet [20] outperforms CNN-based M2Track on all classes except the pedestrian (Table 1) also supports this hypothesis.

References

- [1] Luca Bertinetto, Jack Valmadre, Joao F Henriques, Andrea Vedaldi, and Philip HS Torr. Fully-convolutional siamese networks for object tracking. In *European conference on computer vision*, pages 850–865. Springer, 2016. 3
- [2] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nusenes: A multi-modal dataset for autonomous driving. In *Proceedings of*

- the *IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020. 6
- [3] Boyu Chen, Peixia Li, Lei Bai, Lei Qiao, Qihong Shen, Bo Li, Weihao Gan, Wei Wu, and Wanli Ouyang. Backbone is all your need: A simplified architecture for visual object tracking. *arXiv preprint arXiv:2203.05328*, 2022. 3
- [4] Xin Chen, Bin Yan, Jiawen Zhu, Dong Wang, Xiaoyun Yang, and Huchuan Lu. Transformer tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8126–8135, 2021. 3
- [5] Yubo Cui, Zheng Fang, Jiayao Shan, Zuoxu Gu, and Sifan Zhou. 3d object tracking with transformer. *arXiv preprint arXiv:2110.14921*, 2021. 3, 7
- [6] Xingping Dong, Jianbing Shen, Dongming Wu, Kan Guo, Xiaogang Jin, and Fatih Porikli. Quadruplet network with one-shot learning for fast visual object tracking. *IEEE Transactions on Image Processing*, 28(7):3516–3527, 2019. 2
- [7] 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. 2
- [8] Heng Fan and Haibin Ling. Siamese cascaded region proposal networks for real-time visual tracking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7952–7961, 2019. 3
- [9] Zheng Fang, Sifan Zhou, Yubo Cui, and Sebastian Scherer. 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. 1, 3, 7
- [10] Peng Gao, Teli Ma, Hongsheng Li, Ziyi Lin, Jifeng Dai, and Yu Qiao. Mcmae: Masked convolution meets masked autoencoders. *Advances in Neural Information Processing Systems*, 35:35632–35644, 2022. 2
- [11] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3354–3361. IEEE, 2012. 2, 6
- [12] Silvio Giancola, Jesus Zarzar, and Bernard Ghanem. Leveraging shape completion for 3d siamese tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1359–1368, 2019. 1, 3, 6, 7, 8
- [13] Dongyan Guo, Jun Wang, Ying Cui, Zhenhua Wang, and Shengyong Chen. Siamcar: Siamese fully convolutional classification and regression for visual tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6269–6277, 2020. 3
- [14] Sam Hare, Stuart Golodetz, Amir Saffari, Vibhav Vineet, Ming-Ming Cheng, Stephen L Hicks, and Philip HS Torr. Struck: Structured output tracking with kernels. *IEEE transactions on pattern analysis and machine intelligence*, 38(10):2096–2109, 2015. 2
- [15] Anfeng He, Chong Luo, Xinmei Tian, and Wenjun Zeng. A twofold siamese network for real-time object tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4834–4843, 2018. 2
- [16] David Held, Sebastian Thrun, and Silvio Savarese. Learning to track at 100 fps with deep regression networks. In *European Conference on Computer Vision*, pages 749–765. Springer, 2016. 2
- [17] João F Henriques, Rui Caseiro, Pedro Martins, and Jorge Batista. High-speed tracking with kernelized correlation filters. *IEEE transactions on pattern analysis and machine intelligence*, 37(3):583–596, 2014. 2
- [18] Lianghua Huang, Xin Zhao, and Kaiqi Huang. Got-10k: A large high-diversity benchmark for generic object tracking in the wild. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019. 2
- [19] Le Hui, Lingpeng Wang, Mingmei Cheng, Jin Xie, and Jian Yang. 3d siamese voxel-to-bev tracker for sparse point clouds. *Advances in Neural Information Processing Systems*, 34, 2021. 1, 2, 3, 6, 7, 8, 9
- [20] Le Hui, Lingpeng Wang, Linghua Tang, Kaihao Lan, Jin Xie, and Jian Yang. 3d siamese transformer network for single object tracking on point clouds. *arXiv preprint arXiv:2207.11995*, 2022. 1, 3, 6, 7, 8, 9
- [21] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [22] Bo Li, Wei Wu, Qiang Wang, Fangyi Zhang, Junliang Xing, and Junjie Yan. Siamrpn++: Evolution of siamese visual tracking with very deep networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4282–4291, 2019. 3
- [23] Bo Li, Junjie Yan, Wei Wu, Zheng Zhu, and Xiaolin Hu. High performance visual tracking with siamese region proposal network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8971–8980, 2018. 2, 3
- [24] Rong Li, Anh-Quan Cao, and Raoul de Charette. Coarse3d: Class-prototypes for contrastive learning in weakly-supervised 3d point cloud segmentation. *arXiv preprint arXiv:2210.01784*, 2022. 2
- [25] Yang Li and Jianke Zhu. A scale adaptive kernel correlation filter tracker with feature integration. In *European conference on computer vision*, pages 254–265. Springer, 2014. 2
- [26] Junwei Liang, Lu Jiang, Kevin Murphy, Ting Yu, and Alexander Hauptmann. The garden of forking paths: Towards multi-future trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10508–10518, 2020. 2
- [27] Junwei Liang, Lu Jiang, Juan Carlos Niebles, Alexander G Hauptmann, and Li Fei-Fei. Peeking into the future: Predicting future person activities and locations in videos. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5725–5734, 2019. 2
- [28] Xiankai Lu, Chao Ma, Jianbing Shen, Xiaokang Yang, Ian Reid, and Ming-Hsuan Yang. Deep object tracking with shrinkage loss. *IEEE transactions on pattern analysis and machine intelligence*, 2020. 2
- [29] Jiageng Mao, Xiaogang Wang, and Hongsheng Li. Interpolated convolutional networks for 3d point cloud understanding. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1578–1587, 2019. 2

- [30] Hyeonseob Nam and Bohyung Han. Learning multi-domain convolutional neural networks for visual tracking. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4293–4302, 2016. 2
- [31] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30, 2017. 2, 4, 5
- [32] Haozhe Qi, Chen Feng, Zhiguo Cao, Feng Zhao, and Yang Xiao. 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*, pages 6329–6338, 2020. 1, 3, 6, 7, 8
- [33] Jiayao Shan, Sifan Zhou, Zheng Fang, and Yubo Cui. Ptt: Point-track-transformer module for 3d single object tracking in point clouds. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1310–1316. IEEE, 2021. 1, 3, 7
- [34] Jack Valmadre, Luca Bertinetto, Joao Henriques, Andrea Vedaldi, and Philip HS Torr. End-to-end representation learning for correlation filter based tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2805–2813, 2017. 2
- [35] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 2
- [36] Mengmeng Wang, Yong Liu, and Zeyi Huang. Large margin object tracking with circulant feature maps. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4021–4029, 2017. 2
- [37] Mengmeng Wang, Teli Ma, Xingxing Zuo, Jiajun Lv, and Yong Liu. Correlation pyramid network for 3d single object tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3215–3224, 2023. 3
- [38] Ning Wang, Wengang Zhou, Jie Wang, and Houqiang Li. Transformer meets tracker: Exploiting temporal context for robust visual tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1571–1580, 2021. 3
- [39] Zhoutao Wang, Qian Xie, Yu-Kun Lai, Jing Wu, Kun Long, and Jun Wang. Mlvsnet: Multi-level voting siamese network for 3d visual tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3101–3110, 2021. 1, 3, 7
- [40] Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. Cbam: Convolutional block attention module. In *Proceedings of the European conference on computer vision (ECCV)*, pages 3–19, 2018. 3
- [41] Fei Xie, Chunyu Wang, Guangting Wang, Yue Cao, Wankou Yang, and Wenjun Zeng. Correlation-aware deep tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8751–8760, 2022. 3
- [42] Yinda Xu, Zeyu Wang, Zuoxin Li, Ye Yuan, and Gang Yu. Siamfc++: Towards robust and accurate visual tracking with target estimation guidelines. In *AAAI*, pages 12549–12556, 2020. 2
- [43] Bin Yan, Houwen Peng, Jianlong Fu, Dong Wang, and Huchuan Lu. Learning spatio-temporal transformer for visual tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10448–10457, 2021. 3
- [44] Botao Ye, Hong Chang, Bingpeng Ma, Shiguang Shan, and Xilin Chen. Joint feature learning and relation modeling for tracking: A one-stream framework. In *European Conference on Computer Vision*, pages 341–357. Springer, 2022. 3
- [45] Jesus Zarzar, Silvio Giancola, and Bernard Ghanem. Efficient bird eye view proposals for 3d siamese tracking. *arXiv preprint arXiv:1903.10168*, 2019. 3, 7
- [46] Zhipeng Zhang and Houwen Peng. Deeper and wider siamese networks for real-time visual tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4591–4600, 2019. 2
- [47] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip HS Torr, and Vladlen Koltun. Point transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16259–16268, 2021. 5
- [48] Haojie Zhao, Gang Yang, Dong Wang, and Huchuan Lu. Deep mutual learning for visual object tracking. *Pattern Recognition*, 112:107796, 2021. 2
- [49] Shaochuan Zhao, Tianyang Xu, Xiao-Jun Wu, and Xue-Feng Zhu. Adaptive feature fusion for visual object tracking. *Pattern Recognition*, 111:107679, 2021. 3
- [50] Chaoda Zheng, Xu Yan, Jiantao Gao, Weibing Zhao, Wei Zhang, Zhen Li, and Shuguang Cui. Box-aware feature enhancement for single object tracking on point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13199–13208, 2021. 1, 3, 6, 7, 8
- [51] Chaoda Zheng, Xu Yan, Haiming Zhang, Baoyuan Wang, Shenghui Cheng, Shuguang Cui, and Zhen Li. 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*, pages 8111–8120, 2022. 1, 2, 3, 6, 7, 9
- [52] Changqing Zhou, Zhipeng Luo, Yueru Luo, Tianrui Liu, Liang Pan, Zhongang Cai, Haiyu Zhao, and Shijian Lu. Pptr: 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. 1, 2, 3, 7
- [53] Jiaming Zhou, Kun-Yu Lin, Haoxin Li, and Wei-Shi Zheng. Graph-based high-order relation modeling for long-term action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8984–8993, 2021. 2
- [54] Zhuangwei Zhuang, Rong Li, Kui Jia, Qicheng Wang, Yuanqing Li, and Mingkui Tan. Perception-aware multi-sensor fusion for 3d lidar semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16280–16290, 2021. 2