

Delving into Dynamic Scene Cue-Consistency for Robust 3D Multi-Object Tracking

Haonan Zhang^{1,2*}, Xinyao Wang^{1,*}, Boxi Wu^{1,3†}, Tu Zheng^{2,†}, Wang Yunhua⁴, Zheng Yang²

¹Zhejiang University, Hangzhou, China

²Fabu Inc., Hangzhou, China

³Daerwen AI, Hangzhou, China

⁴Shandong Land-Sea-Nexus Digital Technology Co., Ltd., Shandong, China

Abstract

3D multi-object tracking is a critical and challenging task in the field of autonomous driving. A common paradigm relies on modeling individual object motion, e.g., Kalman filters, to predict trajectories. While effective in simple scenarios, this approach often struggles in crowded environments or with inaccurate detections, as it overlooks the rich geometric relationships between objects. This highlights the need to leverage spatial cues. However, existing geometry-aware methods can be susceptible to interference from irrelevant objects, leading to ambiguous features and incorrect associations. To address this, we propose focusing on cue-consistency: identifying and matching stable spatial patterns over time. We introduce the Dynamic Scene Cue-Consistency Tracker (DSC-Track) to implement this principle. Firstly, we design a unified spatiotemporal encoder using Point Pair Features (PPF) to learn discriminative trajectory embeddings while suppressing interference. Secondly, our cue-consistency transformer module explicitly aligns consistent feature representations between historical tracks and current detections. Finally, a dynamic update mechanism preserves salient spatiotemporal information for stable online tracking. Extensive experiments on the nuScenes and Waymo Open Datasets validate the effectiveness and robustness of our approach. On the nuScenes benchmark, for instance, our method achieves state-of-the-art performance, reaching 73.2% and 70.3% AMOTA on the validation and test sets, respectively.

1 Introduction

Accurate and reliable 3D Multi-Object Tracking (MOT) is essential for ensuring the safety and reliability of autonomous driving systems. With the steady improvement of 3D object detectors (Qi et al. 2018; Zhou and Tuzel 2018; Lang et al. 2019; Yin, Zhou, and Krahenbuhl 2021), the tracking-by-detection paradigm (Zhang et al. 2021; Xingshuo Weng 2020; Zeng et al. 2021) has remained a popular choice due to its ability to leverage rich appearance, motion, and geometric information of objects.

One mainstream category of methods focuses on data association using predefined motion models (Yin, Zhou, and

*Equal contribution. This work was done when Haonan was an intern at Fabu.

†Corresponding authors.

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

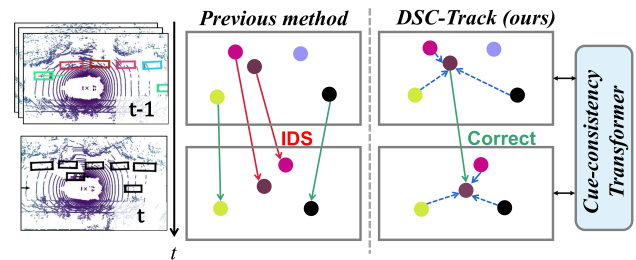


Figure 1: Comparison of a conventional tracking method and our proposed DSC-Track. **Left:** Previous methods, such as those relying on individual object motion, are prone to ID Switches (IDS) when faced with ambiguous associations, like objects in close proximity. **Right:** Our method explicitly models the spatial cue-consistency among all targets using a Transformer, enabling robust and correct associations in challenging scenarios.

Krahenbuhl 2021; Weng et al. 2020) (see Figure 1(Left)). Although these methods have achieved significant progress, their performance remains limited in complex scenarios due to their difficulty in effectively capturing multi-frame temporal information. To address this limitation, another line of work (Kim et al. 2022; Chu et al. 2023) enhances feature representations by fusing geometric information with spatial dependencies of objects. However, this approach introduces new challenges: on one hand, its performance is highly susceptible to inaccuracies from upstream detectors; on the other hand, the indiscriminate utilization of spatial neighborhood nodes leads to temporally inconsistent feature representations, resulting in erroneous associations.

The performance bottleneck of the aforementioned methods lies in constructing robust features for tracklets, which raises two core challenges: 1) How to selectively leverage spatial cues to enhance feature discriminability? 2) How to ensure the consistency and stability of feature representations while incorporating multi-frame temporal information? To tackle these challenges, we propose DSC-Track (see Figure 1(Right)), a method that deeply explores spatial-geometric cues in dynamic scenes to model consistent features between trajectories and detections. Its core idea is to leverage contextual information by searching for and asso-

ciating a set of neighborhood nodes with spatio-temporal consistency for each trajectory in the historical space, rather than relying on prior knowledge (e.g., a fixed search radius). By actively suppressing interference from irrelevant objects and focusing on highly relevant spatial cues, our method significantly enhances the discriminability and stability of feature representations. This leads to more accurate data association and effectively mitigates the challenges posed by dynamic and crowded scenes.

Specifically, DSC-Track consists of the following steps: (1) A memory bank is constructed for each trajectory to store the historical information of the tracked object and the indices of spatially dependent nodes. (2) The unified spatiotemporal encoder aggregates unique feature embeddings of each trajectory from the historical dynamic space through Point Pair Features (PPF) (Drost et al. 2010) as spatial geometric representations, significantly reducing interference from irrelevant nodes. (3) A Transformer-based cue-consistency interaction module leverages spatial cues to extract consistent feature representations between trajectory embeddings and detections. (4) Feature-level affinity matrices constructed from consistency features enable data association, while the spatially dependent nodes in the memory bank are updated based on the latest dependencies of detection nodes mined by the cue-consistency interaction module, achieving robust object tracking in dynamic scenes.

To validate our approach, we conducted extensive experiments on two large-scale benchmarks: the nuScenes (Caesar et al. 2020) and Waymo Open Datasets (Sun et al. 2020). Our ablation studies confirm the effectiveness of each proposed component in improving tracking performance and reducing heuristic dependencies. The key contributions of this work are summarized as follows:

- We present DSC-Track, a novel transformer-based framework that leverages spatiotemporal cues for robust 3D multi-object tracking in dynamic scenes.
- A unified spatio-temporal aggregation module is proposed to capture consistent motion across trajectories from a spatio-temporal perspective, enabling the generation of highly discriminative features.
- A cue-consistent attention module is designed to mine consistent feature pairs from tracks and detections in dense scenes, significantly improving association accuracy.
- Using CenterPoint (Yin, Zhou, and Krahenbuhl 2021) detections as input, our method establishes a new state-of-the-art on the nuScenes benchmark, achieving 73.2% and 70.3% AMOTA on the validation and test sets, respectively, and demonstrates strong, generalizable performance on the Waymo dataset, confirming the robustness of our design.

2 Related Work

Tracking-by-Detection in MOT. The Tracking-by-Detection (TBD) paradigm dominates Multi-Object Tracking (MOT). Early 2D methods evolved from simple motion models like SORT (Bewley et al. 2016) to more advanced techniques incorporating appearance features and

refined association strategies (Wojke, Bewley, and Paulus 2017; Zhang et al. 2021). This philosophy extends to 3D MOT, where methods often pair high-fidelity detectors with similar association techniques (Weng et al. 2020; Zhang et al. 2023). More recent learning-based approaches utilize Graph Neural Networks (GNNs) (Brasó and Leal-Taixé 2020; Zaech et al. 2022) or Transformers (Chu et al. 2023; Ding et al. 2022) to model relationships, yet many still primarily focus on individual object cues, overlooking rich structural information in the spatial context.

Spatial Information in MOT. Leveraging spatial information is crucial for robust tracking. While some methods model spatial relationships using graphs (He et al. 2021; Kim et al. 2022) or Transformers (Chu et al. 2023; Ding et al. 2023), they often face a dilemma: either restricting interactions to local neighborhoods or indiscriminately aggregating global information, which introduces interference from irrelevant objects. In contrast, our method addresses this by adaptively searching for a unique set of spatially consistent dependent nodes for each track. This allows aggregation of highly relevant contextual cues while suppressing interference for more robust association.

3 Approach

3.1 Preliminaries

Problem Statement. Given a set of N 3D bounding box candidates, $\mathbf{B}^t = \{\mathbf{b}_j^t \mid j = 1, \dots, N\}$, detected at the current frame t , and a set of M active tracks, $\mathbf{T}^{t-1} = \{\tau_i^{t-1} \mid i = 1, \dots, M\}$, from the previous frame, our primary goal in 3D Multi-Object Tracking is to establish associations between the active tracks in \mathbf{T}^{t-1} and the new detections in \mathbf{B}^t . These inputs are typically provided by a 3D object detector, such as CenterPoint (Yin, Zhou, and Krahenbuhl 2021).

Detection Node Representation. Each detection candidate \mathbf{b}_i in \mathbf{B}^t is represented by a state vector, following conventions from prior works (Zaech et al. 2022; Ding et al. 2023). Specifically, for the nuScenes dataset, this state is defined as $\mathbf{b}_i = [p_i, \theta_i, s_i, v_i, c_i, o_i] \in R^{17}$. This vector encapsulates the 3D center position $p_i \in R^3$, heading angle $\theta_i \in R$, 3D box size $s_i \in R^3$, velocity $v_i \in R^2$, a one-hot encoded class label $c_i \in R^7$, and the detection confidence score $o_i \in R$.

Trajectory Representation. To maintain temporal context, we represent each active track τ_i using a memory bank. This bank stores two key components: its historical state information, \mathbf{M}_i , and the IDs of its spatially dependent neighboring tracks, \mathcal{K}_i . Formally, a track is defined as $\tau_i = \{\mathbf{M}_i, \mathcal{K}_i\}$. Here, $\mathbf{M}_i \in R^{T_{\max} \times 17}$ is a matrix containing the state vectors of the track over the past T_{\max} time steps. \mathcal{K}_i denotes the set of tracking identifiers for its spatial neighbors. After the association step in each frame, both \mathbf{M}_i and \mathcal{K}_i are updated to reflect the latest state and spatial relationships, with the update mechanism detailed in Section 3.4.

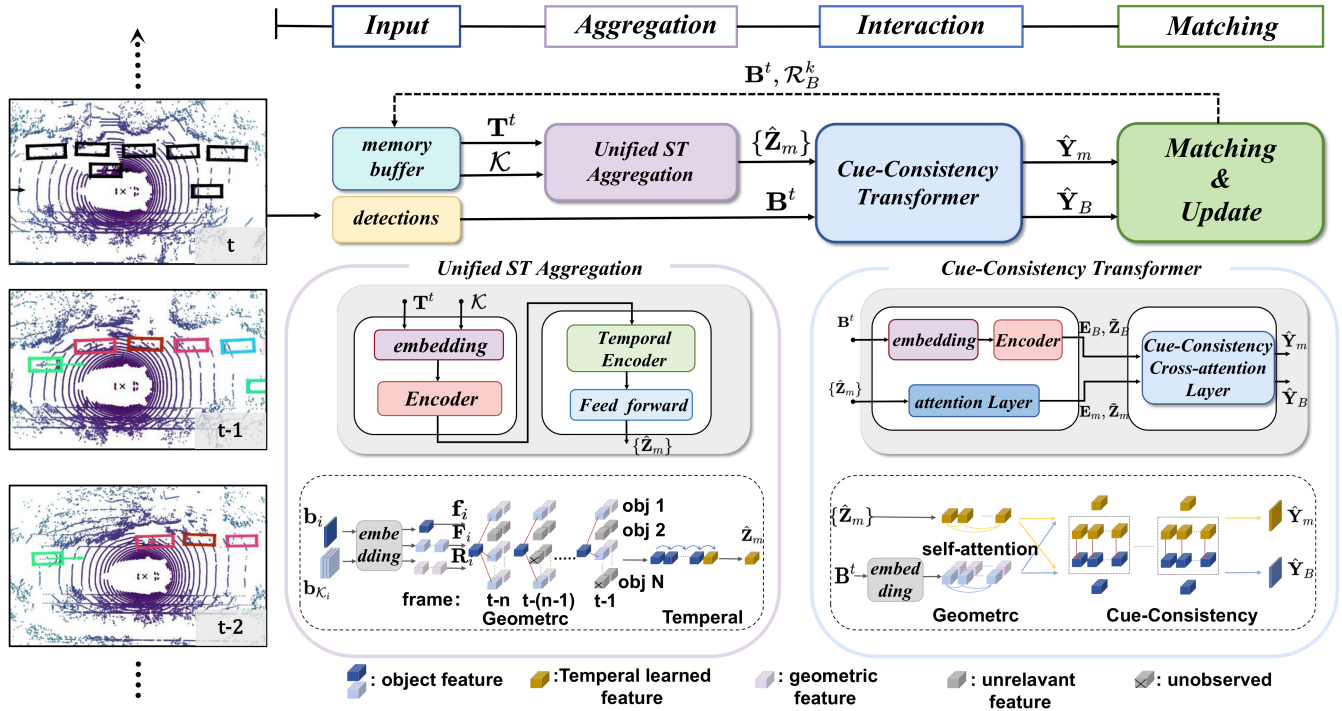


Figure 2: **The overall architecture of our DSC-Track framework.** Our model takes historical track information and new 3D detections as input. (1) The **Unified Spatio-Temporal Aggregation** module first generates a discriminative feature representation, $\{\hat{\mathbf{Z}}_m\}$, for each track by leveraging its historical and spatial context. (2) Then, the **Cue-Consistency Transformer** interacts these track features with detection features (\mathbf{B}^t) to mine consistent cues, yielding enhanced representations for both. (3) Finally, in the **Matching and Update** stage, an affinity matrix is computed from these enhanced features for data association, and the memory buffer is updated for the next frame.

3.2 Unified Spatiotemporal Aggregation

Overview. To generate highly discriminative features for each of the M active tracks in \mathbf{T}^{t-1} , we propose a unified spatio-temporal aggregation module. The core idea is to leverage spatially consistent geometric relationships across frames to produce robust track representations, denoted as $\{\hat{\mathbf{Z}}_m\} \in R^{M \times d}$. As shown in Figure 2, the architecture of this module consists of a Geometric Encoder followed by a Temporal Encoder. This two-stage process is iterated L times to progressively refine the features.

Initial Feature Embedding. Modeling geometric relationships from absolute object states is unreliable, as these features lack rotation-invariance and become unstable when objects maneuver (e.g., a vehicle turning). To address this, we construct a local geometric embedding using the Point Pair Feature (PPF) (Drost et al. 2010). By capturing the relative geometry—defined by the distance and inter-object angles (see Figure 3(Left))—PPF provides a robust representation that remains consistent across rotations and viewpoint changes, ensuring stable geometric modeling.

For a given reference object \mathbf{b}_i and its set of k spatially-dependent neighbors $\{\mathbf{b}_j\}_{j \in \mathcal{K}_i}$, we first extract their positions p and heading angles θ . Each heading angle is converted to a 2D direction vector \mathbf{n} from the Bird’s-Eye View (BEV). The PPF encodes the relationship between the refer-

ence object \mathbf{b}_i and each neighbor \mathbf{b}_j as a 4D vector:

$$\mathbf{e}_{i,j} = (\|\mathbf{d}\|_2, \angle(\mathbf{n}_i, \mathbf{d}), \angle(\mathbf{n}_j, \mathbf{d}), \angle(\mathbf{n}_j, \mathbf{n}_i)) \in R^4, \quad (1)$$

where $\mathbf{d} = p_j - p_i$, and $\angle(\cdot, \cdot)$ computes the angle between two vectors. By concatenating these features for all k neighbors, we form a local geometric matrix $\mathbf{E}_i \in R^{k \times 4}$.

This geometric matrix is then processed to create our final feature triplet. First, \mathbf{E}_i is enhanced with sinusoidal positional encodings (Vaswani et al. 2017) and projected by an MLP to form the relative geometry embedding $\mathbf{R}_i \in R^{k \times d}$. Concurrently, the state vectors of the reference object \mathbf{b}_i and its neighbors $\{\mathbf{b}_j\}_{j \in \mathcal{K}_i}$ are projected by a shared MLP into contextual embeddings: the reference feature $\mathbf{f}_i \in R^{1 \times d}$ and neighbor features $\mathbf{F}_i \in R^{k \times d}$.

These three components form the geometric triplet $\mathcal{G}_i = (\mathbf{f}_i, \mathbf{F}_i, \mathbf{R}_i)$, which serves as the rich, geometrically-aware input for the subsequent encoder.

Geometric Encoder. The geometric encoder processes the geometric triplets frame-by-frame for each track. For a given triplet \mathcal{G}_i^t at frame t , the encoder aggregates the contextual features of neighbors (\mathbf{F}_i) onto the reference feature (\mathbf{f}_i), guided by their rich geometric relationships (\mathbf{R}_i). This produces an updated, geometrically-aware feature $\mathbf{z}_i^t \in R^d$. For simplicity, we omit the frame index t in the following description.

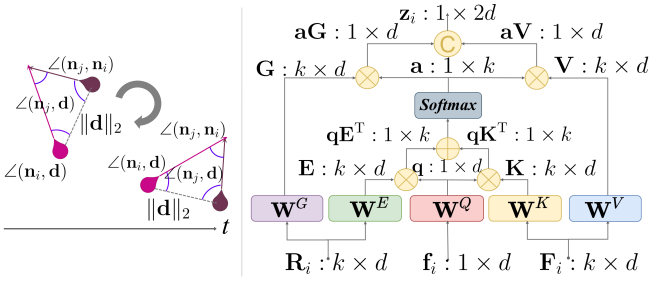


Figure 3: **Left:** Illustration of the rotation-invariance of our Point Pair Feature (PPF). **Right:** Details of the Geometric Inject Attention (GIA) module used for aggregation.

The core of this process is our **Geometric Inject Attention (GIA)** module, inspired by (Yu et al. 2023) and depicted in Figure 3 (right). The GIA module enhances standard attention by injecting geometric information into both the attention score calculation and the value aggregation. The process involves several key steps. **First**, we project the input triplet into components for the attention mechanism. The contextual embeddings are projected into a standard query, key, and value:

$$\mathbf{q} = \mathbf{f}_i \mathbf{W}^q, \quad \mathbf{K} = \mathbf{F}_i \mathbf{W}^K, \quad \mathbf{V} = \mathbf{F}_i \mathbf{W}^V, \quad (2)$$

where $\mathbf{W}^q, \mathbf{W}^K, \mathbf{W}^V \in R^{d \times d}$ are learnable matrices. Concurrently, the relative geometry embedding \mathbf{R}_i is projected into a geometric bias for attention scores, and geometric features for value aggregation:

$$\mathbf{E} = \mathbf{R}_i \mathbf{W}^E, \quad \mathbf{G} = \mathbf{R}_i \mathbf{W}^G. \quad (3)$$

Second, the attention scores are computed by injecting the geometric bias \mathbf{E} into the standard dot-product attention:

$$\mathbf{a} = \text{Softmax} \left(\frac{\mathbf{q} \mathbf{K}^T + \mathbf{q} \mathbf{E}^T}{\sqrt{d}} \right). \quad (4)$$

Finally, the output \mathbf{z}_i is obtained by using the attention scores \mathbf{a} to aggregate both the contextual values \mathbf{V} and the geometric features \mathbf{G} . This is followed by a standard residual connection and a Feed-Forward Network (FFN):

$$\mathbf{z}'_i = \mathbf{f}_i + \text{MLP}_{\text{GIA}}(\text{concat}(\mathbf{a} \mathbf{V}, \mathbf{a} \mathbf{G})), \quad (5)$$

$$\mathbf{z}_i = \text{FFN}(\text{LN}(\mathbf{z}'_i)), \quad (6)$$

where $\text{LN}(\cdot)$ is Layer Normalization (Ba, Kiros, and Hinton 2016) and FFN consists of a two-layer MLP with a residual connection.

Analysis. By explicitly fusing geometric embeddings with contextual information, our method enhances feature representations. At the **intra-frame** level, each trajectory leverages unique spatial dependencies to obtain a highly distinctive feature. This, in turn, enhances robustness at the **inter-frame** level: by producing stable and geometrically-aware features for each frame, this encoder provides a superior input to the subsequent temporal encoder, which can then more effectively suppress noise and ensure stable spatial relationships over time.

Temporal Encoder. After the Geometric Encoder, we obtain a sequence of frame-wise feature representations for each track, denoted as $\mathbf{Z}_i = [\mathbf{z}_i^{t-T_{\max}+1}, \dots, \mathbf{z}_i^t] \in R^{T_{\max} \times d}$. The role of the Transformer-based Temporal Encoder is to aggregate this historical sequence and reason about the track’s latent state.

Inspired by BERT (Devlin 2018; Huang, Yang, and Tsai 2023), we introduce a learnable **track token**, $\mathbf{z}_m \in R^{1 \times d}$, which acts as a summary for the entire trajectory. This token is prepended to the historical sequence, and the combined input is fed into a self-attention module. The query, key, and value matrices are computed as follows:

$$\begin{aligned} \mathbf{Q} &= [\mathbf{z}_m, \mathbf{Z}_i] \mathbf{W}_T^Q, \\ \mathbf{K} &= [\mathbf{z}_m, \mathbf{Z}_i] \mathbf{W}_T^K, \\ \mathbf{V} &= [\mathbf{z}_m, \mathbf{Z}_i] \mathbf{W}_T^V, \end{aligned} \quad (7)$$

where \mathbf{W}_T^* are the learnable projection matrices and $[\cdot, \cdot]$ denotes concatenation along the sequence dimension. To prevent information leakage from the future, a causal mask is applied during the self-attention computation.

The full temporal encoder layer then processes the attention output with a feed-forward network (FFN), including residual connections and layer normalization:

$$[\hat{\mathbf{z}}_m, \hat{\mathbf{Z}}_i] = \text{FFN}(\text{SelfAttn}(\mathbf{Q}, \mathbf{K}, \mathbf{V})), \quad (8)$$

The updated track token, $\hat{\mathbf{z}}_m$, ultimately encapsulates the rich, aggregated spatio-temporal representation of the trajectory, providing a robust feature for the subsequent interaction modules.

3.3 Cue-Consistent Attention Module

Overview. Given the aggregated track features $\{\hat{\mathbf{z}}_m\}$ and new detections, our goal is to robustly associate them by mining their underlying feature consistency. To achieve this, we design a Cue-Consistency Attention Module. As illustrated in Figure 2, this module first enhances both track and detection features independently and then interacts them to produce final, matchable representations. This process can be iterated g times to progressively refine the matching cues.

Self-Information Encoder. First, we independently enhance the features of tracks and detections. For the M input track features, aggregated in a matrix $\hat{\mathbf{Z}}_m \in R^{M \times d}$, we use a vanilla self-attention layer (Vaswani et al. 2017) to mine intra-feature relationships, producing globally-aware features $\tilde{\mathbf{Z}}_m \in R^{M \times d}$.

Concurrently, we process the new detections. For each detection $\mathbf{b}_i \in \mathbf{B}^t$, we apply our geometric encoder (denoted as δ) to aggregate information from its neighborhood. Note that unlike the selective sampling in Section 3.2, here we use **all** other objects as neighbors to form the geometric triplet. This ensures that we do not miss potentially important information, given that detector outputs can be unstable. This process yields the initial detection features $\mathbf{Z}_B = [\delta(\mathbf{b}_1), \dots, \delta(\mathbf{b}_N)] \in R^{N \times d}$. These are then passed through another self-attention layer to yield enhanced features $\tilde{\mathbf{Z}}_B \in R^{N \times d}$.

Crucially, during these two self-attention steps, we cache the resulting attention score matrices, $\mathbf{E}_m \in R^{M \times M}$ and $\mathbf{E}_B \in R^{N \times N}$, for subsequent use in our cue-consistent cross-attention.

Cue-Consistent Cross-Attention. Cross-attention is a typical module for exchanging information between two feature sets (Huang et al. 2021; Wang and Solomon 2019; Sun et al. 2021). Inspired by (Chen et al. 2024), we design a *cue-consistent cross-attention* module to model the consistency between $\tilde{\mathbf{Z}}_m$ and $\tilde{\mathbf{Z}}_B$. For clarity, we demonstrate the process for updating detection features by attending to tracks; the reverse process is symmetric.

The procedure unfolds in three steps. **First**, to expand the structural information for each feature, we extract "cues" guided by the pre-computed self-attention matrices. For each detection i , its cue is constructed from the features of its top- k most similar peer detections, identified via \mathbf{E}_B :

$$\mathbf{C}_{B,i} = \text{Gather}(\tilde{\mathbf{Z}}_B, \text{topk}(\mathbf{E}_B(i, \cdot), k)) \in R^{k \times d}. \quad (9)$$

Here, $\text{topk}(\cdot, k)$ returns the indices of the k largest values, and $\text{Gather}(\cdot)$ collects the features at these indices. Similarly, for each track j , we extract its cue $\mathbf{C}_{m,j} \in R^{k \times d}$ using $\tilde{\mathbf{Z}}_m$ and \mathbf{E}_m .

Second, we compute the attention from a detection i to all tracks $j = 1, \dots, M$. The attention weight is not based on direct feature similarity, but on the consistency score s_{ij} between their respective cues, $\mathbf{C}_{B,i}$ and $\mathbf{C}_{m,j}$. We model this score by performing a feature-wise dot-product between the two cue matrices and summing the result:

$$s_{ij} = \sum_{l=1}^k \frac{(\mathbf{C}_{B,i}[l] \mathbf{W}^Q)(\mathbf{C}_{m,j}[l] \mathbf{W}^K)^T}{\sqrt{d}}, \quad (10)$$

where $\mathbf{C}[l]$ denotes the l -th feature vector in the cue matrix, and $\mathbf{W}^Q, \mathbf{W}^K$ are learnable projection matrices. This process learns to measure the consistency confidence between the two neighborhood structures.

Finally, the consistency scores $\{s_{ij}\}_{j=1}^M$ are normalized using a softmax to produce the attention weights. The updated detection feature $\mathbf{y}_{B,i}$ is then the weighted average of all track features:

$$\mathbf{y}_{B,i} = \sum_{j=1}^M \text{softmax}_j(s_{ij}) \cdot (\tilde{\mathbf{z}}_{m,j} \mathbf{W}^V). \quad (11)$$

The full matrix of updated features \mathbf{Y}_B is further sent into a feed-forward layer to obtain the enhanced message $\hat{\mathbf{Y}}_B$. The cue-consistent cross-attention is also applied in the reverse direction, yielding $\hat{\mathbf{Y}}_m$.

Computational Cost Reduction. It is important to note that computing cue-consistent cross-attention for all pairs is computationally demanding. To reduce this burden, we utilize prior knowledge from online MOT. We first predict the potential locations of historical trajectories in the current frame by applying an FFN: $p_m = \text{FFN}(\{\tilde{\mathbf{Z}}_m\})$. Interactions are then established only between objects of the same class, and are further pruned based on a class-specific distance threshold. This threshold is calculated from dataset statistics of maximum speed per class, following (Zaech et al. 2022).

3.4 Feature Matching and Update

Overview. After extracting the enhanced features $\hat{\mathbf{Y}}_m$ and $\hat{\mathbf{Y}}_B$, we first compute a feature-based affinity matrix to associate tracks with detections. Then, for each successful match, we update the track's stored neighborhood information for the next frame.

Feature-based Matching. To compute the affinity between all track-detection pairs, we expand the feature matrices $\hat{\mathbf{Y}}_m \in R^{M \times C}$ and $\hat{\mathbf{Y}}_B \in R^{N \times C}$ to a compatible shape of $R^{M \times N \times C}$ by tiling. An MLP then processes their element-wise difference to predict a match score:

$$\mathbf{A} = \text{Sigmoid}(\text{MLP}(\hat{\mathbf{Y}}_m^e - \hat{\mathbf{Y}}_B^e)), \quad (12)$$

where $\mathbf{A} \in R^{M \times N}$ is the final affinity matrix, whose entries represent matching probabilities.

Neighborhood Update. To keep the context of each track current, its neighborhood information must be updated over time. When a track j is associated with a detection i , we update its stored set of neighbor IDs with those of detection i , which are defined by the cue $\mathbf{C}_{B,i}$ from Equation 9.

Loss Function. The model is trained with a composite loss. For association, we use the Focal Loss (Lin et al. 2017) (\mathcal{L}_a) with $\alpha = -1, \gamma = 1$. For position regression, we use a smooth ℓ_1 loss (\mathcal{L}_p). The total loss is:

$$\mathcal{L} = \mathcal{L}_a + \lambda_p \mathcal{L}_p, \quad (13)$$

where the balancing weight λ_p is set to 0.5.

4 Experiment

4.1 Experiment Setup

Dataset. We conduct experiments on two large-scale autonomous driving benchmarks. The Waymo Open Dataset (Sun et al. 2020) provides 1150 sequences (798 training, 202 validation, 150 testing) of 20-second driving data with 3D labels for three classes: Vehicle, Pedestrian, and Cyclist. The nuScenes dataset (Caesar et al. 2020), our primary benchmark, comprises 1000 sequences of similar length. It is distinguished by a richer sensor suite (32-beam LiDAR, RADAR, six cameras) and provides 3D tracking annotations at 2Hz for seven object classes, making it ideal for leveraging multi-modal data.

Metrics. We follow the official nuScenes tracking benchmark protocol for evaluation. The primary metrics are Average Multi-Object Tracking Accuracy (AMOTA) and Average Multi-Object Tracking Precision (AMOTP) (Weng et al. 2020). We also report secondary metrics from CLEAR MOT (Bernardin and Stiefelhagen 2008), including Multi-Object Tracking Accuracy (MOTA), ID Switches (IDS).

Implementation details. We construct the training dataset using outputs from CenterPoint (Yin, Zhou, and Krahenbuhl 2021). Following 3DMOTFormer (Ding et al. 2023), we trained our tracker for 16 epochs on randomly sampled mini-sequences of length $T=10$. For the memory buffer, we set the maximum temporal length $T_{max}=7$ and select the top

Method	Additional Cues	Average Metrics				Class-specific AMOTA						
		AMOTA \uparrow	AMOTP \downarrow	MOTA \uparrow	IDS \downarrow	car	ped	bicycle	bus	motor	trailer	truck
CenterPoint	–	0.638	0.555	0.537	760	0.829	0.767	0.321	0.711	0.591	0.651	0.599
CBMOT	–	0.649	0.592	0.545	557	0.828	0.794	0.372	0.704	0.592	0.667	0.587
SimpleTrack \ddagger	–	0.668	0.550	0.566	575	0.823	0.796	0.407	0.715	0.674	0.673	0.587
ImmortalTracker	–	0.677	0.599	0.572	320	0.833	0.816	0.416	0.716	0.689	0.675	0.596
PolarMOT-offline \dagger	–	0.664	0.566	0.561	242	0.853	0.806	0.349	0.708	0.656	0.673	0.602
3DMOTFormer	–	0.682	0.496	0.556	438	0.821	0.807	0.374	0.749	0.705	0.696	0.626
NEBP	3D appearance	0.683	0.624	0.584	227	0.835	0.802	0.447	0.708	0.698	0.69	0.598
ShaSTA	3D appearance	0.696	0.540	0.578	473	0.838	0.81	0.41	0.733	0.727	0.704	0.65
DSC-Track(ours)	–	0.703	0.476	0.575	301	0.858	0.830	0.401	0.755	0.703	0.719	0.653

Table 1: Results on nuScenes test set using CenterPoint detections. \dagger denotes offline methods, \ddagger denotes using 10Hz data. The best results are highlighted in bold

Method	AMOTA \uparrow	AMOTP \downarrow	MOTA \uparrow	IDS \downarrow
CenterPoint	0.665	0.567	0.562	562
SimpleTrack	0.696	0.547	0.602	–
ImmortalTracker	0.702	–	0.601	–
PolarMOT-offline	0.711	–	–	213
3DMOTFormer	0.712	0.515	0.607	341
DSC-Track(ours)	0.732	0.498	0.625	298

Table 2: Comparison with state-of-the-art tracking-by-detection approaches on the nuScenes validation split.

Method	Vehicle	Pedestrian	Cyclist
	MOTA \uparrow / IDS $\% \downarrow$		
CenterPoint	55.1 / 0.26	54.9 / 1.13	57.4 / 0.83
SpOT	55.7 / 0.18	60.5 / 0.56	– / –
TrajectoryFormer	59.7 / 0.19	61.0 / 0.37	60.6 / 0.70
DSC-Track(ours)	60.5 / 0.11	61.1 / 0.27	60.9 / 0.18

Table 3: Tracking results (MOTA / IDS%) on Waymo validation split.

$k = 3$ spatially dependent objects to capture the historical information for each tracklet. In the cue-consistent attention module, we set top $k = 3$ to perform the cue-consistent cross attention. For optimization, we use the AdamW (Loshchilov 2017) optimizer with an initial learning rate of $2e^{-4}$ and a cosine decay schedule with power set to 0.8. All experiments are trained on eight NVIDIA 4090 GPUs with a batch size of one per GPU.

4.2 Main Results

To validate our approach, we compare DSC-Track against state-of-the-art methods on the nuScenes and Waymo benchmarks. On nuScenes, we compare against two categories of baselines: methods leveraging additional **3D appearance** cues like ShaSTA (Sadjadpour et al. 2023)

and NEBP (Liang and Meyer 2023), and those without, including strong competitors like 3DMOTFormer (Ding et al. 2023), PolarMOT (Kim et al. 2022), ImmortalTracker (Wang et al. 2021), SimpleTrack (Pang, Li, and Wang 2022), and CBMOT (Benbarka, Schröder, and Zell 2021). On Waymo, we compare against recent strong methods such as SpOT (Stearns et al. 2022) and TrajectoryFormer (Chen et al. 2023). For fair comparison, all methods are built upon detections from the widely-used CenterPoint (Yin, Zhou, and Krahenbuhl 2021) detector.

Results on the nuScenes Test Set. Our primary results on the challenging nuScenes test set (Table 1) show that DSC-Track establishes a new state-of-the-art with 70.3% AMOTA. Notably, this is achieved without relying on any appearance features, yet our method surpasses the best appearance-based method, ShaSTA, by 0.7% and outperforms 3DMOTFormer by a significant 2.1% in AMOTA. This superiority extends to other crucial metrics; it achieves the best AMOTP (0.476), indicating superior localization precision, and its IDS score marks a dramatic 31.3% reduction in identity switches compared to 3DMOTFormer. The class-level analysis further underscores its robustness, as it secures the highest AMOTA for 5 out of 7 categories, including the critical car, pedestrian, and bus classes.

Validation on nuScenes and Generalization to Waymo.

This strong performance is mirrored on the nuScenes validation set (Table 2), where DSC-Track again sets a new state-of-the-art with 73.2% AMOTA, surpassing 3DMOTFormer by 2.0%. To evaluate generalizability, we also test on the Waymo Open Dataset (Table 3), where our method excels at tracking stability, securing the best or tied-for-best IDS scores across all major categories: Vehicle, Pedestrian, and Cyclist.

Runtime Speed. Our approach runs at a rate of 26.77 FPS on a single NVIDIA 4090 GPU, which is efficient for real-time multi-object tracking.

Method	AMOTA	AMOTP	MOTA	IDS
w/o Geometric Encoder	71.8	0.510	61.3	450
w/o Temporal Encoder	67.5	0.580	57.0	1015
w/o Cue-Consistent Attn	70.5	0.505	60.1	950
Ours	73.2	0.498	62.5	298

Table 4: Analysis of different components using the nuScenes validation set.

Method	AMOTA	AMOTP	MOTA	IDS
w/o PPF	72.1	0.510	61.9	385
w/o Contextual	72.0	0.512	61.7	460
Ours	73.2	0.498	62.5	298

Table 5: Ablation study on the the features of geometric encoders.

4.3 Ablation Study and Analysis

Analysis of Different Components. Our ablation study in Table 4 confirms the indispensability of each module. Removing the Temporal Encoder is most detrimental, causing a 5.7% AMOTA drop and a massive surge in IDS, underscoring the necessity of historical context. Ablating the Cue-Consistent Attention leads to a 2.7% AMOTA loss and a threefold increase in IDS, validating our explicit matching mechanism. Finally, removing the Geometric Encoder results in a 1.4% AMOTA drop, proving the importance of spatial relationship encoding. These results affirm our model’s performance stems from the synergy of its components.

Analysis of Initial Feature Embedding. As shown in Table 5, our initial feature embedding design relies on the synergy between relative geometry (PPF) and contextual features. Removing the PPF-based geometry (w/o PPF) impairs pose-sensitivity, resulting in a 1.1% AMOTA drop and a 29% IDS increase. Conversely, removing contextual features (w/o Contextual) is even more detrimental, causing a 1.2% AMOTA loss and a 54% surge in IDS due to the lack of global scene context. These results confirm that both information sources are complementary and essential for constructing robust discriminative features.

Analysis of Cue-Consistent Cross-Attention. We compare our Cue-Consistent Attention with two alternatives in Table 6. A Vanilla Cross-Attention struggles without explicit guidance, leading to a 0.4% AMOTA drop and a 31% IDS increase. A simpler Max Pooling aggregation performs worse, with AMOTA decreasing by 0.9%, as it discards contextual information. This superior performance validates our design: by using pre-computed self-attention scores to guide the interaction, it robustly aggregates consistent features for more accurate associations.

Method	AMOTA	AMOTP	MOTA	IDS
Max Pooling	72.3	0.515	61.4	480
Vanilla Cross-Attn	72.8	0.508	61.8	390
Ours	73.2	0.498	62.5	298

Table 6: Ablation study on different interaction mechanisms for cue-consistency.




Figure 4: Qualitative results of DSC-Track on the nuScenes validation set. Our tracker successfully handles a long-term occlusion during a turn by leveraging stable geometric cues from the environment (e.g., the roadside), correctly re-identifying the target (ID 6) upon reappearance.

4.4 Qualitative Results

Figure 4 showcases the robustness of DSC-Track in a challenging scenario with long-term occlusion during a turn. While conventional trackers often fail due to non-linear motion and prolonged target absence, our method successfully bridges the gap. As shown, despite the target vehicle (ID 6) being fully occluded, our approach leverages stable geometric cues from the environment (e.g., the roadside relationship). This allows our tracker to maintain the target’s state during occlusion and robustly re-identify it upon reappearance, ensuring trajectory integrity.

5 Conclusion

In this paper, we propose **DSC-Track** to address data association in 3D MOT. Based on cue-consistency, our method leverages rotation-invariant Point Pair Features (PPF) and a cue-consistent attention strategy to match trajectories and detections by their neighborhood structures, rather than individual features. Experiments show DSC-Track achieves state-of-the-art performance and significantly reduces identity switches in complex scenarios. Our work validates modeling higher-order relational consistency over object-level cues as a promising direction for robust tracking.

Acknowledgments

This research was supported by The National Nature Science Foundation of China (Grant Nos: 62402417, 62273301, 62273302), in part by "Pioneer" and "Leading Goose" R&D Program of Zhejiang (Grant No. 2025C02026), in part by the Key R&D Program of Ningbo (Grant Nos: 2024Z115, 2025Z035), in part by Yongjiang Talent Introduction Programme (Grant No: 2023A-197-G).

References

- Ba, J.; Kiros, J. R.; and Hinton, G. E. 2016. Layer Normalization. *ArXiv*, abs/1607.06450.
- Benbarka, N.; Schröder, J.; and Zell, A. 2021. Score refinement for confidence-based 3D multi-object tracking. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 8083–8090. IEEE.
- Bernardin, K.; and Stiefelhagen, R. 2008. Evaluating Multiple Object Tracking Performance: The CLEAR MOT Metrics. *EURASIP Journal on Image and Video Processing*, 2008: 1–10.
- Bewley, A.; Ge, Z.; Ott, L.; Ramos, F. T.; and Upcroft, B. 2016. Simple online and realtime tracking. *IEEE International Conference on Image Processing (ICIP)*, 3464–3468.
- Brasó, G.; and Leal-Taixé, L. 2020. Learning a neural solver for multiple object tracking. In *IEEE/CVF conference on computer vision and pattern recognition*, 6247–6257.
- 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 *IEEE/CVF conference on computer vision and pattern recognition*, 11621–11631.
- Chen, H.; Yan, P.; Xiang, S.; and Tan, Y. 2024. Dynamic Cues-Assisted Transformer for Robust Point Cloud Registration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 21698–21707.
- Chen, X.; Shi, S.; Zhang, C.; Zhu, B.; Wang, Q.; Cheung, K. C.; See, S.; and Li, H. 2023. Trajectoryformer: 3d object tracking transformer with predictive trajectory hypotheses. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 18527–18536.
- Chu, P.; Wang, J.; You, Q.; Ling, H.; and Liu, Z. 2023. Transmot: Spatial-temporal graph transformer for multiple object tracking. In *Proceedings of the IEEE/CVF Winter Conference on applications of computer vision*, 4870–4880.
- Devlin, J. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Ding, S.; Rehder, E.; Schneider, L.; Cordts, M.; and Gall, J. 2022. End-to-End Single Shot Detector Using Graph-Based Learnable Duplicate Removal. In *DAGM German Conference on Pattern Recognition*, 375–389. Springer.
- Ding, S.; Rehder, E.; Schneider, L.; Cordts, M.; and Gall, J. 2023. 3DMOTFormer: Graph Transformer for Online 3D Multi-Object Tracking. In *IEEE/CVF International Conference on Computer Vision*, 9784–9794.
- Drost, B.; Ulrich, M.; Navab, N.; and Ilic, S. 2010. Model globally, match locally: Efficient and robust 3D object recognition. In *2010 IEEE computer society conference on computer vision and pattern recognition*, 998–1005. Ieee.
- He, J.; Huang, Z.; Wang, N.; and Zhang, Z. 2021. Learnable graph matching: Incorporating graph partitioning with deep feature learning for multiple object tracking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5299–5309.
- Huang, K.-C.; Yang, M.-H.; and Tsai, Y.-H. 2023. Delving into motion-aware matching for monocular 3d object tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6909–6918.
- Huang, S.; Gojcic, Z.; Usvyatsov, M.; Wieser, A.; and Schindler, K. 2021. Predator: Registration of 3d point clouds with low overlap. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, 4267–4276.
- Kim, A.; Brasó, G.; Ošep, A.; and Leal-Taixé, L. 2022. PolarMOT: How far can geometric relations take us in 3D multi-object tracking? In *European Conference on Computer Vision*, 41–58. Springer.
- Lang, A. H.; Vora, S.; Caesar, H.; Zhou, L.; Yang, J.; and Beijbom, O. 2019. Pointpillars: Fast encoders for object detection from point clouds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 12697–12705.
- Liang, M.; and Meyer, F. 2023. Neural enhanced belief propagation for multiobject tracking. *IEEE Transactions on Signal Processing*, 72: 15–30.
- Lin, T.-Y.; Goyal, P.; Girshick, R. B.; He, K.; and Dollár, P. 2017. Focal Loss for Dense Object Detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42: 318–327.
- Loshchilov, I. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Pang, Z.; Li, Z.; and Wang, N. 2022. Simpletrack: Understanding and rethinking 3d multi-object tracking. In *European Conference on Computer Vision*, 680–696. Springer.
- Qi, C. R.; Liu, W.; Wu, C.; Su, H.; and Guibas, L. J. 2018. Frustum pointnets for 3d object detection from rgb-d data. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 918–927.
- Sadjadpour, T.; Li, J.; Ambrus, R.; and Bohg, J. 2023. Shasta: Modeling shape and spatio-temporal affinities for 3d multi-object tracking. *IEEE Robotics and Automation Letters*, 9(5): 4273–4280.
- Stearns, C.; Rempe, D.; Li, J.; Ambrus, R.; Zakharov, S.; Guizilini, V.; Yang, Y.; and Guibas, L. J. 2022. Spot: Spatiotemporal modeling for 3d object tracking. In *European Conference on Computer Vision*, 639–656. Springer.
- Sun, J.; Shen, Z.; Wang, Y.; Bao, H.; and Zhou, X. 2021. LoFTR: Detector-free local feature matching with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8922–8931.

Sun, P.; Kretschmar, H.; Dotiwalla, X.; Chouard, A.; Patnaik, V.; Tsui, P.; Guo, J.; Zhou, Y.; Chai, Y.; Caine, B.; 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.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.

Wang, Q.; Chen, Y.; Pang, Z.; Wang, N.; and Zhang, Z. 2021. Immortal tracker: Tracklet never dies. *ArXiv*, abs/2111.13672.

Wang, Y.; and Solomon, J. M. 2019. Deep closest point: Learning representations for point cloud registration. In *Proceedings of the IEEE/CVF international conference on computer vision*, 3523–3532.

Weng, X.; Wang, J.; Held, D.; and Kitani, K. 2020. 3d multi-object tracking: A baseline and new evaluation metrics. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 10359–10366. IEEE.

Wojke, N.; Bewley, A.; and Paulus, D. 2017. Simple online and realtime tracking with a deep association metric. *IEEE International Conference on Image Processing (ICIP)*, 3645–3649.

Xinshuo Weng, Y. M. K. K., Yongxin Wang. 2020. GNN3DMOT: Graph Neural Network for 3D Multi-Object Tracking with Multi-Feature Learning. In *IEEE conference on computer vision and pattern recognition*.

Yin, T.; Zhou, X.; and Krahenbuhl, P. 2021. Center-based 3d object detection and tracking. In *IEEE/CVF conference on computer vision and pattern recognition*, 11784–11793.

Yu, H.; Qin, Z.; Hou, J.; Saleh, M.; Li, D.; Busam, B.; and Ilic, S. 2023. Rotation-invariant transformer for point cloud matching. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5384–5393.

Zaech, J.-N.; Liniger, A.; Dai, D.; Danelljan, M.; and Van Gool, L. 2022. Learnable online graph representations for 3d multi-object tracking. *IEEE Robotics and Automation Letters*, 7(2): 5103–5110.

Zeng, Y.; Ma, C.; Zhu, M.; Fan, Z.; and Yang, X. 2021. Cross-Modal 3D Object Detection and Tracking for Auto-Driving. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 3850–3857.

Zhang, Y.; Sun, P.; Jiang, Y.; Yu, D.; Yuan, Z.; Luo, P.; Liu, W.; and Wang, X. 2021. ByteTrack: Multi-Object Tracking by Associating Every Detection Box. In *European Conference on Computer Vision*.

Zhang, Y.; Wang, X.; Ye, X.; Zhang, W.; Lu, J.; Tan, X.; Ding, E.; Sun, P.; and Wang, J. 2023. ByteTrackV2: 2D and 3D Multi-Object Tracking by Associating Every Detection Box. *ArXiv*, abs/2303.15334.

Zhou, Y.; and Tuzel, O. 2018. Voxelnet: End-to-end learning for point cloud based 3d object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4490–4499.