

# CompTrack: Information Bottleneck-Guided Low-Rank Dynamic Token Compression for Point Cloud Tracking

Sifan Zhou<sup>1,2</sup>, Yichao Cao<sup>3</sup>, Jiahao Nie<sup>4</sup>, Yuqian Fu<sup>5</sup>, Ziyu Zhao<sup>1,2</sup>, Xiaobo Lu<sup>1,2,\*</sup>, Shuo Wang<sup>6</sup>

<sup>1</sup>School of Automation, Southeast University, Nanjing, China.

<sup>2</sup>Key Laboratory of Measurement and Control of Complex Systems of Engineering, Ministry of Education, Nanjing, China.

<sup>3</sup>Central South University, Changsha, China.

<sup>4</sup>Zhejiang University of Finance and Economics, Hangzhou, China.

<sup>5</sup>INSAIT, Sofia University "St. Kliment Ohridski", Sofia, Bulgaria.

<sup>6</sup>Mininglamp Technology, Beijing, China.

sifanjay@gmail.com, xblu2013@126.com, wangshuo.e@mininglamp.com

## Abstract

3D single object tracking (SOT) in LiDAR point clouds is a critical task in for computer vision and autonomous driving. Despite great success having been achieved, the inherent sparsity of point clouds introduces a dual-redundancy challenge that limits existing trackers: (1) vast spatial redundancy from background noise impairs accuracy, and (2) informational redundancy within the foreground hinders efficiency. To tackle these issues, we propose **CompTrack**, an novel end-to-end framework that systematically eliminates both forms of redundancy in point clouds. First, CompTrack incorporates a *Spatial Foreground Predictor (SFP)* module to filter out irrelevant background noise based on information entropy, addressing spatial redundancy. Subsequently, its core is an *Information Bottleneck-guided Dynamic Token Compression (IB-DTC)* module that eliminates the informational redundancy within the foreground. Theoretically grounded in low-rank approximation, this module leverages an online SVD analysis to adaptively compress the redundant foreground into a compact and highly informative set of proxy tokens. Extensive experiments on KITTI, nuScenes and Waymo datasets demonstrate that CompTrack achieves top-performing tracking performance with superior efficiency, running at a real-time 90 FPS on a single RTX 3090 GPU.

## Introduction

Single object tracking (SOT) based on LiDAR point clouds is a fundamental task in 3D computer vision, with broad applications in autonomous driving and robotics (Zhou et al. 2025a; Brödermann et al. 2025; Fan et al. 2025). Given the annotated target in the first frame of a point cloud sequence, the goal of this task is to continuously localize the same object in subsequent frames (Zheng et al. 2022; Xu et al. 2024). Despite recent progress, accurately and efficiently track objects in sparse point clouds remains a key challenge, stemming from unaddressed data redundancy.

Early approaches for 3D SOT primarily rely on point-based representations. As a pioneering method, SC3D (Giancola et al. 2019) introduces appearance matching to the

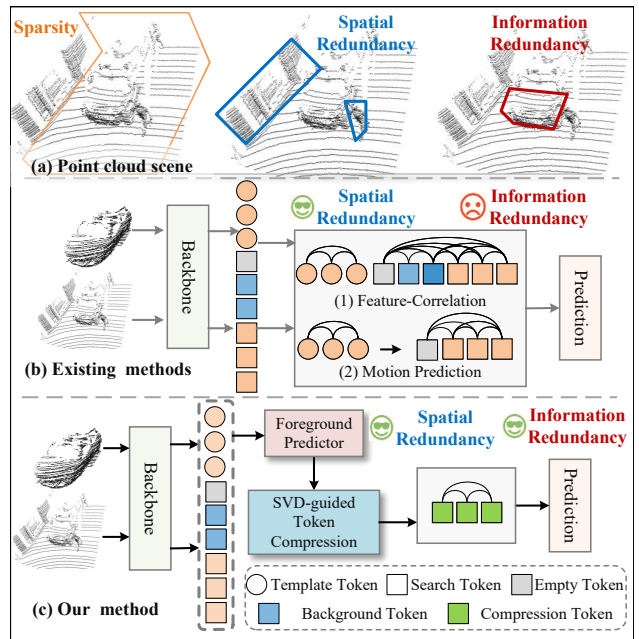


Figure 1: (a) The inherent sparsity of point clouds introduces dual challenges: spatial redundancy from irrelevant background and informational redundancy from repetitive geometries in foreground. (b) Existing SOT methods overlook the information redundancy, which limits their efficiency. (c) Our proposed **CompTrack** framework tackles both spatial and informational redundancy.

3D SOT by identifying the most similar region in subsequent frames based on the given template. Based on this, P2B (Qi et al. 2020) advances appearance matching by integrating a region proposal network (RPN) to generate high-quality candidates. Due to its strong performance, a wide range of follow-up works are proposed, such as PTT (Shan et al. 2021) and MBPTrack (Xu et al. 2023b). Recognizing the limitations of appearance-matching, which can be unreliable with sparse or ambiguous features, motion-centric paradigm (Zheng et al. 2022, 2023; Nie et al. 2025) emerged. These methods instead formulate the tracking as the task of

\*Corresponding author: Xiaobo Lu.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

modeling of target’s motion between consecutive frames.

However, the methods above largely ignore two distinct forms of redundancy that arise from the inherent sparsity of LiDAR data. As shown in Fig. 1 (a), this challenge manifests on two distinct levels. (1) **Spatial redundancy**, caused by the vast number of irrelevant background and empty points that inundate the few features representing the actual target. This creates a severe signal-to-noise problem and leads to computationally wasteful processing. (2) **Informational redundancy**, which stems from the fact that not all points on an object are equally informative within the foreground points. Consider a vehicle (Fig. 1 (a)): points on a large, flat surface like a hood offer ambiguous localization cues, as the local geometry remains unchanged under translation. In contrast, points on corners or edges—where multiple surfaces intersect—provide unique structural information and serve as a reliable indicator of the object location. This is similar to the *aperture problem* in optical flow and has long been pointed out in 2D image recognition (Harris, Stephens et al. 1988). Consequently, the foreground representation is dominated by a large number of these less informative, highly correlated points from simple surfaces. This results in significant informational redundancy and a low-rank feature structure. In turn, existing methods (Fig. 1 (b)) primarily address only spatial redundancy, inherently leaving the challenge of informational redundancy unsolved.

To tackle this dual challenge, we propose **CompTrack**, a novel framework designed to eliminate both the spatial and information redundancy. As show in Fig. 1 (c), CompTrack first employs a *Spatial Foreground Predictor (SPF)* to address spatial redundancy by filtering the vast amount of background noise. Secondly, its core is an *Information Bottleneck-guided Dynamic Token Compression (IB-DTC)* module that resolves informational redundancy. Grounded in the principle of Information Bottleneck and optimal low-rank approximation, this module adaptively compresses the redundant foreground into a compact and highly informative set of proxy tokens for precise and efficient tracking. By systematically removing both forms of redundancy, CompTrack achieves an exceptional balance between performance and latency. Our CompTrack establishes new state-of-the-art performance on the large-scale nuScenes (Caesar et al. 2020) and Waymo (Sun et al. 2020) benchmarks, while achieving highly competitive results on the KITTI (Geiger, Lenz, and Urtasun 2012). Meanwhile, it also runs at a high real-time speed of 90 FPS on a single RTX 3090 GPU, which is 1.4X faster than previous leading methods (Nie et al. 2025). The main contributions can be summarized as follows:

1. We propose **CompTrack**, a novel end-to-end framework that, for the first time, tackles the dual challenges of spatial and information redundancy inherent to sparse point clouds during object tracking.
2. We design a **Spatial Foreground Predictor** that effectively eliminates spatial redundancy. It was trained with a center-guided Gaussian circle, and justified by information entropy analysis showing noise filtering.
3. We introduce a theoretical **Information Bottleneck-**

**guided Dynamic Token Compression** module that is theoretically grounded in the Information Bottleneck principle. It uniquely synergizes an online SVD analysis to dynamically determine the optimal compression ratio with a learnable query mechanism that performs task-specific adaptation, resolving informational redundancy.

4. Extensive experiments on three popular benchmarks demonstrate that CompTrack achieves leading performance, operating at high 90 FPS and demonstrating a superior trade-off between accuracy and speed.

## Related Work

**3D Point Cloud Object Tracking.** Early methods in 3D SOT primarily adopt the appearance matching paradigm, where the goal is to compare features of the template and candidates to identify the most representative instance. As a pioneer, SC3D (Giancola et al. 2019) introduces the first Siamese architecture in 3D tracking. P2B (Qi et al. 2020) and 3D-SiamRPN (Fang et al. 2020) leverage Region Proposal Network (RPN) to generate high-quality 3D proposals in an end-to-end manner. Those inspired a wave of follow-up works (Nie et al. 2023a; Guo et al. 2022; Cui et al. 2021). For instance, BAT (Zheng et al. 2021) augmented correlation learning by encoding the size priors, while methods like PTT (Shan et al. 2021, 2022), and STNet (Hui et al. 2022), GLT-T (Nie et al. 2023b) explored various attention mechanisms to improve feature representation. Recent works such as CXTrack (Xu et al. 2023a) and MBPTrack (Xu et al. 2023b) emphasized the importance of context and memory, respectively. Meanwhile, M<sup>2</sup>Track (Zheng et al. 2022, 2023) series proposes motion-centric frameworks that track objects without relying on appearance similarity. Instead, they predict object future positions based on motion priors, which is particularly beneficial in scenarios where appearance information is ambiguous. Despite their progress, existing appearance- and motion-based trackers neglect the inherent sparsity of point clouds, leading to struggle with background noise and foreground redundancy.

**Efficient Visual Tracking.** The pursuit of efficiency is a significant trend in the broader Tracking domain, driven by the demand for high-performance, low-latency tracking on resource-constrained devices. In 2D SOT, a new wave of efficient algorithms (Blatter et al. 2023; Tan et al. 2025) has emerged to address this. Early approaches such as ECO (Danelljan et al. 2017) and ATOM (Danelljan et al. 2019) demonstrated real-time tracking capabilities on edge devices, but fell short in terms of accuracy compared to recent state-of-the-art methods. Recently, efficient trackers such as LightTrack (Yan et al. 2021), which leverages neural architecture search (NAS), and FEAR (Borsuk et al. 2022), which introduces a dual-template representation and pixel-wise fusion block, have developed between trade-off between speed and accuracy. HiT (Kang et al. 2025) proposes a bridge module to integrate high-level semantics with fine-grained details, utilizing large-stride down-sampling backbones in efficient tracking. LoRAT (Lin et al. 2024) leverages LoRA (Hu et al. 2022) to fine-tune partial parameters without extra latency. While the 2D domain has made significant strides, efficiency has been a less explored frontier in

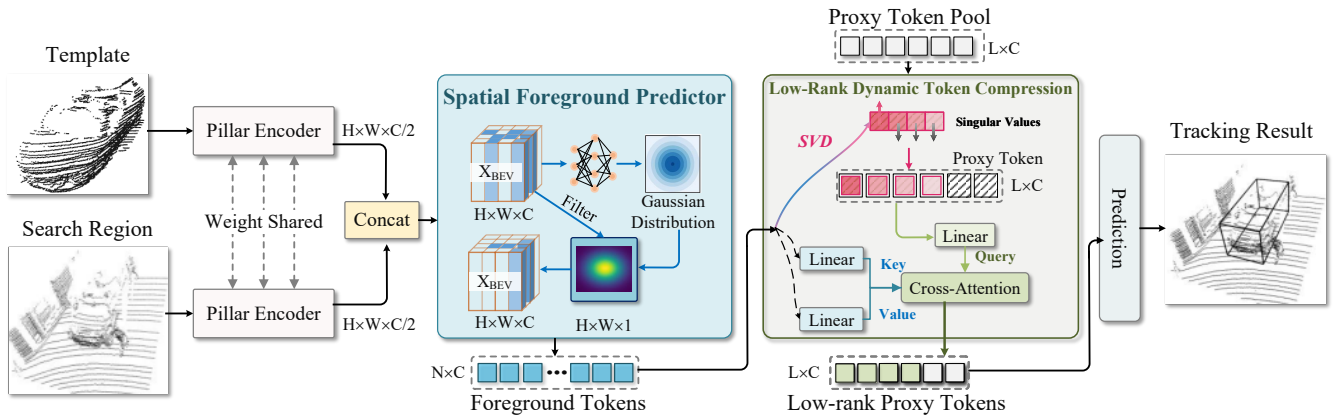


Figure 2: **Overall architecture of our proposed CompTrack.** It consists of two main designs: (1) *Spatial Foreground Predictor (SFP)* that filters irrelevant background to decrease the spatial redundancy, and (2) *Information Bottleneck-guided Low-rank Dynamic Token Compression (IB-DTC)* module that compresses the foreground into a more compact, low-rank representation.

3D SOT. We argue that the inherent sparsity of point clouds presents a unique opportunity to bridge this gap. As a result, we propose the CompTrack, enabling a better trade-off between efficiency and accuracy on 3D tracking.

## Methodology

### CompTrack Architecture

Given a template point cloud  $\mathcal{P}^t \in \mathbb{R}^{N_t \times 3}$  and its corresponding 3D bounding box (BBox)  $B_t = (x_t, y_t, z_t, w_t, h_t, l_t, \theta_t)$  in the first frame, where  $(x_t, y_t, z_t)$  and  $(w_t, h_t, l_t)$  denote the center coordinate and size, and  $\theta_t$  is the rotation angle. LiDAR-based 3D single object tracking (3D SOT) aims to locate the object within the search region  $\mathcal{P}^s \in \mathbb{R}^{N_s \times 3}$  and output a 3D BBox  $B_s$  frame by frame. Note that for the consistent size of the given target in all frames, we can output only 4 parameters to represent  $B_s$ . Our propose **CompTrack** is designed to address this task by tackling the dual challenges of **spatial** and **informational redundancy** inherent to sparse LiDAR data. As depicted in Fig 2, CompTrack is composed of two main stages: (1) *Spatial Foreground Predictor (SFP)* that filters irrelevant background to decrease the spatial redundancy, and (2) *Information Bottleneck-guided Dynamic Token Compression (IB-DTC)* module that compresses the key foreground into a more compact, high-rank representation.

**Pillar Encoder.** We transform the raw point clouds into Bird’s-Eye-View (BEV) feature maps for computational efficiency. Specifically, we follow PillarHist (Zhou et al. 2025b) to process irregular points due to its preservation of fine-grained geometries and efficiency. The input points  $\mathcal{P} \in \mathbb{R}^{N \times 3}$  are scattered to form a 2D pseudo-image as a BEV feature. This approach circumvents the need for computationally expensive 3D convolutions used in voxel-based methods (Lu et al. 2024). The resulting BEV features for the template and search region are denoted as  $F_t$  and  $F_s \in \mathbb{R}^{H \times W \times C}$ . Details can refer to (Zhou et al. 2025b).

### Spatial Foreground Predictor (SFP)

As discussed in Sec. 1 and illustrated in Fig. 3, the inherent sparsity of point clouds leads to significant spatial redun-

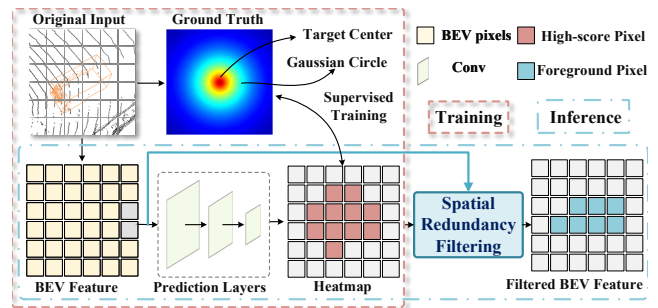


Figure 3: Illustration of the spatial foreground predictor (SFP). SFP removes the spatial redundancy by filtering irrelevant background information.

dancy in the BEV representation. This can be formally understood from an information-theoretic perspective. Given a low occupancy probability  $p \ll 1$  for any BEV pillar, the total entropy (information content) of the BEV map is:

$$H(\mathbf{X}) = HW [H_b(p) + p H_{fg}]. \quad (1)$$

where  $H_b(p)$  is the binary entropy of occupancy and  $H_{fg}$  is the entropy of foreground. For small  $p$ , empty or background pillars contribute negligible information. Consequently, filtering these locations is theoretically **information-lossless** and yields a cleaner feature for tracking.

To practically realize this information-theoretic filtering, we propose the **Spatial Foreground Predictor (SFP)**, a lightweight module designed to generate a spatial attention map that assigns an importance score to each location in the BEV grid. The SFP, denoted as  $\mathbf{F}_{pred}$  with parameters  $\Theta_{pred}$ , is a lightweight CNN with group convolutions (Zhang et al. 2018), which results in a negligible computational overhead. SFP module takes the concatenated BEV features  $\mathbf{X}_{BEV} = \text{cat}(\mathbf{F}_t, \mathbf{F}_s)$  as input and outputs a spatial-wise importance heatmap  $\mathbf{Y}_{pred} \in [0, 1]^{H \times W}$ . The refined search feature map,  $\hat{\mathbf{F}}_s$ , is then obtained via element-wise modulation:

$$\mathbf{Y}_{pred} = \mathbf{F}_{pred}(\mathbf{X}_{BEV}; \Theta_{pred}), \quad (2)$$

$$\hat{\mathbf{F}}_s = \mathbf{F}_s \odot \mathbf{Y}_{pred}. \quad (3)$$




Figure 4: Illustration of the proposed information bottleneck-guided dynamic token compression.

This operation dynamically enhances target-relevant foreground features while suppressing background noise with minimal computational overhead.

**SFP Training.** As discussed in (Zhao et al. 2023), the center of the bounding box is of high importance, and the importance spreads to the local region. Naturally, inspired by CenterPoint (Yin, Zhou, and Krahenbuhl 2021), we supervise the predictor’s training by generating a ground-truth heatmap  $\mathbf{M}_{gt}$ . Instead of a binary mask,  $\mathbf{M}_{gt}$  is formed by rendering a 2D Gaussian circle with a peak centered at each ground-truth box center  $(u, v)_i$ , which can be formulated as:

$$\mathbf{M}_{gt} = \sum_b \mathcal{G}((u, v), \sigma), \quad (4)$$

where  $b$  is the ground-truth bounding box, and  $\mathcal{G}$  is the 2D gaussian function with radius  $\sigma$  adaptively determined by the box size. The predictor is then trained by minimizing the Mean Squared Error (MSE) loss between the predicted heatmap and the ground-truth:  $\mathbf{L}_{pred} = \text{MSE}(\mathbf{Y}_{pred}, \mathbf{M}_{gt})$ .

### Information Bottleneck-guided Dynamic Token Compression (IB-DTC)

Having filtered the *spatial redundancy*, we next tackle the more challenging of *informational redundancy* within the foreground. As discussed in sec , not all foreground points are equally informative for localization. Points on geometrically simple surfaces (e.g., a vehicle’s hood) provide ambiguous cues, whereas points at structural intersections (e.g., corners) are highly informative. This disparity leads to a foreground feature matrix,  $\mathbf{X}_{fg}$ , that is redundant and possesses a *low effective rank*. To this end, we propose the IB-DTC module. Its central goal is to dynamically compress the sparse yet redundant foreground sequence, denoted as  $\mathbf{X}_{fg} \in \mathbb{R}^{N \times C}$ , into a compact, information-dense sequence  $\mathbf{X}_{proxy} \in \mathbb{R}^{K \times C}$  (where  $K \ll N$ ), while preserving the most discriminative characteristics essential for robust tracking. This process is implemented via a learnable, end-to-end differentiable mechanism.

**Information Bottleneck Objective and Low-Rank Approximation.** For the foreground token set  $\mathbf{X}_{fg}$ , our compression objective aligns naturally with the *Information Bottleneck* (IB) principle, which seeks a representation  $\mathbf{X}_{comp}$ :

$$\min_g I(\mathbf{X}_{fg}; \mathbf{X}_{proxy}) \quad \text{s.t.} \quad I(\mathbf{X}_{proxy}; \mathbf{y}) \geq I_0, \quad (5)$$

where  $\mathbf{y}$  is the network output, denotes the target motion state to be predicted (e.g., center offset, orientation). The IB principle directs us to discard all information in  $\mathbf{X}_{fg}$  that is

not predictive of  $\mathbf{y}$ . In other words, we seek to retain only the subset of information in  $\mathbf{X}_{fg}$  that is predictive of the state  $\mathbf{y}$ , and discard irrelevant or redundant components.

Directly optimizing the IB objective is intractable. We therefore employ a practical and powerful surrogate: optimal low-rank approximation. The key insight is that the informational redundancy in  $\mathbf{X}_{fg}$  corresponds to the low-variance components in its feature space. According to the Eckart-Young Theorem (Eckart and Young 1936), the optimal rank- $K$  approximation of  $\mathbf{X}_{fg}$ , denoted as  $\mathbf{X}_{proxy}$  is found by truncating its Singular Value Decomposition (SVD) (Golub and Van Loan 2013) to its top  $K$  components:

$$\mathbf{X}_{proxy} = U \Sigma V^T, \quad (6)$$

where  $U \in \mathbb{R}^{N \times N}$ ,  $\Sigma \in \mathbb{R}^{N \times C}$  (containing the top  $K$  largest singular values), and  $V \in \mathbb{R}^{C \times C}$  are the truncated versions of  $U, \Sigma, V$ , respectively. The theorem guarantees that this  $\mathbf{X}_{proxy}$  minimizes the approximation error. The approximation error is quantified by the sum of the squares of the discarded singular values:

$$\|\mathbf{X}_{fg} - \mathbf{X}_{proxy}\|_F^2 = \sum_{i=K+1}^N \sigma_i^2. \quad (7)$$

Since  $\mathbf{X}_{fg}$  has a low effective rank, its singular values  $\sigma_i$  decay rapidly. Therefore, if we select the fixed compression length  $K$  such that  $K \ll N$ , then the approximation error becomes negligible. This mathematically demonstrates that  $\mathbf{X}_{proxy}$  can capture nearly all the essential information and underlying structure of the original high-dimensional, sparse  $\mathbf{X}_{fg}$ , making the compression process *information-nearly-lossless*. This provides a rigorous theoretical foundation for our learnable compression module.

### SVD-based Dynamic Token Compression (DTC).

While SVD provides an optimal theoretical solution for compression, its non-differentiable nature prevents its direct use in an end-to-end framework. To bridge this gap, we propose SVD-Guided dynamic token compression module, which integrates the principled nature of SVD with the power of a learnable attention mechanism. First, we create a holistic token representation by infusing positional encoding,  $\mathbf{X}'_{fg} = \mathbf{X}_{fg} + \mathbf{PE}_{orig}$ , to ensure subsequent decisions are aware of both semantics and geometry while containing the spatial position information. Our dynamic compression is then realized through the following process:

- 1. Online Rank Estimation:** For each input sequence  $\mathbf{X}'_{fg} \in \mathbb{R}^{N \times C}$ , we perform a fast, non-backpropagated SVD to obtain its singular values  $\sigma_i$ . Guided by the information bottleneck principle, we compute the cumulative energy of singular values and determine the

effective rank  $\mathbf{K}$  as the minimum number of components required to preserve a predefined threshold  $\tau$  of the total variance. Specifically,  $\mathbf{K}$  is derived by identifying the smallest integer such that:

$$\sum_{i=1}^{\mathbf{K}} \sigma_i^2 \geq \tau \sum_{j=1}^{\mathbf{N}} \sigma_j^2, \quad (8)$$

where  $\tau$  is the energy retention threshold (see Appendix). This process dynamically yields two key outputs for each sample: the effective rank  $\mathbf{K}$  and the corresponding optimal basis  $\mathbf{Q}_{\text{SVD}}$  (the first  $\mathbf{K}$  rows of  $\mathbf{V}^T$ ). The rapid energy decay of its singular components concentrates the foreground’s intrinsic information into a small set, enabling a high compression ratio while preserving nearly all essential information. Notably, this SVD step operates on the compact foreground feature and only computes the singular values, making it highly efficient. In practice, its latency is negligible (e.g.,  $<1$  ms on a 3090 GPU).

2. **Dynamic Query Learning based Singular Value:** Our module maintains a fixed pool of  $\mathbf{L}$  learnable *Compression Queries* ( $\mathbf{Q}_{\text{learn}} \in \mathbb{R}^{\mathbf{L} \times \mathbf{C}}$ ). We dynamically select the first  $\mathbf{K}$  of these learnable queries from this pool. This selection of the top- $\mathbf{K}$  queries encourages the model to learn on an ordered basis within the learnable query pool, where the first queries are trained to capture the most common and significant aspects of target objects. The final active queries,  $\mathbf{Q}_{\text{act}} \in \mathbb{R}^{\mathbf{K} \times \mathbf{C}}$ , are then formed by informing this learnable basis with the SVD prior:

$$\mathbf{Q}_{\text{act}} = \mathbf{S}_K \mathbf{Q}_{\text{learn}} + \mathbf{Q}_{\text{SVD}}. \quad (9)$$

where  $\mathbf{S}_K \mathbf{Q}_{\text{learn}}$  select the first  $\mathbf{K}$  rows from the learnable query pool. This formulation allows the learnable queries  $\mathbf{Q}_{\text{learn}}$  to act as task-specific adaptations on top of the strong, data-dependent prior provided by  $\mathbf{Q}_{\text{SVD}}$ .

3. **Guided Cross-Attention:** The final compressed proxy token sequence,  $\mathbf{X}_p \in \mathbb{R}^{\mathbf{K} \times \mathbf{C}}$ , is produced by performing cross-attention between  $\mathbf{K}$  active queries  $\mathbf{Q}_{\text{act}}$  and the input tokens from  $\mathbf{X}'_{fg}$ :

$$\mathbf{X}_p = \text{SoftM} \left( \frac{\mathbf{Q}_{\text{act}} \mathbf{W}_q (\mathbf{X}'_{fg} \mathbf{W}_k)^T}{\sqrt{\mathbf{C}}} \right) \mathbf{X}'_{fg} \mathbf{W}_v. \quad (10)$$

This hybrid approach is end-to-end trainable because the non-differentiable SVD step is only used to determine an integer index for slicing, while the gradient flows through the learnable queries and the cross-attention module. By doing so, our model’s compression ratio is not determined by a learned heuristic but is dynamically guided by the intrinsic rank of the foreground representation. The effectiveness of this SVD-guided formulation is validated in our ablation studies (Tab. 5), showing that combining the SVD prior with learnable queries outperforms using either component alone.

**Training with Adaptive Masking.** In SVD-guided dynamic token compression module, a key challenge is that the number of active tokens,  $\mathbf{K}$ , is dynamically determined for each sample, while efficient batch training requires fixed-size tensors. To resolve this, we adopt an *Adaptive Masking*

strategy where tensor dimensions remain fixed at the maximum length,  $\mathbf{L}$ , during training, but only the dynamically selected tokens contribute to the final loss. The process is as follows: (1) *Adaptive K Determination:* In each forward pass, the effective rank  $\mathbf{K}$  for each sample in the batch is determined via our online SVD estimation. (2) *Mask Generation* A binary mask  $\mathbf{M} \in \{0, 1\}^{\mathbf{L}}$  is generated for each sample, where the first  $\mathbf{K}$  entries are 1 and the rest are 0. (3) *Mask Application* This mask is passed as the subsequent `token_mask` to all subsequent self-attention layers. This logically nullifies the contribution of the “inactive” queries (from  $\mathbf{K} + 1$  to  $\mathbf{L}$ ) by ensuring their attention weights become zero after the softmax operation. This approach allows the gradient to flow only through the  $\mathbf{K}$  active queries that were adaptively selected for that specific sample, while maintaining a consistent tensor shape for batch processing.

**Prediction Head and Training Loss.** Adopting parallel regression branches from (Ma et al. 2023), we directly predict target parameters  $(x, y, z, \theta)$  using the final backbone features, eliminating multi-scale aggregation used in (Nie et al. 2023a). The entire CompTrack network is trained end-to-end with a composite loss function that supervises both the foreground prediction and the final tracking task: The total loss,  $\mathbf{L}_{\text{total}}$ , is a weighted sum of two main components: (1) **Prediction Loss** ( $\mathbf{L}_{\text{pred}}$ ): As described in Eq 4, this loss supervises the Spatial Foreground Predictor (SFP). (2) **Tracking Loss** ( $\mathbf{L}_{\text{track}}$ ): This loss supervises the final prediction head that operates on the compressed proxy tokens. It consists of a classification component to identify the target and a regression component to refine the 3D bounding box. The training loss is adopted from (Hui et al. 2021):  $\mathbf{L}_{\text{track}} = \lambda_1 \mathcal{L}_{(x,y)} + \lambda_2 \mathcal{L}_z + \lambda_3 \mathcal{L}_{\text{rot}}$ , where  $\lambda_1, \lambda_2$  and  $\lambda_3$  are hyper-parameters to balance different losses. Details can refer to (Hui et al. 2021). The final loss is formulated as:

$$\mathbf{L}_{\text{total}} = \theta_1 \mathbf{L}_{\text{pred}} + \theta_2 \mathbf{L}_{\text{track}}, \quad (11)$$

where  $\theta_1$  and  $\theta_2$  are hyper-parameters that balance the two tasks. Note that our SVD-guided compression module does not require a separate sparsity regularizer, as the compression ratio is directly guided by the data’s intrinsic rank.

## Experiments

**Implementation Details.** We follow the common setup (Qi et al. 2020; Zheng et al. 2022) and conduct extensive experiments on KITTI (Geiger, Lenz, and Urtasun 2012), nuScenes (Caesar et al. 2020) and Waymo Open Dataset (WOD) (Sun et al. 2020). The evaluation metrics is followed the common setup (Shan et al. 2021; Nie et al. 2023b; Xu et al. 2023b) to report *Success* and *Precision* based on one pass evaluation (OPE) (Wu, Lim, and Yang 2013; Kristan et al. 2016). More implementation details are in appendix.

### Comparison with State-of-the-art Trackers

**Results on KITTI.** We present comprehensive accuracy and speed comparisons on KITTI (Geiger, Lenz, and Urtasun 2012) dataset. As shown in Tab. 1, CompTrack achieves a highly competitive mean Success/Precision of 71.4% / 89.3%, ranking it among the top performers. Its key advantage lies in exceptional efficiency: while matching the

Tracker	Publish						Computation Efficiency		
		Car [6,424]	Pedestrian [6,088]	Van [1,248]	Cyclist [308]	Mean [14,068]	FLOPs	FPS	Device
SC3D (Giancola et al. 2019)	CVPR'19	41.3 / 57.9	18.2 / 37.8	40.4 / 47.0	41.5 / 70.4	31.2 / 48.5	19.80 G	2	GTx 1080Ti
P2B (Qi et al. 2020)	CVPR'20	56.2 / 72.8	28.7 / 49.6	40.8 / 48.4	32.1 / 44.7	42.4 / 60.0	4.30 G	40	GTx 1080Ti
PTT (Shan et al. 2021)	IROS'21	67.8 / 81.8	44.9 / 72.0	43.6 / 52.5	37.2 / 47.3	55.1 / 74.2	-	40	GTx 1080Ti
BAT (Zheng et al. 2021)	ICCV'21	60.5 / 77.7	42.1 / 70.1	52.4 / 67.0	33.7 / 45.4	51.2 / 72.8	2.77 G	57	RTx 2080
V2B (Hui et al. 2021)	NeurIPS'21	70.5 / 81.3	48.3 / 73.5	50.1 / 58.0	40.8 / 49.7	58.4 / 75.2	5.57 G	37	TITAN RTX
PTTR (Zhou et al. 2022)	CVPR'22	65.2 / 77.4	50.9 / 81.6	52.5 / 61.8	65.1 / 90.5	57.9 / 78.2	2.61 G	50	Tesla V100
M <sup>2</sup> Track (Zheng et al. 2022)	CVPR'22	65.5 / 80.8	61.5 / 88.2	53.8 / 70.7	73.2 / 93.5	62.9 / 83.4	2.54 G	57	Tesla V100
STNet (Hui et al. 2022)	ECCV'22	72.1 / 84.0	49.9 / 77.2	58.0 / 70.6	73.5 / 93.7	61.3 / 80.1	3.14 G	35	TITAN RTX
GLT-T (Nie et al. 2023b)	AAAI'23	68.2 / 82.1	52.4 / 78.8	52.6 / 62.9	68.9 / 92.1	60.1 / 79.3	3.87 G	30	GTx 1080Ti
CXTrack (Xu et al. 2023a)	CVPR'23	69.1 / 81.6	67.0 / 91.5	60.0 / 71.8	74.2 / 94.3	67.5 / 85.3	4.63 G	34	RTx 3090
MBPTrack (Xu et al. 2023b)	ICCV'23	73.4 / 84.8	68.6 / 93.9	61.3 / 72.7	76.7 / 94.3	70.3 / 87.9	2.88 G	50	RTx 3090
SyncTrack (Ma et al. 2023)	ICCV'23	73.3 / 85.0	54.7 / 80.5	60.3 / 70.0	73.1 / 93.8	64.1 / 81.9	2.51 G	45	TITAN RTX
M <sup>2</sup> Track++ (Zheng et al. 2023)	TPAMI'23	71.1 / 82.7	61.8 / 88.7	62.8 / 78.5	75.9 / 94.0	66.5 / 85.2	2.54 G	57	Tesla V100
SCVTrack (Zhang et al. 2024)	AAAI'24	68.7 / 81.9	62.0 / 89.1	58.6 / 72.8	<b>77.4</b> / 94.4	65.2 / 84.6	-	31	RTx 3090
PTTR++ (Luo et al. 2024)	TPAMI'24	73.4 / 84.5	55.2 / 84.7	55.1 / 62.2	71.6 / 92.8	63.9 / 82.8	-	43	Tesla V100
P2P <sup>†</sup> (Nie et al. 2025)	IJCV'25	<b>73.6 / 85.7</b>	<b>69.6 / 94.0</b>	<b>70.3 / 83.9</b>	75.5 / 94.6	<b>71.7 / 89.4</b>	1.23 G	65	RTx 3090
<b>CompTrack (Ours)</b>	-	<u>73.4 / 85.2</u>	<u>69.5 / 94.7</u>	<u>68.5 / 82.5</u>	76.0 / <b>94.8</b>	<u>71.4 / 89.3</u>	<b>0.94 G</b>	<b>90</b>	RTx 3090

Table 1: Comparisons with state-of-the-art methods on KITTI dataset (Geiger, Lenz, and Urtasun 2012). The upper and lower parts include two-stream and one-stream trackers, respectively. *Success / Precision* are used for evaluation. **Bold** and underline denote the best result and the second-best one, respectively. † means our reimplement based on official code.

Tracker	Car [64,159]	Pedestrian [33,227]	Truck [13,587]	Trailer [3,352]	Bus [2,953]	Mean [117,278]
SC3D (Giancola et al. 2019)	22.31 / 21.93	11.29 / 12.65	30.67 / 27.73	35.28 / 28.12	29.35 / 24.08	20.70 / 20.20
P2B (Qi et al. 2020)	38.81 / 43.18	28.39 / 52.24	42.95 / 41.59	48.96 / 40.05	32.95 / 27.41	36.48 / 45.08
PTT (Shan et al. 2021)	41.22 / 45.26	19.33 / 32.03	50.23 / 48.56	51.70 / 46.50	39.40 / 36.70	36.33 / 41.72
BAT (Zheng et al. 2021)	40.73 / 43.29	28.83 / 53.32	45.34 / 42.58	52.59 / 44.89	35.44 / 28.01	38.10 / 45.71
M <sup>2</sup> Track (Zheng et al. 2022)	55.85 / 65.09	32.10 / 60.92	57.36 / 59.54	57.61 / 58.26	51.39 / 51.44	49.23 / 62.73
PTTR (Zhou et al. 2022)	51.89 / 58.61	29.90 / 45.09	45.30 / 44.74	45.87 / 38.36	43.14 / 37.74	44.50 / 52.07
GLT-T (Nie et al. 2023b)	48.52 / 54.29	31.74 / 56.49	52.74 / 51.43	57.60 / 52.01	44.55 / 40.69	44.42 / 54.33
PTTR++ (Luo et al. 2024)	59.96 / 66.73	32.49 / 50.50	59.85 / 61.20	54.51 / 50.28	53.98 / 51.22	51.86 / 60.63
MBPTrack (Xu et al. 2023b)	62.47 / 70.41	45.32 / 74.03	62.18 / 63.31	65.14 / 61.33	55.41 / 51.76	57.48 / 69.88
P2P <sup>†</sup> (Nie et al. 2025)	<u>64.61 / 71.98</u>	<u>45.64 / 74.62</u>	<u>64.42 / 65.37</u>	<u>70.23 / 66.08</u>	<u>58.54 / 56.13</u>	<u>59.22 / 71.19</u>
<b>CompTrack (Ours)</b>	<b>65.70 / 73.50</b>	<b>47.86 / 77.52</b>	<b>68.19 / 69.78</b>	<b>72.89 / 68.11</b>	<b>61.74 / 58.88</b>	<b>61.04 / 73.68</b>

Table 2: Comparisons with state-of-the-art methods on nuScenes dataset (Caesar et al. 2020). *Success / Precision* are used for evaluation. **Bold** and underline denote the best result and the second-best one, respectively.

accuracy of leading methods like P2P, our model operates at a high speed of 90 FPS with significantly lower FLOPs. This speedup is a direct result of our IB-DTC design, which avoids the costly processing of redundant data by distilling the foreground into a compact set of low-rank tokens, proving the design to be both powerful and efficient.

**Results on nuScenes.** On the challenging, sparser and large-scale nuScenes (Caesar et al. 2020) dataset, Tab. 2 shows that our CompTrack sets a new state-of-the-art with a mean Success/Precision of 61.04% / 73.68%, outperforming all prior trackers across all categories. This strong performance in nuScenes’ sparse and complex environments underscores the robustness of our core design. By dynamically filtering and compressing the foreground into a compact, low-rank token representation, CompTrack maintains high accuracy where other methods falter, validating CompTrack’s potential for large-scale practical deployment.

**Results on Waymo.** We follow common setup (Zheng et al. 2022; Hui et al. 2021; Nie et al. 2025) and test the generalization of CompTrack by evaluating our KITTI-trained models directly on the large-scale WOD (Sun et al. 2020) dataset, with results in Tab. 3. CompTrack demonstrates superior generalization, particularly in the Pedestrian category, where

it establishes a new state-of-the-art with a mean Success/Precision of 39.0% / 62.7%. While memory-based methods like MBPTrack show an edge on Vehicle tracking at the cost of speed, our method’s strong performance on pedestrians highlights the robustness of our dynamic token compression. This mechanism learns to distill the most critical features from sparse inputs, ensuring excellent generalization to unseen data and difficult object classes.

### Ablation Study

Due to the limited scale and diversity of the KITTI, to comprehensively evaluate our method, we follow the recent protocol (Hu et al. 2025; Nie et al. 2025) and perform ablations in the large-scale nuScenes (Caesar et al. 2020) dataset. More ablation studies can refer to Appendix.

**Ablation of SFP and IB-DTC Designs.** We conduct ablation studies to validate the performance of each proposed design. As shown in Tab. 4, our baseline model (A), which disables both modules, achieves a mean Success/Precision of 59.38%/71.63% at 48 FPS. By individually integrating the Spatial Foreground Predictor (SFP) in setting (B) or the Information Bottleneck-guided Dynamic Token Compression (IB-DTC) in setting (C), both accuracy and speed are notably improved. This confirms the distinct benefits of filter-

Tracker	Mean	Vehicle			Pedestrian				
		Easy [67,832]	Medium [61,252]	Hard [56,647]	Mean [185,731]	Easy [85,280]	Medium [82,253]	Hard [74,219]	Mean [241,752]
P2B (Qi et al. 2020)	33.0 / 43.8	57.1 / 65.4	52.0 / 60.7	47.9 / 58.5	52.6 / 61.7	18.1 / 30.8	17.8 / 30.0	17.7 / 29.3	17.9 / 30.1
BAT (Zheng et al. 2021)	34.1 / 44.4	61.0 / 68.3	53.3 / 60.9	48.9 / 57.8	54.7 / 62.7	19.3 / 32.6	17.8 / 29.8	17.2 / 28.3	18.2 / 30.3
V2B (Hui et al. 2021)	38.4 / 50.1	64.5 / 71.5	55.1 / 63.2	52.0 / 62.0	57.6 / 65.9	27.9 / 43.9	22.5 / 36.2	20.1 / 33.1	23.7 / 37.9
STNet (Hui et al. 2022)	40.4 / 52.1	65.9 / 72.7	57.5 / 66.0	54.6 / 64.7	59.7 / 68.0	29.2 / 45.3	24.7 / 38.2	22.2 / 35.8	25.5 / 39.9
M <sup>2</sup> Track (Zheng et al. 2022)	44.6 / 58.2	68.1 / 75.3	58.6 / 66.6	55.4 / 64.9	61.1 / 69.3	35.5 / 54.2	30.7 / 48.4	29.3 / 45.9	32.0 / 49.7
CXTrack (Xu et al. 2023a)	42.2 / 56.7	63.9 / 71.1	54.2 / 62.7	52.1 / 63.7	57.1 / 66.1	35.4 / 55.3	29.7 / 47.9	26.3 / 44.4	30.7 / 49.4
MBPTrack (Xu et al. 2023b)	46.0 / 61.0	<b>68.5 / 77.1</b>	<b>58.4 / 68.1</b>	<b>57.6 / 69.7</b>	<b>61.9 / 71.9</b>	37.5 / 57.0	33.0 / 51.9	30.0 / 48.8	33.7 / 52.7
P2P <sup>†</sup> (Nie et al. 2025)	<u>47.2 / 62.9</u>	66.2 / 73.8	57.8 / 67.0	56.8 / 68.1	60.0 / 69.1	<u>43.7 / 65.2</u>	<u>36.4 / 57.1</u>	<u>31.3 / 51.0</u>	<u>37.4 / 58.1</u>
<b>CompTrack (Ours)</b>	<b>48.6 / 65.7</b>	<u>68.2 / 72.8</u>	<u>57.4 / 66.0</u>	<u>57.2 / 69.4</u>	<u>61.2 / 69.6</u>	<b>43.8 / 68.6</b>	<b>38.7 / 62.6</b>	<b>33.8 / 56.0</b>	<b>39.0 / 62.7</b>

Table 3: Comparisons with state-of-the-art methods on Waymo Open Dataset (Sun et al. 2020). *Success / Precision* are used for evaluation. **Bold** and underline denote the best result and the second-best one, respectively.

Setting	SFP	IB-DTC	FPS	Mean	Car	Pedestrian	Truck	Trailer	Bus
(A)	✗	✗	48	59.38 / 71.63	64.73 / 71.90	45.70 / 74.98	64.86 / 65.80	70.67 / 66.56	59.30 / 56.60
(B)	✓	✗	55	60.01 / 72.20	65.01 / 72.25	46.21 / 75.88	67.50 / 67.51	71.05 / 67.03	59.55 / 57.12
(C)	✗	✓	75	59.95 / 72.18	65.01 / 72.25	46.21 / 75.88	65.50 / 66.51	71.05 / 67.03	59.55 / 57.12
(D)	✓	✓	90	<b>61.04 / 73.68</b>	<b>65.70 / 73.50</b>	<b>47.86 / 77.52</b>	<b>68.19 / 69.78</b>	<b>72.89 / 68.11</b>	<b>61.74 / 58.88</b>

Table 4: Ablation of proposed SFP and IB-DTC design. Baseline from the reproduced P2P (Nie et al. 2025) with replacing the voxel-based backbone with self-attention (Liu et al. 2021).

ing spatial redundancy and compressing informational redundancy. Our full model (D), which combines both modules, achieves the best performance and the highest speed (90 FPS), demonstrating that they are complementary.

Query Formulation	FPS	Mean	Car	Pedestrian
(A) Learnable-Only	88	60.70 / 73.25	65.31 / 73.02	47.45 / 76.91
(B) SVD-Only	<b>91</b>	60.15 / 72.50	64.95 / 72.33	46.91 / 75.85
(C) IB-DTC (Ours)	90	<b>61.04 / 73.68</b>	<b>65.70 / 73.50</b>	<b>47.86 / 77.52</b>

Table 5: Ablation study of our SVD-guided query formulation. We compare our full hybrid model against variants that use only learnable queries or only the SVD-derived basis.

**Ablation of SVD-guided Token Generation.** We ablate our dynamic token compression in Table 5. The *Learnable-Only* (A) establishes a strong baseline, demonstrating the capability of end-to-end trained queries. Using the non-learnable *SVD-Only* (B) is also effective, confirming the benefits of SVD priors. However, our (C) *IB-DTC*, which synergizes both, achieves the best performance. This confirms that the SVD provides a meaningful contribution by guiding the learnable queries with a strong prior.

**Feature and Tracking Visualization.** As shown in Fig. 5, we visualize tracking results on the nuScenes (Caesar et al. 2020). It can be seen that whether in dense (Fig. 5 (a-1)), or sparse scenarios (Fig. 5 (a-2)), our CompTrack is able to tightly track the target compared to previous SOTA P2P (Nie et al. 2025), which highlights the CompTrack’s superior performance. Additionally, we plot the feature maps with SPF and IB-DTC designs ((Fig. 5 (b))). The visualization results show the effectiveness of proposed CompTrack.

## Conclusion

In this paper, we propose **CompTrack**, a novel end-to-end framework for 3D SOT that tackles the challenges of spatial and informational redundancy of point clouds.

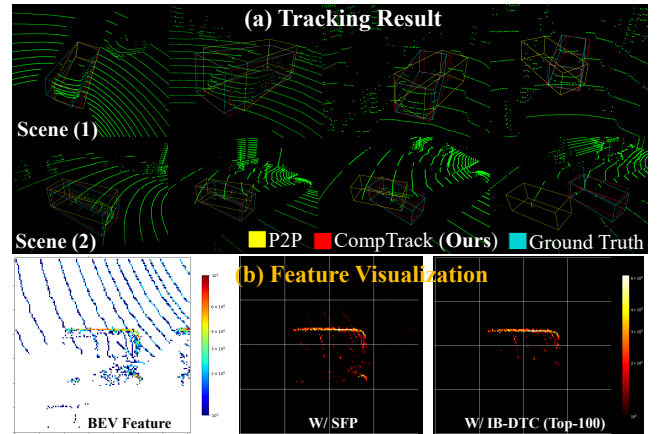


Figure 5: The visualization of tracking and feature maps.

Our CompTrack first employs a Spatial Foreground Predictor to filter background noise, followed by an Information Bottleneck-guided Dynamic Token Compression module that adaptively compresses the foreground into a low-rank and compact representation. By reducing the background and foreground redundancy, our method achieves an exceptional balance between accuracy and efficiency. Extensive experiments demonstrate that CompTrack has leading performance on three well-known benchmarks, striking a better trade-off between efficiency and accuracy in latency-sensitive tracking tasks.

**Limitation.** Although effective, CompTrack still struggles in extremely sparse and occlusion scenes, where the object is partially visible and lacks geometry details. Addressing this requires more information from other modalities (Hu et al. 2025) or fusing the temporal context (Fan et al. 2025). Besides, we believe that model quantization techniques (Zhou et al. 2024; Wang et al. 2025; Xu et al. 2025) are orthogonal with information compression to improve model efficiency.

## Acknowledgements

This work was supported by the National Natural Science Foundation of China (No.62271143), Frontier Technologies R&D Program of Jiangsu (No. BF2024060) and the Big Data Computing Center of Southeast University.

## References

- Blatter, P.; Kanakis, M.; Danelljan, M.; and Van Gool, L. 2023. Efficient visual tracking with exemplar transformers. In *Proceedings of the IEEE/CVF Winter conference on applications of computer vision*, 1571–1581.
- Borsuk, V.; Vei, R.; Kupyn, O.; Martyniuk, T.; Krashenyi, I.; and Matas, J. 2022. FEAR: Fast, efficient, accurate and robust visual tracker. In *European conference on computer vision*, 644–663. Springer.
- Brödermann, T.; Sakaridis, C.; Fu, Y.; and Van Gool, L. 2025. Cafuser: Condition-aware multimodal fusion for robust semantic perception of driving scenes. *IEEE Robotics and Automation Letters*.
- Caesar, H.; Bankiti, V.; Lang, A. H.; Vora, S.; Liong, V. E.; Xu, Q.; Krishnan, A.; Pan, Y.; Baldan, G.; and Beijbom, O. 2020. nuscenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11621–11631.
- Cui, Y.; Fang, Z.; Shan, J.; Gu, Z.; and Zhou, S. 2021. 3d object tracking with transformer. *British Machine Vision Conference*, 1445–1458.
- Danelljan, M.; Bhat, G.; Khan, F. S.; and Felsberg, M. 2019. Atom: Accurate tracking by overlap maximization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4660–4669.
- Danelljan, M.; Bhat, G.; Shahbaz Khan, F.; and Felsberg, M. 2017. Eco: Efficient convolution operators for tracking. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 6638–6646.
- Eckart, C.; and Young, G. 1936. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3): 211–218.
- Fan, B.; Zhou, S.; Li, J.; Zhao, S.; Cao, M.; and Wang, Q. 2025. Beyond Frame-wise Tracking: A Trajectory-based Paradigm for Efficient Point Cloud Tracking. *arXiv preprint arXiv:2509.11453*.
- Fang, Z.; Zhou, S.; Cui, Y.; and Scherer, S. 2020. 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.
- Geiger, A.; Lenz, P.; and Urtasun, R. 2012. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE conference on computer vision and pattern recognition*, 3354–3361. IEEE.
- Giancola, S.; Zarzar, J.; Ghanem, B.; Giancola, S.; and Zarzar, J. 2019. Leveraging shape completion for 3d siamese tracking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1359–1368.
- Golub, G. H.; and Van Loan, C. F. 2013. *Matrix computations*. JHU press.
- Guo, Z.; Mao, Y.; Zhou, W.; Wang, M.; and Li, H. 2022. CMT: Context-Matching-Guided Transformer for 3D Tracking in Point Clouds. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXII*, 95–111. Springer.
- Harris, C.; Stephens, M.; et al. 1988. A combined corner and edge detector. In *Alvey vision conference*, volume 15, 10–5244. Manchester, UK.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2): 3.
- Hu, Z.; Zhou, S.; Yuan, Z.; Yang, D.; Zhao, S.; and Liang, C.-j. 2025. Mvctrack: Boosting 3d point cloud tracking via multimodal-guided virtual cues. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, 3745–3751. IEEE.
- Hui, L.; Wang, L.; Cheng, M.; Xie, J.; and Yang, J. 2021. 3D Siamese voxel-to-BEV tracker for sparse point clouds. *Advances in Neural Information Processing Systems*, 34: 28714–28727.
- Hui, L.; Wang, L.; Tang, L.; Lan, K.; Xie, J.; and Yang, J. 2022. 3d siamese transformer network for single object tracking on point clouds. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part II*, 293–310. Springer.
- Kang, B.; Chen, X.; Zhao, J.; Bo, C.; Wang, D.; and Lu, H. 2025. Exploiting Lightweight Hierarchical ViT and Dynamic Framework for Efficient Visual Tracking. *arXiv preprint arXiv:2506.20381*.
- Kristan, M.; Matas, J.; Leonardis, A.; Vojř, T.; Pflugfelder, R.; Fernandez, G.; Nebel, G.; Porikli, F.; and Čehovin, L. 2016. A novel performance evaluation methodology for single-target trackers. *IEEE transactions on pattern analysis and machine intelligence*, 38(11): 2137–2155.
- Lin, L.; Fan, H.; Zhang, Z.; Wang, Y.; Xu, Y.; and Ling, H. 2024. Tracking meets lora: Faster training, larger model, stronger performance. In *European Conference on Computer Vision*, 300–318. Springer.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*.
- Lu, Y.; Nie, J.; He, Z.; Gu, H.; and Lv, X. 2024. VoxelTrack: Exploring Voxel Representation for 3D Point Cloud Object Tracking. *arXiv preprint arXiv:2408.02263*.
- Luo, Z.; Zhou, C.; Pan, L.; Zhang, G.; Liu, T.; Luo, Y.; Zhao, H.; Liu, Z.; and Lu, S. 2024. Exploring point-bev fusion for 3d point cloud object tracking with transformer. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Ma, T.; Wang, M.; Xiao, J.; Wu, H.; and Liu, Y. 2023. Synchronize Feature Extracting and Matching: A Single Branch Framework for 3D Object Tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9953–9963.
- Nie, J.; He, Z.; Yang, Y.; Bao, Z.; Gao, M.; and Zhang, J. 2023a. OSP2B: One-Stage Point-to-Box Network for 3D Siamese Tracking. In *Proceedings of the Thirty-Second*

- International Joint Conference on Artificial Intelligence*, 1285–1293.
- Nie, J.; He, Z.; Yang, Y.; Gao, M.; and Zhang, J. 2023b. GLT-T: Global-Local Transformer Voting for 3D Single Object Tracking in Point Clouds. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 1957–1965.
- Nie, J.; Xie, F.; Zhou, S.; Zhou, X.; Chae, D.-K.; and He, Z. 2025. P2P: Part-to-Part Motion Cues Guide a Strong Tracking Framework for LiDAR Point Clouds. *International Journal of Computer Vision*, 1–17.
- Qi, H.; Feng, C.; Cao, Z.; Zhao, F.; and Xiao, Y. 2020. 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*, 6329–6338.
- Shan, J.; Zhou, S.; Cui, Y.; and Fang, Z. 2022. Real-time 3D single object tracking with transformer. *IEEE Transactions on Multimedia*, 25: 2339–2353.
- Shan, J.; Zhou, S.; Fang, Z.; and Cui, Y. 2021. 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)*, 1310–1316.
- Sun, P.; Kretschmar, H.; Dotiwala, X.; Chouard, A.; Patnaik, V.; Tsui, P.; Guo, J.; Zhou, Y.; et al. 2020. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2446–2454.
- Tan, Y.; Wu, Z.; Fu, Y.; Zhou, Z.; Sun, G.; Zamfir, E.; Ma, C.; Paudel, D.; Van Gool, L.; and Timofte, R. 2025. Xtrack: Multimodal training boosts rgb-x video object trackers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5734–5744.
- Wang, J.; Wang, Y.; Zhao, S.; and Zhou, S. 2025. Point4Bit: Post Training 4-bit Quantization for Point Cloud 3D Detection. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Wu, Y.; Lim, J.; and Yang, M.-H. 2013. Online object tracking: A benchmark. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2411–2418.
- Xu, T.-X.; Guo, Y.-C.; Lai, Y.-K.; and Zhang, S.-H. 2023a. CXTrack: Improving 3D Point Cloud Tracking With Contextual Information. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 1084–1093.
- Xu, T.-X.; Guo, Y.-C.; Lai, Y.-K.; and Zhang, S.-H. 2023b. MBPTrack: Improving 3D Point Cloud Tracking with Memory Networks and Box Priors. *arXiv preprint arXiv:2303.05071*.
- Xu, W.; Zhou, S.; Xiong, J.; Zhao, Z.; and Yuan, Z. 2024. PillarTrack: Boosting Pillar Representation for Transformer-based 3D Single Object Tracking on Point Clouds. *arXiv preprint arXiv:2404.07495*.
- Xu, Z.; Yue, Y.; Hu, X.; Yang, D.; Yuan, Z.; Jiang, Z.; Chen, Z.; Zhou, S.; et al. 2025. MambaQuant: Quantizing the Mamba Family with Variance Aligned Rotation Methods. *The Thirteenth International Conference on Learning Representations*.
- Yan, B.; Peng, H.; Wu, K.; Wang, D.; Fu, J.; and Lu, H. 2021. Lighttrack: Finding lightweight neural networks for object tracking via one-shot architecture search. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 15180–15189.
- Yin, T.; Zhou, X.; and Krahenbuhl, P. 2021. Center-based 3d object detection and tracking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11784–11793.
- Zhang, J.; Zhou, Z.; Lu, G.; Tian, J.; and Pei, W. 2024. Robust 3D Tracking with Quality-Aware Shape Completion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 7160–7168.
- Zhang, X.; Zhou, X.; Lin, M.; and Sun, J. 2018. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 6848–6856.
- Zhao, T.; Ning, X.; Hong, K.; Qiu, Z.; Lu, P.; Zhao, Y.; Zhang, L.; Zhou, L.; Dai, G.; Yang, H.; et al. 2023. Ada3d: Exploiting the spatial redundancy with adaptive inference for efficient 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 17728–17738.
- Zheng, C.; Yan, X.; Gao, J.; Zhao, W.; Zhang, W.; Li, Z.; and Cui, S. 2021. Box-aware feature enhancement for single object tracking on point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 13199–13208.
- Zheng, C.; Yan, X.; Zhang, H.; Wang, B.; Cheng, S.; Cui, S.; and Li, Z. 2022. 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*, 8111–8120.
- Zheng, C.; Yan, X.; Zhang, H.; Wang, B.; Cheng, S.; Cui, S.; and Li, Z. 2023. An Effective Motion-Centric Paradigm for 3D Single Object Tracking in Point Clouds. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Zhou, C.; Luo, Z.; Luo, Y.; Liu, T.; Pan, L.; Cai, Z.; Zhao, H.; and Lu, S. 2022. Pptr: Relational 3d point cloud object tracking with transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8531–8540.
- Zhou, S.; Li, L.; Zhang, X.; Zhang, B.; Bai, S.; Sun, M.; Zhao, Z.; Lu, X.; and Chu, X. 2024. LiDAR-PTQ: Post-Training Quantization for Point Cloud 3D Object Detection. In *The Twelfth International Conference on Learning Representations (ICLR)*.
- Zhou, S.; Nie, J.; Zhao, Z.; Cao, Y.; and Lu, X. 2025a. FocusTrack: One-Stage Focus-and-Suppress Framework for 3D Point Cloud Object Tracking. In *Proceedings of the 33rd ACM International Conference on Multimedia*, 7366–7375. New York, NY, USA.
- Zhou, S.; Yuan, Z.; Yang, D.; Hu, X.; Qian, J.; and Zhao, Z. 2025b. Pillarhist: A quantization-aware pillar feature encoder based on height-aware histogram. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 27336–27345.