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

Beyond 3D Siamese Tracking: A Motion-Centric Paradigm for 3D Single Object Tracking in Point Clouds

Chaoda Zheng¹²³, Xu Yan¹²³, Haiming Zhang¹²³, Baoyuan Wang⁴, Shenghui Cheng⁵, Shuguang Cui¹²³, Zhen Li^{123*} ¹The Chinese University of Hong Kong (Shenzhen) ²The Future Network of Intelligence Institute ³Shenzhen Research Institute of Big Data ⁴Xiaobing.AI ⁵Westlake University {chaodazheng@link., xuyanl@link., haimingzhang@link., lizhen@}cuhk.edu.cn

Abstract

3D single object tracking (3D SOT) in LiDAR point clouds plays a crucial role in autonomous driving. Current approaches all follow the Siamese paradigm based on appearance matching. However, LiDAR point clouds are usually textureless and incomplete, which hinders effective appearance matching. Besides, previous methods greatly overlook the critical motion clues among targets. In this work, beyond 3D Siamese tracking, we introduce a motion-centric paradigm to handle 3D SOT from a new perspective. Following this paradigm, we propose a matching-free two-stage tracker M^2 -Track. At the 1^{st} -stage, M^2 -Track localizes the target within successive frames via **m**otion transformation. Then it refines the target box through **m**otion-assisted shape completion at the 2^{nd} -stage. Extensive experiments confirm that M^2 -Track significantly outperforms previous state-of-the-arts on three large-scale datasets while running at 57FPS (\sim 8%, \sim 17% and \sim 22% precision gains on KITTI, NuScenes, and Waymo Open Dataset respectively). Further analysis verifies each component's effectiveness and shows the motioncentric paradigm's promising potential when combined with appearance matching. Code will be made available at https://github.com/Ghostish/Open3DSOT.

1. Introduction

Single Object Tracking (SOT) is a basic computer vision problem with various applications, such as autonomous driving [22, 38, 39] and surveillance system [32]. Its goal is to keep track of a specific target across a video sequence, given only its initial state (appearance and location).

Existing LiDAR-based SOT methods [8,10,23,26,43,44] all follow the Siamese paradigm, which has been widely



Figure 1. **Top.** Previous Siamese approaches obtain a canonical target template using the previous target box and search for the target in the current frame according to the matching similarity, which is sensitive to distractors. **Bottom**. Our motion-centric paradigm learns the relative target motion from two consecutive frames and then robustly localizes the target in the current frame via motion transformation.

adopted in 2D SOT since it strikes a balance between performance and speed. During the tracking, a Siamese model searches for the target in the candidate region with an **appearance matching** technique, which relies on the features of the target template and the search area extracted by a shared backbone (see Fig.1(a)).

Though the appearance matching for 3D SOT shows satisfactory results on KITTI dataset [9], we observe that KITTI has the following proprieties: i) the target's motion between two consecutive frames is minor, which ensures no drastic appearance change; ii) there are few/no distractors in the surrounding of the target. However, the above characteristics do not hold in natural scenes. Due to self-occlusion, significant appearance changes may occur in consecutive

^{*} Corresponding author

LiDAR views when objects move fast, or the hardware only supports a low frame sampling rate. Besides, the negative samples grow significantly in dense traffic scenes. In these scenarios, it is not easy to locate a target based on its appearance alone (even for human beings).

Is the appearance matching the only solution for LiDAR SOT? Actually, Motion Matters. Since the task deals with a dynamic scene across a video sequence, the target's movements among successive frames are critical for effective tracking. Knowing this, researchers have proposed various 2D Trackers to temporally aggregate information from previous frames [2, 34]. However, the motion information is rarely explicitly modeled since it is hard to be estimated under the perspective distortion. Fortunately, 3D scenes keep intact information about the object motion, which can be easily inferred from the relationships among annotated 3D bounding boxes (BBoxes)¹. Although 3D motion matters for tracking, previous approaches have greatly overlooked it. Due to the Siamese paradigm, previous methods have to transform the target template (initialized by the object point cloud in the first target 3D BBox and updated with the last prediction) from the world coordinate system to its own object coordinate system. This transformation ensures that the shared backbone extracts a canonical target feature, but it adversely breaks the motion connection between consecutive frames.

Based on the above observations, we propose to tackle 3D SOT from a different perspective instead of sticking to the Siamese paradigm. For the first time, we introduce a new motion-centric paradigm that localizes the target in sequential frames without appearance matching by explicitly modeling the target motion between successive frames (Fig. 1(b)). Following this paradigm, we design a novel two-stage tracker M²-Track (Fig. 2). During the tracking, the 1^{st} -stage aims at generating the target BBox by predicting the inter-frame relative target motion. Utilizing all the information from the 1^{st} -stage, the 2^{nd} -stage refines the BBox using a denser target point cloud, which is aggregated from two partial target views using their relative motion. We evaluate our model on KITTI [9], NuScenes [3] and Waymo Open Dataset (WOD) [29], where NuScenes and WOD cover a wide variety of real-world environments and are challenging for their dense traffics. The experiment results demonstrate that our model outperforms the existing methods by a large margin while running faster than the previous top-performer [44]. Besides, the performance gap becomes even more significant when more distractors exist in the scenes. Furthermore, we demonstrate that our method can directly benefit from appearance matching when integrated with existing methods.

In summary, our main contributions are as follows: 1) A

novel motion-centric paradigm for real-time LiDAR SOT, which is free of appearance matching. **2**) A specific secondstage pipeline named M^2 -Track that leverages the motionmodeling and motion-assisted shape completion. **3**) Stateof-the-art online tracking performance with significant improvement on three widely adopted datasets (*i.e.*, KITTI, NuScenes and Waymo Open Dataset).

2. Related Work

Single Object Tracking. A majority of approaches are built for camera systems and take 2D RGB images as input [1, 14, 15, 34, 37, 45]. Although achieving promising results, they face great challenges when dealing with low light conditions or textureless objects. In contrast, LiDARs are insensitive to texture and robust to light variations, making them a suitable complement to cameras. This inspires a new trend of SOT approaches [8, 10, 23, 26, 43, 44] which operate on 3D LiDAR point clouds. These 3D methods all inherit the Siamese paradigm based on appearance matching. As a pioneer, [10] uses the Kalman filter to heuristically sample a bunch of target proposals, which are then compared with the target template based on their feature similarities. The proposal which has the highest similarity with the target template is selected as the tracking result. Since heuristic sampling is time-consuming and inhibits end-to-end training, [23, 43] propose to use a Region Proposal Network (RPN) to generate high-quality target proposals efficiently. Unlike [43] which uses an off-the-shelf 2D RPN operating on bird's eye view (BEV), [23] adapts SiamRPN [15] to 3D point clouds by integrating a pointwise correlation operator with a point-based RPN [19]. The promising improvement brought by [23] inspires a series of follow-up works [8, 13, 26, 44]. They focus on either improving the point-wise correlation operator [44] by feature enhancement, or refining the point-based RPN [8, 13, 26] with more sophisticated structures.

The appearance matching achieves excellent success in 2D SOT because images provide rich texture, which helps the model distinguish the target from its surrounding. However, LiDAR point clouds only contain geometric appearances that lack texture information. Besides, objects in LiDAR sweeps are usually sparse and incomplete. These bring considerable ambiguities which hinder effective appearance matching. Unlike existing 3D approaches, our work no more uses any appearance matching. Instead, we examine a new motion-centric paradigm and show its great potential for 3D SOT.

3D Multi-object Tracking / Detection. In parallel with 3D SOT, 3D multi-object tracking (MOT) focuses on tracking multiple objects simultaneously. Unlike SOT where the user can specify a target of interest, MOT relies on an independent detector [19, 27, 40] to extract potential targets, which obstructs its application for unfamiliar objects

¹This is greatly held for rigid objects (*e.g.*, cars), and it is approximately true for non-rigid objects (*e.g.*, pedestrian).



Figure 2. The overall architecture of M^2 -Track. Given two consecutive point clouds and the possible target BBox at the previous frame, M^2 -Track first segments the target points from their surroundings via joint spatial-temporal learning. At the 1^{st} stage, the model takes in the target points and obtains a coarse BBox at the current frame via motion prediction and transformation. The coarse BBox is further refined at the 2^{nd} stage using motion-assisted shape completion. A detailed illustration with data flows is presented in the supplementary.

(categories unknown by the detector). Current 3D MOT methods predominantly follow the "tracking-by-detection" paradigm, which first detects objects at each frame and then heuristically associates detected BBoxes based on objects' motion or appearance [4, 16, 35, 42]. Recently, [17] proposes to jointly perform detection and tracking by combining object detection and motion association into a unified pipeline. Our motion-centric tracker draws inspiration from the motion-based association in MOT. But unlike MOT, which applies motion estimation on detection results, our approach does not depend on any detector and can leverage the motion prediction to refine the target BBox further.

Spatial-temporal Learning on Point Clouds. Our method utilizes spatial-temporal learning to infer relative motion from multiple frames. Inspired by recent advances in natural language processing [5, 30, 33], there emerges methods that adapt LSTM [12], GRU [41], or Transformer [7] to model point cloud videos. However, their heavy structures make them impractical to be integrated with other downstream tasks, especially for real-time applications. Another trend forms a spatial-temporal (ST) point cloud by merging multiple point clouds with a temporal channel added to each point [11, 22, 24]. Treating the temporal channel as an additional feature (like RGB or reflectance), one can process such an ST point cloud using any 3D backbones [20, 21] without structural modifications. We adopt this strategy to process successive frames for simplicity and efficiency.

3. Methodology

3.1. Problem Statement

Given the initial state of a target, our goal is to localize the target in each frame of an input sequence of a dynamic 3D scene. The frame at timestamp t is a LiDAR point clouds $\mathcal{P}_t \in \mathbb{R}^{N_t \times 3}$ with N_t points and 3 channels, where the point channels encode the xyz global coordinate. The initial state of the target is given as its 3D BBox at the first frame \mathcal{P}_1 . A 3D BBox $\mathcal{B}_t \in \mathbb{R}^7$ is parameterized by its center (xyz coordinate), orientation (heading angle θ around the up-axis), and size (width, length, and height). For the tracking task, we further assume that the size of a target remains unchanged across frames even for non-rigid objects (for a non-rigid object, its BBox size is defined by its maximum extent in the scene). For each frame \mathcal{P}_t , the tracker outputs the amodal 3D BBox of the target with access to only history frames { \mathcal{P}_i }^t_{i=1}.

3.2. Motion-centric Paradigm

Given an input LiDAR sequence and the 3D BBox of the target in the first frame, the motion-centric tracker aims to localize the target frame by frame using explicit motion modeling. At timestamp t (t > 1), the target BBox \mathcal{B}_{t-1} at frame t - 1 is known (either given as the initial state or predicted by the tracker). Having two consecutive frames \mathcal{P}_t and \mathcal{P}_{t-1} as well as the target BBox \mathcal{B}_{t-1} in \mathcal{P}_{t-1} , the tracker predicts the **relative target motion (RTM)** between the successive two frames. We only consider 4DOF instead of 6DOF RTM since the targets are always aligned with the ground plane (no roll and pitch). Specifically, a 4DOF RTM $\mathcal{M}_{t-1,t} \in \mathbb{R}^4$ is defined between two target BBoxes in frame t and t - 1, and contains the translation offset ($\Delta x, \Delta y, \Delta z$) and the yaw offset $\Delta \theta$. We can formulate this process as a function \mathcal{F} :

$$\mathcal{F}: \mathbb{R}^{N_t \times C} \times \mathbb{R}^{N_{t-1} \times C} \times \mathbb{R}^7 \mapsto \mathbb{R}^4, \mathcal{F}(\mathcal{P}_t, \mathcal{P}_{t-1}, \mathcal{B}_{t-1}) \mapsto (\Delta x, \Delta y, \Delta z, \Delta \theta);$$
(1)

Having the predicted RTM $\mathcal{M}_{t-1,t}$, one can easily obtain the target BBox in \mathcal{P}_t using rigid body transformation:

$$\mathcal{B}_t = \operatorname{Transform}(\mathcal{B}_{t-1}, \mathcal{M}_{t-1,t}).$$
(2)



Figure 3. **Stage I.** Taking in the segmented target points $\tilde{\mathcal{P}}_{t-1,t}$ and the target BBox \mathcal{B}_{t-1} at the previous frame, the model outputs a relative target motion state (including an RTM $\mathcal{M}_{t-1,t}$ and 2D binary motion state logits), a refined target BBox $\tilde{\mathcal{B}}_{t-1}$ at the previous frame, and a coarse target BBox \mathcal{B}_t at the current frame.

3.3. *M*²**-Track: Motion-Centric Tracking Pipeline**

Following the motion-centric paradigm, we design a two-stage motion-centric tracking pipeline M^2 -Track (illustrated in Fig.2). M^2 -Track first coarsely localizes the target through target segmentation and **m**otion transformation at the 1^{st} stage, and then refines the BBox at the 2^{nd} stage using **m**otion-assisted shape completion. More details of each module are given as follows.

Target segmentation with spatial-temporal learning

To learn the relative target motion, we first need to segment the target points from their surrounding. Taking as inputs two consecutive frames \mathcal{P}_t and \mathcal{P}_{t-1} together with the target BBox \mathcal{B}_{t-1} , we do this by exploiting the spatial-temporal relation between the two frames (illustrated in the first part of Fig.2). Similar to [18, 24], we construct a spatial-temporal point cloud $\mathcal{P}_{t-1,t} \in \mathbb{R}^{(N_{t-1}+N_t)\times 4} = \{p_i = (x_i, y_i, z_i, t_i)\}_{i=1}^{N_{t-1}+N_t}$ from \mathcal{P}_{t-1} and \mathcal{P}_t by adding a temporal channel to each point and then merging them together. Since there are multiple objects in a scene, we have to specify the target of interest according to \mathcal{B}_{t-1} . To this end, we create a prior-targetness map $\mathcal{S}_{t-1,t} \in \mathbb{R}^{N_{t-1}+N_t}$ to indicate target location in $\mathcal{P}_{t-1,t}$, where $s_i \in \mathcal{S}_{t-1,t}$ is defined as:

$$s_{i} = \begin{cases} 0 & \text{if } p_{i} \text{ is in } \mathcal{P}_{t-1} \text{ and } p_{i} \text{ is not in } \mathcal{B}_{t-1} \\ 1 & \text{if } p_{i} \text{ is in } \mathcal{P}_{t-1} \text{ and } p_{i} \text{ is in } \mathcal{B}_{t-1} \\ 0.5 & \text{if } p_{i} \text{ is in } \mathcal{P}_{t} \end{cases}$$
(3)

Intuitively, one can regard s_i as the prior-confidence of p_i being a target point. For a point in \mathcal{P}_{t-1} , we set its confidence according to its location with respect to \mathcal{B}_{t-1} . Since the target state in \mathcal{P}_t is unknown, we set a median score 0.5 for each point in \mathcal{P}_t . Note that $\mathcal{S}_{t-1,t}$ is not 100% correct for points in \mathcal{P}_{t-1} since \mathcal{B}_{t-1} could be the previous output by the tracker. After that, we form a 5D point cloud by concatenating $\mathcal{P}_{t-1,t}$ and $\mathcal{S}_{t-1,t}$ along the channel axis, and use a PointNet [20] segmentation network to obtain the target mask, which is finally used to extract a spatial-temporal



Figure 4. Stage II. Taking the segmented target points $\tilde{\mathcal{P}}_{t-1,t}$ and the coarse target BBox \mathcal{B}_t as inputs, the model regresses a refined target BBox $\tilde{\mathcal{B}}_t$ on a denser point cloud, which is merged from two partial target point clouds according to their relative motion state.

target point cloud $\widetilde{\mathcal{P}}_{t-1,t} \in \mathbb{R}^{(M_{t-1}+M_t)\times 4}$, where M_{t-1} and M_t are the numbers of target points in frame (t-1) and t respectively.

Stage I: Motion-Centric BBox prediction

As shown in Fig. 3, we encode the spatial-temporal target point clouds $\tilde{\mathcal{P}}_{t-1,t}$ into an embedding using another Point-Net encoder. A multi-layer perceptron (MLP) is applied on top of the embedding to obtain the motion state of the target, which includes a 4D RTM $\mathcal{M}_{t-1,t}$ and 2D binary classification logits indicating whether the target is dynamic. To reduce accumulation errors while performing frame-by-frame tracking, we generate a refined previous target BBox $\tilde{\mathcal{B}}_{t-1}$ by predicting its RTM with respect to \mathcal{B}_{t-1} through another MLP (More details are presented in the supplementary). Finally, we can get the current target BBox \mathcal{B}_t by applying Eqn. 2 on $\mathcal{M}_{t-1,t}$ and $\tilde{\mathcal{B}}_{t-1}$ if the target is classified as dynamic. Otherwise, we simply set \mathcal{B}_t as $\tilde{\mathcal{B}}_{t-1}$.

Stage II: BBox refinement with shape completion

Inspired by two-stage detection networks [27, 28], we improve the quality of the 1st-stage BBox \mathcal{B}_t by additionally regressing a relative offset, which can be regarded as an RTM between \mathcal{B}_t and the refined BBox $\tilde{\mathcal{B}}_t$. Unlike detection networks, we refine the BBox via a novel **motionassisted shape completion** strategy. Due to self-occlusion and sensor movements, LiDAR point clouds suffer from great incompleteness, which hinders precise BBox regression. To mitigate this, we form a denser target point cloud by using the predicted motion state to aggregate the targets from two successive frames. According to the temporal channel, two target point clouds $\tilde{\mathcal{P}}_{t-1} \in \mathbb{R}^{M_t \times 3}$ from different timestamps are extracted

from $\widetilde{\mathcal{P}}_{t-1,t}$. Based on the motion state, we transform $\widetilde{\mathcal{P}}_{t-1}$ to the current timestamp using $\mathcal{M}_{t-1,t}$ if the target is dynamic. The transformed point cloud (identical as $\widetilde{\mathcal{P}}_{t-1}$ if the target is static) is merged with $\widetilde{\mathcal{P}}_t$ to form a denser point cloud $\hat{\mathcal{P}}_t \in \mathbb{R}^{(M_{t-1}+M_t)\times 3}$. Similar to [22, 27], we transform $\hat{\mathcal{P}}_t$ from the world coordinate system to the canonical coordinate system defined by \mathcal{B}_t . We apply a PointNet on the canonical $\hat{\mathcal{P}}_t$ to regress another RTM with respect to \mathcal{B}_t . Finally, the refined target BBox $\widetilde{\mathcal{B}}_t$ is obtained by applying Eqn. 2 on the regressed RTM and \mathcal{B}_t .

3.4. Box-aware Feature Enhancement

As shown in [44], LiDAR SOT directly benefits from the part-aware and size-aware information, which can be depicted by point-to-box relation. To achieve better target segmentation, we construct a distance map $C_{t-1} \in \mathbb{R}^{N_{t-1} \times 9}$ by computing the pairwise Euclidean distance between \mathcal{P}_{t-1} and 9 key points of \mathcal{B}_{t-1} (eight corners and one center arranged in a predefined order with respect to the canonical box coordinate system). After that, we extend C_{t-1} to size $(N_{t-1}+N_t) \times 9$ with zero-padding (for points in \mathcal{P}_t) and additionally concatenate it with $\mathcal{P}_{t-1,t}$ and $\mathcal{S}_{t-1,t}$. The overall box-aware features are then sent to the PointNet segmentation.

3.5. Implementation Details

Loss Functions. The loss function contains classification losses and regression losses, which is defined as L = $\lambda_1 L_{\text{cls}_\text{target}} + \lambda_2 L_{\text{cls}_\text{motion}} + \lambda_3 (L_{\text{reg}_\text{motion}} + L_{\text{reg}_\text{refine}_\text{prev}} +$ $L_{\text{reg-1st}} + L_{\text{reg-2nd}}$). $L_{\text{cls}_\text{target}}$ and $L_{\text{cls}_\text{motion}}$ are standard cross-entropy losses for target segmentation and motion state classification at the 1st-stage (Points are considered as the target if they are inside the target BBoxes; A target is regarded as dynamic if its center moves more than 0.15 meter between two frames). All regression losses are defined as the Huber loss [25] between the predicted and ground-truth RTMs (inferred from ground-truth target BBoxes), where $L_{\text{reg_motion}}$ is for the RTM between targets in the two frames; $L_{\text{reg_refine_prev}}$ is for the RTM between the predicted and the ground-truth BBoxes at timestamp (t-1); $L_{\text{reg_lst}}$ / $L_{\text{reg_lnd}}$ is for the RTM between the 1^{st} / 2^{nd} -stage and ground-truth BBoxes. We empirically set $\lambda_1 = \lambda_2 = 0.1$ and $\lambda_3 = 1$. Input & Motion Augmentation. Since SOT only takes care of one target in a scene, we only need to consider a subregion where the target may appear. For two consecutive frames at (t-1) and t timestamp, we choose the subregion by enlarging the target BBox at (t-1) timestamp by 2 meters. We then sample 1024 points from the subregion respectively at (t-1) and t timestamp to form \mathcal{P}_{t-1} and \mathcal{P}_t . To simulate testing errors during the training, we feed the model a perturbed BBox by adding a slight random

shift to the ground-truth target BBox at (t-1) timestamp.

To encourage the model to learn various motions during the

Table 1. Comparison among our M^2 -Track and the state-of-theart methods on the KITTI datasets. *Mean* shows the average result weighed by frame numbers. **Bold** and <u>underline</u> denote the best performance and the second-best performance, respectively. Improvements over previous state-of-the-arts are shown in *Italic*.

			D 1	1 7	G II I	
	Category	Car	Pedestrian	Van	Cyclist	Mean
	Frame Number	6424	6088	1248	308	14068
	SC3D [10]	41.3	18.2	40.4	41.5	31.2
	SC3D-RPN [43]	36.3	17.9	-	43.2	-
	P2B [23]	56.2	28.7	40.8	32.1	42.4
ess	3DSiamRPN [8]	58.2	35.2	45.6	36.1	46.6
ŏ	LTTR [6]	65.0	33.2	35.8	66.2	48.7
Su	PTT [26]	<u>67.8</u>	44.9	43.6	37.2	55.1
	V2B [13]	70.5	<u>48.3</u>	50.1	40.8	<u>58.4</u>
	BAT [44]	65.4	45.7	<u>52.4</u>	33.7	55.0
	M^2 -Track (Ours)	65.5	61.5	53.8	73.2	62.9
	Improvement	↓5.0	<i>↑13.2</i>	<i>↑1.4</i>	<i>↑7.0</i>	<i>↑</i> 4.5
	SC3D [10]	57.9	37.8	47.0	70.4	48.5
	SC3D-RPN [43]	51.0	47.8	-	81.2	-
_	P2B [23]	72.8	49.6	48.4	44.7	60.0
ior	3DSiamRPN [8]	76.2	56.2	52.8	49.0	64.9
Precis	LTTR [6]	77.1	56.8	45.6	<u>89.9</u>	65.8
	PTT [26]	81.8	72.0	52.5	47.3	74.2
	V2B [13]	<u>81.3</u>	73.5	58.0	49.7	<u>75.2</u>
	BAT [44]	78.9	<u>74.5</u>	<u>67.0</u>	45.4	<u>75.2</u>
	M^2 -Track (Ours)	80.8	88.2	70.7	93.5	83.4
	Improvement	↓1.0	<i>↑13.7</i>	<i>†3.7</i>	<i>↑3.6</i>	<u>↑8.2</u>

training, we randomly flip both targets' points and BBoxes in their horizontal axes and rotate them around their *up*-axes by *Uniform* $[-10^{\circ}, 10^{\circ}]$. We also randomly translate the targets by offsets drawn from *Uniform* [-0.3, 0.3] meter.

4. Experiments

4.1. Experiment Setups

Datasets. We extensively evaluate our approach on three large-scale datasets: KITTI [9], NuScenes [3] and Waymo Open Dataset (WOD) [29]. We follow [10] to adapt these datasets for 3D SOT by extracting the tracklets of annotated tracked instances from each of the scenes. KITTI contains 21 training sequences and 29 test sequences. We follow previous works [10, 23, 44] to split the training set into train/val/test splits due to the inaccessibility of the test labels. NuScenes contains 1000 scenes, which are divided into 700/150/150 scenes for train/val/test. Officially, the train set is further evenly split into "train_track" and "train_detect" to remedy overfitting. Following [44], we train our model with "train_track" split and test it on the val set. WOD includes 1150 scenes with 798 for training, 202 for validation, and 150 for testing. We do training and testing respectively on the training and validation set. Note that NuScenes and WOD are much more challenging than KITTI due to larger data volumes and complexities. The

	Dataset	NuScenes					Waymo Open Dataset			
_	Category Frame Number	Car 64,159	Pedestrian 33,227	Truck 13,587	Trailer <i>3,352</i>	Bus 2,953	Mean 117,278	Vehicle 1,057,651	Pedestrian 510,533	Mean 1,568,184
Success	SC3D [10] P2B [23] BAT [44]	22.31 38.81 <u>40.73</u>	11.29 28.39 <u>28.83</u>	30.67 42.95 <u>45.34</u>	35.28 48.96 <u>52.59</u>	29.35 32.95 <u>35.44</u>	20.70 36.48 <u>38.10</u>	28.32 35.62	15.60 <u>22.05</u>	24.18 <u>31.20</u>
	M ² -Track (Ours) Improvement	55.85 <i>↑15.12</i>	32.10 <i>↑</i> 3.27	57.36 ↑12.02	57.61 ↑5.02	51.39 ↑15.95	49.23 ↑11.14	43.62 ↑8.00	42.10 ↑20.05	43.13 ↑11.92
Precision	SC3D [10] P2B [23] BAT [44]	21.93 43.18 43.29	12.65 52.24 53.32	27.73 41.59 <u>42.58</u>	28.12 40.05 <u>44.89</u>	24.08 27.41 <u>28.01</u>	20.20 45.08 <u>45.71</u>	35.41 44.15	29.56 <u>36.79</u>	- 33.51 <u>41.75</u>
	M ² -Track (Ours) Improvement	65.09 ↑21.80	60.92 ↑7.60	59.54 ↑16.96	58.26 <i>↑13.37</i>	51.44 ↑2 <i>3.43</i>	62.73 ↑17.02	61.64 ↑17.49	67.31 <i>↑30.52</i>	63.48 ↑21.73

Table 2. Comparison of M^2 -Track against state-of-the-arts on the NuScenes and Waymo Open Dataset.



Figure 5. Visualization results. Top: Distractor case in KITTI. Middle: Large motion case in KITTI. Bottom: Case in NuScenes.

LiDAR sequences are sampled at 10Hz for both KITTI and WOD. Though NuScenes samples at 20Hz, it only provides the annotations at 2Hz. Since only annotated keyframes are considered, such a lower frequency for keyframes introduces additional difficulties for NuScenes. Evaluation Metrics. We evaluate the models using the One Pass Evaluation (OPE) [36]. It defines overlap as the Intersection Over Union (IOU) between the predicted and ground-truth BBox, and defines error as the distance between two BBox centers. We report the Success and Precision of each model in the following experiments. Success is the Area Under the

Curve (AUC) with the *overlap* threshold varying from 0 to 1. Precision is the AUC with the error threshold from 0 to 2 meters.

4.2. Comparison with State-of-the-arts

Results on KITTI. We compare M^2 -Track with seven topperformance approaches [6, 8, 10, 13, 23, 26, 43, 44], which have published results on KITTI. As shown in Tab. 1, our method benefits both rigid and non-rigid object tracking, outperforming current approaches under all categories except Car, where PTT [26] and V2B [13] surpass us

Method	Success	Precision		
BAT [44] w/o aug	65.37	78.88		
BAT [44] w/ aug	63.59 ↓ 1.78	76.99 ↓ 1.89		
P2B [23] w/o aug	56.20	72.80		
P2B [23] w/ aug	55.21 ↓ 0.99	71.51 ↓ 1.29		
M^2 -Track w/o aug M^2 -Track w/ aug	$65.29 \\ 65.49 \uparrow 0.20$	77.12 80.81 † 3.69		

Table 3. Influence of Motion Augmentation. "aug" stands for motion augmentation.

The lack of car distractors in the by minor margins. scenes makes our improvement over previous appearancematching-based methods minor for cars. But our improvement for pedestrians is significant (13.2%/13.7% in terms of success/precision) because pedestrian distractors are widespread in the scenes (see the supplementary for more details about distractors). Besides, methods using point-based RPN [8, 23, 26, 44] all perform badly on cyclists, which are relative small in size but usually move fast across time. The second row in Fig. 5 shows the case in which a cyclist moves rapidly across frames. Our method perfectly keeps track of the target while BAT almost fails. To handle such fast-moving objects, [6, 43] leverage BEVbased RPN to generate high-recall proposals from a larger search region. In contrast, we handle this simply by motion modeling without sophisticated architectures.

Results on NuScenes & WOD. We select three representative open-source works: SC3D [10], P2B [23] and BAT [44] as our competitors on NuScenes and WOD. The results on NuScenes except for the Pedestrian class are provided by [44]. We use the published codes of the competitors to obtain other results absent in [44]. SC3D [10] is omitted for WOD comparison due to its costly training time. As shown in Tab. 2, M^2 -Track exceeds all the competitors under all categories, mostly by a large margin. On such two challenging datasets with pervasive distractors and drastic appearance changes, the performance gap between previous approaches and M^2 -Track becomes even larger (e.g., more than 30% precision gain on Waymo Pedestrian). Note that for large objects (*i.e.*, Truck, Trailer, and Bus), even if the predicted centers are far from the target (reflected from lower precision), the output BBoxes of the previous model may still overlap with the ground truth (results in higher success). In contrast, the motion modeling helps to improve not only the success but also the precision by a large margin (e.g., +23.43% gain on Bus) for large objects. Visualization results are provided in Fig. 5 and the supplementary.

4.3. Analysis Experiments

In this section, we extensively analyze M^2 -Track with a series of experiments. Firstly, we compare the behaviors of

Table 4. Integration with Appearance Matching.

Method	Success	Precision
PTT [26]	67.80	81.80
V2B [13]	70.50	81.30
M^2 -Track	65.49	80.81
M^2 -Track + BAT [44]	69.22 ↑ 3.73	81.09 ↑ 0.28
M^2 -Track + P2B [23]	70.21 ↑ 4.72	81.80 ↑ 0.99



Figure 6. Robustness analysis with variant numbers of distractors.

 M^2 -Track and previous appearance-matching-based methods in different setups. Afterward, we equip M^2 -Track with the previous appearance matching approaches to show its potential. Finally, we study the effectiveness of each component in M^2 -Track. All the experiments are conducted on the *Car category of KITTI* unless otherwise stated.

Robustness to Distractors. Though achieving promising improvement on NuScenes and WOD, M^2 -Track brings little improvement on the Car of KITTI. To explain this, we look at the scenes of three datasets and find that the surroundings of most cars in KITTI are free of distractors, which are pervasive in NuScenes and WOD (see the supplementary). Although appearance-matching-based methods are sensitive to distractors, they provide more precise results than our motion-based approach in distractor-free scenarios. But as the number of distractors increases, these methods suffer from noticeable performance degradation due to ambiguities from the distractors. To verify this hypothesis, we randomly add K car instances to each scene of KITTI, and then re-train and evaluate different models using this synthesis dataset. As shown in Fig. 6, M^2 -Track consistently outperforms the other two matching-based methods in scenes with more distractors, and the performance gap grows as K increases. Thanks to the box-awareness, BAT [44] can aid such ambiguities to some extent. But our performance is more stable than BAT's when more distractors are added. Besides, the first row in Fig. 5 shows that, when the number of points decreases due to occlusion, BAT is misled by a distractor and then tracks off course, while

Table 5. Results of M^2 -Track when different modules are ablated. The last row denotes the full model. Bold denotes the largest change.

Box Aware	Prev Box Refinement	Motion Classification	Stage-II	Kitti		NuScenes	
Enhancement				Success	Precision	Success	Precision
	1	1	1	62.00 \ 3.49	76.15 ↓ 4.66	53.68 ↓ 2.17	62.47 \ 2.62
1		1	1	64.23 \ 1.26	78.12 ↓ 2.69	54.70 \ 1.15	61.94 J.15
1	1		1	65.74 10.25	80.29 ↓ 0.52	54.88 ↓ 0.97	64.40 U .69
1	1	1		61.29 4.20	77.31 4 3.50	54.66 1.99	64.15 J 0.94
✓	1	1	1	65.49	80.81	55.85	65.09

 M^2 -Track keeps holding tight to the ground truth. All these observations demonstrate the robustness of our approach.

Influence of Motion Augmentation. We improve the performance of M^2 -Track using the motion augmentation in training, which is not adopted in previous approaches. For a fair comparison, we re-train BAT [44] and P2B [23] using the same configurations in their open-source projects except additionally adding motion augmentation. Tab. 3 shows that motion augmentation instead has an adverse effect on both BAT and P2B. Our model benefits from motion augmentation since it explicitly models target motion and is robust to distractors. In contrast, motion augmentation may move a target closer to its potential distractors and thus harm those appearance-matching-based approaches.

Combine with Appearance Matching. Although our motion-centric model outperforms previous methods from various aspects, appearance-matching-based approaches still show their advantage when dealing with distractorfree scenarios. To combine the advantages of both motionbased and matching-based methods, we apply BAT/P2B as a "re-tracker" to fine-tune the results of M^2 -Track. Specifically, we directly utilize BAT/P2B to search for the target in a small neighborhood of the M^2 -Track's output. Tab. 4 confirms that M^2 -Track can further benefit from appearance matching, even under this naive combination. On KITTI Car, both combined models outperform the topranking PTT [26] by noticeable margins. We believe that one can further boost 3D SOT by combining motion-based and matching-based paradigms with a more delicate design. Ablations. In Tab. 5, we conduct an exhaustive ablation study on both KITTI and NuScenes to understand the components of M^2 -Track. Specifically, we respectively ablate the box-aware feature enhancement, previous BBox refinement, binary motion classification and 2^{nd} stage from M^2 -Track. In general, the effectiveness of the components varies across the datasets, but removing any one of them causes performance degradation. The only exception is the binary motion classification used in the 1^{st} stage, which causes a slight drop on KITTI in terms of success. We suppose this is due to the lack of static objects for KITTI's cars, which results in a biased classifier. Besides, Tab. 5 shows that M^2 -Track keeps performing competitively even with module ablated, especially on NuScenes. This reflects that

the main improvement of M^2 -Track is from the motioncentric paradigm instead of the specific pipeline design.

4.4. Running Overheads

 M^2 -Track achieves exciting performance with just a simple PointNet [20]. Compared with other hierarchical backbones (*e.g.*, [21]) used in previous works, PointNet saves more computational overheads since it does not perform any sampling or grouping operations, which are not only time-consuming but also memory-intensive. Therefore, M^2 -Track runs $1.67 \times$ faster as the previous topperformer BAT [44] (only consider model forwarding time) but saves **31.1%** memory footprint. Using a more advanced backbone (*e.g.*, [21,31]) may further boost the performance but inevitably slows down the running speed. Since we focus on online tracking, we prefer a simpler backbone to balance performance and efficiency.

5. Conclusions

In this work, we revisit 3D SOT in LiDAR point clouds and propose to handle it with a new motion-centric paradigm, which is proven to be an excellent complement to the matching-based Siamese paradigm. In addition to the new paradigm, we propose a specific motion-centric tracking pipeline M^2 -Track, which significantly outperforms the state-of-the-arts from various aspects. Extensive analysis confirms that the motion-centric model is robust to distractors and appearance changes and can directly benefit from previous matching-based trackers. We believe that the motion-centric paradigm can serve as a primary principle to guide future architecture designs.

Acknowledgment

This work was supported in part by NSFC-Youth 61902335, by Key Area R&D Program of Guangdong Province with grant No.2018B030338001, by the National Key R&D Program of China with grant No.2018YFB1800800, by Shenzhen Outstanding Talents Training Fund, by Guangdong Research Project No.2017ZT07X152, by Guangdong Regional Joint Fund-Key Projects 2019B1515120039, by the NSFC 61931024&81922046, by helixon biotechnology company Fund and CCF-Tencent Open Fund.

References

- Goutam Bhat, Martin Danelljan, Luc Van Gool, and Radu Timofte. Learning discriminative model prediction for tracking. In *Int. Conf. Comput. Vis.*, pages 6182–6191, 2019. 2
- [2] Goutam Bhat, Martin Danelljan, Luc Van Gool, and Radu Timofte. Know your surroundings: Exploiting scene information for object tracking. In *Eur. Conf. Comput. Vis.*, pages 205–221. Springer, 2020. 2
- [3] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 11621–11631, 2020. 2, 5
- [4] Hsu-kuang Chiu, Antonio Prioletti, Jie Li, and Jeannette Bohg. Probabilistic 3d multi-object tracking for autonomous driving. arXiv preprint arXiv:2001.05673, 2020. 3
- [5] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv* preprint arXiv:1412.3555, 2014. 3
- [6] Yubo Cui, Zheng Fang, Jiayao Shan, Zuoxu Gu, and Sifan Zhou. 3d object tracking with transformer. arXiv preprint arXiv:2110.14921, 2021. 5, 6, 7
- [7] Hehe Fan, Yi Yang, and Mohan Kankanhalli. Point 4d transformer networks for spatio-temporal modeling in point cloud videos. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 14204–14213, 2021. 3
- [8] 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, 2, 5, 6, 7
- [9] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3354–3361, 2012. 1, 2, 5
- [10] Silvio Giancola, Jesus Zarzar, and Bernard Ghanem. Leveraging shape completion for 3d siamese tracking. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1359–1368, June 2019. 1, 2, 5, 6, 7
- [11] Peiyun Hu, Jason Ziglar, David Held, and Deva Ramanan. What you see is what you get: Exploiting visibility for 3d object detection. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 11001–11009, 2020. 3
- [12] Rui Huang, Wanyue Zhang, Abhijit Kundu, Caroline Pantofaru, David A Ross, Thomas Funkhouser, and Alireza Fathi. An lstm approach to temporal 3d object

detection in lidar point clouds. In *Eur. Conf. Comput. Vis.*, pages 266–282. Springer, 2020. 3

- [13] Le Hui, Lingpeng Wang, Mingmei Cheng, Jin Xie, and Jian Yang. 3d siamese voxel-to-bev tracker for sparse point clouds. *Adv. Neural Inform. Process. Syst.*, 34, 2021. 2, 5, 6, 7
- [14] Bo Li, Wei Wu, Qiang Wang, Fangyi Zhang, Junliang Xing, and Junjie Yan. Siamrpn++: Evolution of siamese visual tracking with very deep networks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4282– 4291, 2019. 2
- [15] Bo Li, Junjie Yan, Wei Wu, Zheng Zhu, and Xiaolin Hu. High performance visual tracking with siamese region proposal network. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 8971–8980, 2018. 2
- [16] Ming Liang, Bin Yang, Wenyuan Zeng, Yun Chen, Rui Hu, Sergio Casas, and Raquel Urtasun. Pnpnet: End-to-end perception and prediction with tracking in the loop. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 11553–11562, 2020. 3
- [17] Chenxu Luo, Xiaodong Yang, and Alan Yuille. Exploring simple 3d multi-object tracking for autonomous driving. In *Int. Conf. Comput. Vis.*, pages 10488–10497, 2021. 3
- [18] AJ Piergiovanni, Vincent Casser, Michael S Ryoo, and Anelia Angelova. 4d-net for learned multi-modal alignment. In *Int. Conf. Comput. Vis.*, pages 15435– 15445, 2021. 4
- [19] Charles R Qi, Or Litany, Kaiming He, and Leonidas J Guibas. Deep hough voting for 3d object detection in point clouds. In *Int. Conf. Comput. Vis.*, pages 9277– 9286, 2019. 2
- [20] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 652–660, 2017. 3, 4, 8
- [21] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In Adv. Neural Inform. Process. Syst., pages 5099–5108, 2017. 3, 8
- [22] Charles R Qi, Yin Zhou, Mahyar Najibi, Pei Sun, Khoa Vo, Boyang Deng, and Dragomir Anguelov. Offboard 3d object detection from point cloud sequences. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 6134–6144, 2021. 1, 3, 5
- [23] Haozhe Qi, Chen Feng, Zhiguo Cao, Feng Zhao, and Yang Xiao. P2b: Point-to-box network for 3d object tracking in point clouds. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 6329–6338, 2020. 1, 2, 5, 6, 7, 8

- [24] Davis Rempe, Tolga Birdal, Yongheng Zhao, Zan Gojcic, Srinath Sridhar, and Leonidas J Guibas. Caspr: Learning canonical spatiotemporal point cloud representations. Adv. Neural Inform. Process. Syst., 2020. 3, 4
- [25] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Adv. Neural Inform. Process. Syst.*, 28:91–99, 2015. 5
- [26] Jiayao Shan, Sifan Zhou, Zheng Fang, and Yubo Cui. Ptt: Point-track-transformer module for 3d single object tracking in point clouds. *arXiv preprint arXiv:2108.06455*, 2021. 1, 2, 5, 6, 7, 8
- [27] Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. Pointrenn: 3d object proposal generation and detection from point cloud. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 770–779, 2019. 2, 4, 5
- [28] Shaoshuai Shi, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. From points to parts: 3d object detection from point cloud with part-aware and part-aggregation network. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2020. 4
- [29] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2446–2454, 2020. 2, 5
- [30] Martin Sundermeyer, Ralf Schlüter, and Hermann Ney. Lstm neural networks for language modeling. In *Thirteenth annual conference of the international speech communication association*, 2012. 3
- [31] Hugues Thomas, Charles R Qi, Jean-Emmanuel Deschaud, Beatriz Marcotegui, François Goulette, and Leonidas J Guibas. Kpconv: Flexible and deformable convolution for point clouds. In *Int. Conf. Comput. Vis.*, pages 6411–6420, 2019. 8
- [32] Simen Thys, Wiebe Van Ranst, and Toon Goedemé. Fooling automated surveillance cameras: adversarial patches to attack person detection. In *IEEE Conf. Comput. Vis. Pattern Recog. Worksh.*, pages 0– 0, 2019. 1
- [33] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In Adv. Neural Inform. Process. Syst., pages 5998– 6008, 2017. 3
- [34] Ning Wang, Wengang Zhou, Jie Wang, and Houqiang Li. Transformer meets tracker: Exploiting temporal context for robust visual tracking. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1571–1580, 2021. 2

- [35] Xinshuo Weng, Jianren Wang, David Held, and Kris Kitani. 3D Multi-Object Tracking: A Baseline and New Evaluation Metrics. *IROS*, 2020. 3
- [36] Yi Wu, Jongwoo Lim, and Ming-Hsuan Yang. Online object tracking: A benchmark. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2411–2418, 2013. 6
- [37] 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
- [38] Xu Yan, Jiantao Gao, Jie Li, Ruimao Zhang, Zhen Li, Rui Huang, and Shuguang Cui. Sparse single sweep lidar point cloud segmentation via learning contextual shape priors from scene completion. In AAAI, volume 35, pages 3101–3109, 2021. 1
- [39] Xu Yan, Chaoda Zheng, Zhen Li, Sheng Wang, and Shuguang Cui. Pointasnl: Robust point clouds processing using nonlocal neural networks with adaptive sampling. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5589–5598, 2020. 1
- [40] Yan Yan, Yuxing Mao, and Bo Li. Second: Sparsely embedded convolutional detection. Sensors, 18(10):3337, 2018. 2
- [41] Junbo Yin, Jianbing Shen, Chenye Guan, Dingfu Zhou, and Ruigang Yang. Lidar-based online 3d video object detection with graph-based message passing and spatiotemporal transformer attention. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 11495– 11504, 2020. 3
- [42] Tianwei Yin, Xingyi Zhou, and Philipp Krähenbühl. Center-based 3d object detection and tracking. *IEEE Conf. Comput. Vis. Pattern Recog.*, 2021. 3
- [43] Jesus Zarzar, Silvio Giancola, and Bernard Ghanem. Efficient bird eye view proposals for 3d siamese tracking. *arXiv preprint arXiv:1903.10168*, 2019. 1, 2, 5, 6, 7
- [44] 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 *Int. Conf. Comput. Vis.*, pages 13199– 13208, 2021. 1, 2, 5, 6, 7, 8
- [45] Zheng Zhu, Qiang Wang, Bo Li, Wei Wu, Junjie Yan, and Weiming Hu. Distractor-aware siamese networks for visual object tracking. In *Eur. Conf. Comput. Vis.*, pages 101–117, 2018. 2