

# ShaSTA: Modeling Shape and Spatio-Temporal Affinities for 3D Multi-Object Tracking

Tara Sadjadpour<sup>1</sup>, Jie Li<sup>2</sup>, Rares Ambrus<sup>3</sup>, and Jeannette Bohg<sup>1</sup>

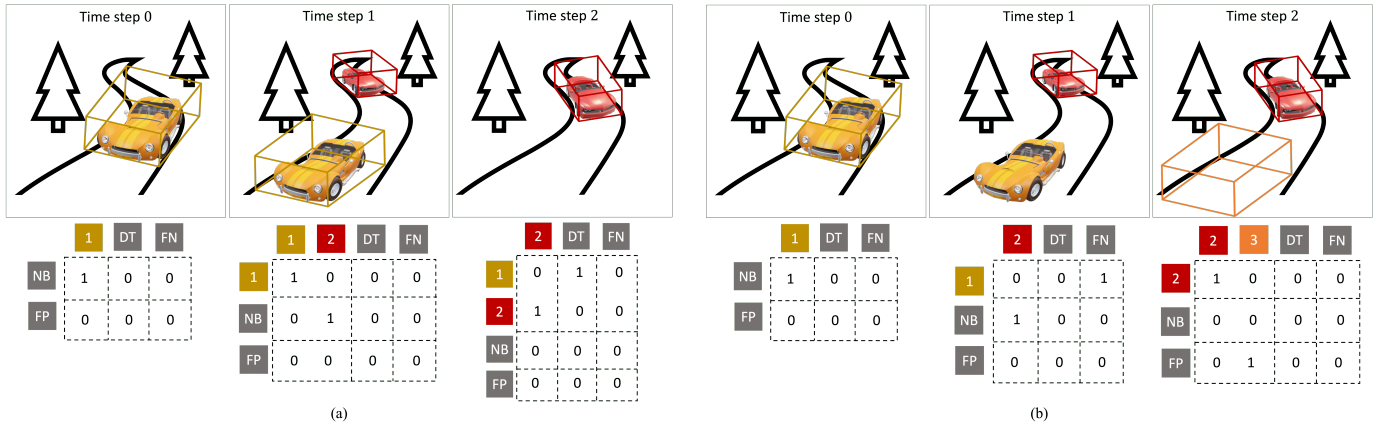


Fig. 1: Examples for tracking scenarios (top) and associated ground-truth affinity matrices (bottom) that our algorithm, ShaSTA, must estimate. In all affinity matrices, the rows correlate with previous frame tracks that are represented with single-history bounding box detections, while the columns correlate with current frame detections. The affinity matrices are augmented with two rows for *newborn track* (NB) and *false-positive* (FP) anchors that the current frame detections can match with, and two columns for *dead track* (DT) and *false-negative* (FN) anchors that the previous frame tracks can match with. These four augmentations are learned representations that capture the essence of these four detection and track types for each frame. The bounding box color corresponds to the track ID. At time step  $t = 0$ , both figures (a) and (b) start with affinity matrix  $A_0$ , where the yellow car is detected and is matched with NB, thus initializing Track 1. For Figure (a), at  $t = 1$ ,  $A_1$  shows that the yellow car is detected again and matched with the previous frame's Track 1, while the red car is detected as a newborn with Track ID 2. Then, at  $t = 2$ , there is only one detection for the red car, so it is matched with the previous frame's Track 2, while Track 1 is matched with DT in  $A_2$ . In Figure (b), at time step  $t = 1$ , the red car is detected as a newborn with Track ID 2, while the yellow car is not detected at all so Track 1 from  $t = 1$  is matched with FN. Finally, at  $t = 2$ , the red car is detected again and matched with Track 2, while a detection is found in a region with no object so it is matched with FP in  $A_2$ .

**Abstract**—Multi-object tracking (MOT) is a cornerstone capability of any robotic system. Tracking quality is largely dependent on the quality of input detections. In many applications, such as autonomous driving, it is preferable to over-detect objects to avoid catastrophic outcomes due to missed detections. As a result, current state-of-the-art 3D detectors produce high rates of false-positives to ensure a low number of false-negatives. This can negatively affect tracking by making data association and track lifecycle management more challenging. Additionally, occasional false-negative detections due to difficult scenarios like occlusions can harm tracking performance. To address these issues in a unified framework, we propose *ShaSTA* which learns shape and spatio-temporal affinities between tracks and detections in consecutive frames. The affinity is a probabilistic matching that leads to robust data association, track lifecycle management, false-positive elimination, false-negative propagation, and sequential track confidence refinement. We offer the first self-contained framework that addresses all aspects of the 3D MOT problem. We quantitatively evaluate ShaSTA on the nuScenes tracking benchmark with 5 metrics, including the most common tracking accuracy metric called AMOTA, to demonstrate how

ShaSTA may impact the ultimate goal of an autonomous mobile agent. ShaSTA achieves 1st place amongst LiDAR-only trackers that use CenterPoint detections. The open-source code for reproducing and extending our work can be found [here](#).

**Index Terms**—Visual Tracking, Deep Learning for Visual Perception, Computer Vision for Transportation

## I. INTRODUCTION

WE address the problem of 3D multi-object tracking (MOT) using LiDAR point cloud data as input. 3D MOT is an essential capability for autonomous agents to navigate effectively in their environments. For example, autonomous cars need to understand the motion of surrounding traffic agents to drive safely. Online tracking is achieved by (1) accurately matching uncertain detections to existing tracks and (2) determining when to birth and kill tracks of unmatched observations. These two processes are referred to as *data association* and *track lifecycle management*, respectively.

Most recent approaches [1–11] toward 3D MOT in autonomous driving apply the tracking-by-detection paradigm whereby off-the-shelf detections are available from a state-of-the-art 3D detector. Though this framework has been successful, there exists a practical gap between the objectives of 3D detection and tracking. Since missing a detection can lead to catastrophic events in self-driving, 3D detectors are evaluated on the full recall spectrum of detections (e.g., mAP [12]) and are deployed with the goal of having high recall. As a result, they tend to generate high rates of false-positives to ensure a low number of false-negatives [7, 13]. These redundant

Manuscript received: May 9, 2023; Revised: August 14, 2023; Accepted: September 12, 2023.

This paper was recommended for publication by Editor Cesar Cadena Lerma upon evaluation of the Associate Editor and Reviewers' comments. Toyota Research Institute provided funds to support this work.

<sup>1</sup>Tara Sadjadpour and Jeannette Bohg are with School of Engineering, Computer Science Department, Stanford University, United States {tsadja, bohg}@stanford.edu

<sup>2</sup>Jie Li is with NVIDIA, Santa Clara, California, United States jieli@nvidia.com

<sup>3</sup>Rares Ambrus is with Toyota Research Institute, Los Altos, California, United States rares.ambrus@tri.global

Digital Object Identifier (DOI): see top of this page.

detections lead to highly cluttered outputs that make data association and track lifecycle management more challenging. Additionally, occasional false-negative detections due to factors such as heavy occlusion can cause significant fragmentations in the tracks or lead to the premature termination of tracks.

In this work, we address the task of LiDAR-only tracking. In contrast to camera-only and camera-fusion tracking [3, 9–11, 14], LiDAR tracking offers many practical advantages, balancing high accuracy with a less complex system configuration and lower computational demands. Still, LiDAR-only systems face significant challenges due to the sparse, unstructured nature of LiDAR point clouds. Recent works in LiDAR-only tracking-by-detection [4–6, 8] have attempted to develop learning-based and statistical approaches to better handle data association, false-positive elimination, false-negatives propagation, and track lifecycle management. However, these approaches do not effectively leverage spatio-temporal and shape information from the LiDAR sensor input, and in some cases, only rely on detected bounding box parameters for spatio-temporal information.

Our **main contribution** is a 3D MOT framework called *ShaSTA* that models **shape** and **spatio-temporal affinities** between tracks and detections in consecutive frames. By better understanding objects' shapes and spatio-temporal contexts, *ShaSTA* improves data association, false-positive (FP) detection elimination, false-negative (FN) track propagation, newborn (NB) track initialization, dead track (DT) termination, and track confidence refinement. At its core, *ShaSTA* learns an affinity matrix that encodes shape and spatio-temporal information to obtain a probabilistic matching between detections and tracks. *ShaSTA* also learns representations of FP, NB, FN, and DT *anchors* for each frame. By computing the affinity between detections and tracks to these learned representations, we can classify detections as newborn tracks or false-positives, while tracks can be classified as dead tracks or false-negatives. Once we use these matches to form tracks, *ShaSTA* continues to leverage the affinity estimation with a novel sequential track confidence refinement technique that represents track score refinement as a time-dependent problem and leverages information extracted from shape and spatio-temporal encodings to obtain significant improvements in overall tracking accuracy.

This overall approach not only aids in tracking performance but also addresses concerns for real-world deployment and downstream decision-making tasks by eliminating false-positive tracks and minimizing false-negative tracks. *ShaSTA* offers an efficient and practical implementation that can re-use the LiDAR backbone of most off-the-shelf LiDAR-based 3D detectors. Without loss of generality, we choose the LiDAR backbone from CenterPoint [7] for a fair comparison against most previous works. We demonstrate that our approach achieves state-of-the-art results on the nuScenes tracking benchmark [15] where we achieve significant gains in overall tracking accuracy and precision, while improving the raw number of true-positive, false-positive, and false-negative tracks. We also include ablation studies to analyze the contribution of each component.

## II. RELATED WORK

### A. LiDAR-based 3D Detection

Most recent 3D MOT works follow the tracking-by-detection paradigm [1, 3, 5, 8, 10, 16], making the quality of 3D detectors a key component to tracking performance [3, 17, 18]. LiDAR-based 3D detectors extend the image-based 2D detection algorithms [19, 20] to 3D. The typical detector starts with an encoder [21–23] applied to unstructured point cloud data to extract an intermediate 3D feature grid [13] or bird's-eye-view (BEV) feature map [7, 24]. Then, a decoder is applied to extract objectness [25] and detailed object attributes [7, 26]. Our approach proposes to leverage the intermediate BEV feature map for robust multi-object tracking, which can be generalized to most modern 3D detectors. CenterPoint detections have become a standardized 3D detection in the 3D MOT community, since the authors of the work released the official detection files for all data splits. As of this writing, CenterPoint has the highest mAP of all publicly-released 3D detections. Thus, in this work, we apply CenterPoint [7] for fair comparison against existing 3D MOT techniques.

### B. 3D Multi-Object Tracking

With the advances in 3D detection, the 3D MOT community has been focusing on four directions to establish robust tracking within the tracking-by-detection paradigm: motion prediction, data association, track lifecycle management, and confidence refinement.

In 3D MOT, abstracted detection predictions are often used for motion prediction and data association, which is completed using pair-wise feature distance between tracks and detections [1, 2, 7]. AB3DMOT [1] provides a baseline to combine Kalman Filters and Intersection Over Union (IoU) association for 3D MOT. Subsequent works in this line explore variations of data association metrics, such as the Mahalanobis distance to leverage uncertainty measurements in the Kalman Filter [2]. As an alternative, CenterPoint [7] uses high-fidelity instance velocity predictions as a motion model to complete data association. Unlike these works that only use low-dimensional bounding boxes and explicit motion models for tracking, we propose to leverage intermediate representations from a detection network for richer information about each object, while also modeling the global relationship of all the objects in a scene to improve data association.

Additionally, track lifecycle management has been reasonably successful with heuristics that are tuned to each dataset [1, 2, 7], but these hand-tuned techniques have shortcomings in more complex situations. Recently, OGR3MOT [8] has proposed combining tracking and predictive modeling through a unified graph representation that learns data associations and track lifecycle management, FP elimination, and FN interpolation. Additionally, [5] proposes a statistical approach for data association under a Bayesian estimation framework. This work was later extended with Neural-Enhanced Belief Propagation (NEBP) [6] to include a learning-based aspect to the framework for FP suppression. Though these techniques yield competitive results compared to heuristics, their main

shortfall is their inability to capture object shape information. Unlike these techniques, our approach leverages raw LiDAR data to capture shape information about objects in our scene, which we can use to further improve tracking performance.

Finally, track confidence is a less-explored but important aspect of the tracking problem. Track confidence indicates the quality of a track relative to tracks for the same object class within a given frame. Most 3D MOT algorithms assign the track confidence to be the same as that of the confidence score that comes with off-the-shelf detections matched to the track [1, 2, 7, 8]. In a recent effort, [4] used an existing detection score refinement technique from [27] to obtain new tracking scores and improve the overall tracking performance. Since [4] directly applies a detection score refinement technique to the track score, it overlooks the fact that tracks are time-dependent quantities that cannot be treated in isolation for the current time step. In this paper, we argue that track score refinement should be a sequential process that reflects changing environmental factors. Thus, we propose a first-of-its-kind approach that represents track score refinement as a sequential problem and leverages shape and spatio-temporal information encoded in our learned affinity matrix to significantly improve tracking accuracy.

Another line of 3D MOT algorithms focuses on using camera-LiDAR fusion to learn data association, track lifecycle management, and FP suppression [3, 9–11]. However, in this work we focus on LiDAR-only 3D MOT due to the practical advantages it offers, namely cost efficiency and reduced calibration and synchronization issues with fewer sensors.

### III. SHASTA

Our algorithm operates under the tracking-by-detection paradigm and is visualized in Fig. 2. ShaSTA extracts shape and spatio-temporal relationships between consecutive frames to learn affinity matrices that capture associations between detections and tracks. ShaSTA learns bounding box and shape representations for FP, NB, FN, and DT *anchors*, where each anchor captures the common features between the set of detections and tracks that falls under each of these four categories in a given frame. Thus, ShaSTA’s affinity matrix not only estimates the probabilities of previous frame tracks matching with current frame detections, but also considers previous frame tracks matching with DT or FN anchors and current frame detections matching with NB or FP anchors. Unlike past techniques that only use low-dimensional bounding box representations, ShaSTA leverages shape and spatio-temporal information from the LiDAR sensor input, resulting in a robust data association technique that effectively accounts for FP detection elimination, FN track propagation, NB track initialization, and DT termination.

#### A. Affinity Matrix Usage for Online Tracking

Allowing for up to  $N_{max}$  detections per frame, ShaSTA estimates the affinity matrix  $A \in \mathbb{R}^{(N_{max}+2) \times (N_{max}+2)}$ . As shown in Figure 1, the rows of our affinity matrix represent the previous frame tracks and the columns represent our current frame detections. The two augmented columns are for the DT

and FN anchors, so previous frame tracks can match with DT or FN. Similarly, because current frame detections can match with NB or FP, the affinity matrix has two augmented rows for the NB and FP anchors. Though there can be no more than one match for each bounding box in the current and previous frames, the same is not true for our learned anchors, i.e. more than one current frame bounding box can match with the NB or FP anchors, respectively. To handle this situation, we create a forward matching affinity matrix  $A_{fm} \in \mathbb{R}^{N_{max} \times (N_{max}+2)}$  by removing the two row augmentations from  $A$  and applying a row-wise softmax so that more than one previous frame track can have sufficient probability to match with the DT and FN anchors, respectively:

$$A_{fm} = \text{softmax}_{\text{row}}(A_1). \quad (1)$$

Using similar logic, we create a backward matching affinity matrix  $A_{bm} \in \mathbb{R}^{(N_{max}+2) \times N_{max}}$  by removing the two column augmentations from  $A$  and applying a column-wise softmax so that more than one current frame detection can have sufficient probability to match with the NB and FP anchors, respectively:

$$A_{bm} = \text{softmax}_{\text{col}}(A_2). \quad (2)$$

Here, forward matching indicates that we want to find the best current frame match for each previous frame track, while backward matching refers to finding the best previous frame match for each current frame detection.

To form tracks, we combine the affinity matrix outputs and the matching algorithm from [7] as follows. We take  $A_{fm}$  and label any previous frame track that has a probability above the threshold  $\tau_{dt}$  in the DT column as a DT, and we forward propagate any previous frame track into the current frame if it has a probability above the threshold  $\tau_{fn}$  in the FN column. Therefore, in addition to provided detections, *forward propagation* creates a new detection from an existing track for the current frame to handle occluded objects that off-the-shelf detections missed. We create a box in the current frame  $b'_t := (x', y', z, w, l, h, r_y)$  by moving the 2D center coordinate of the existing track  $b_{t-1} := (x, y, z, w, l, h, r_y)$  to the next frame based on the estimated velocity provided by our off-the-shelf 3D detector and  $\Delta t$  between the two frames, i.e.  $x' = x + v_x \Delta t$  and  $y' = y + v_y \Delta t$ . All previous frame tracks that do not qualify as FN or DT are kept as is. Analogously, we take  $A_{bm}$  and remove any current frame detection that has a probability above the threshold  $\tau_{fp}$  in the FP row and we label any current frame detection that has a value above the threshold  $\tau_{nb}$  in the NB row as an NB. All current frame detections that do not qualify as FP or NB are kept as is. See Section V-B for details on the threshold values we choose.

With these detections, we then run the greedy algorithm from [7]. In its original form, the greedy algorithm takes all unmatched detections and initializes NBs with them. However, we check if an unmatched detection is labeled as an NB in the affinity matrix based on the threshold  $\tau_{nb}$  and whether it is outside the maximum distance threshold with respect to other tracks to initialize it as an NB. Otherwise, it is discarded. Additionally, we check if an unmatched track is labeled as a DT based on the threshold  $\tau_{dt}$  and whether it is outside the maximum distance threshold with respect to other

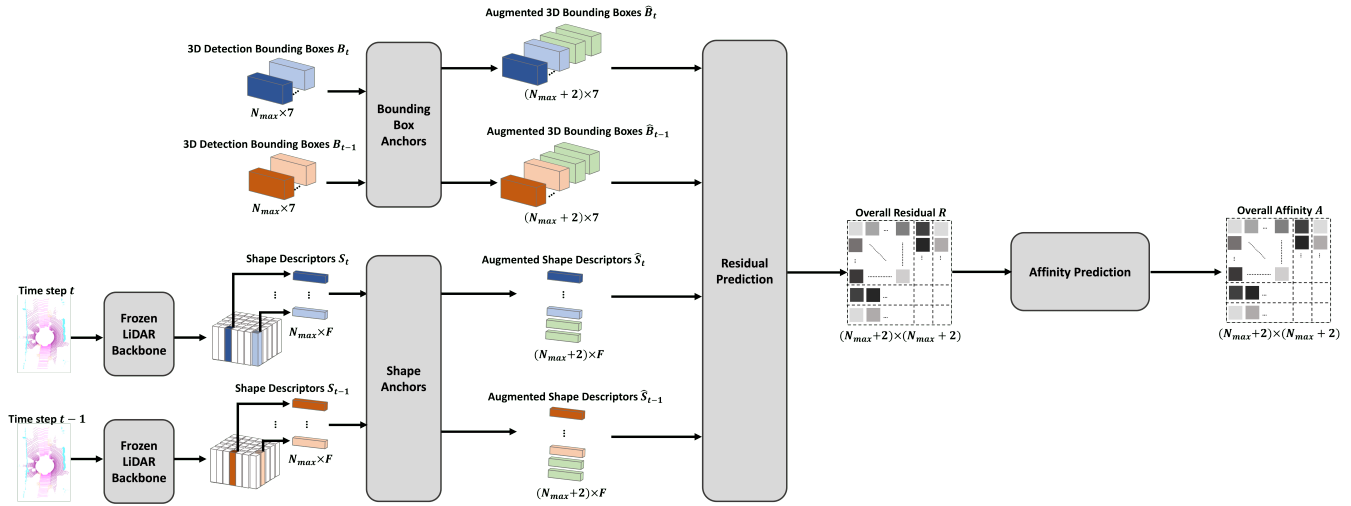


Fig. 2: Algorithm Flowchart. Going from top to bottom and left to right, ShaSTA uses  $N_{max}$  low-dimensional bounding boxes from the current and previous frames to learn bounding box representations for the FP, NB, FN, and DT anchors. The FP and NB anchors are appended to previous frame bounding boxes to form the augmented bounding boxes  $\hat{B}_{t-1}$ , while the FN and DT anchors are added to the current frame detections to form the augmented bounding boxes  $\hat{B}_t$ . Using the pre-trained frozen LiDAR backbone from our off-the-shelf detector, we extract shape descriptors that are then used to learn shape representations for our anchors. The learned shape representations for the anchors are appended to create augmented shape descriptors. These augmented bounding boxes and shape descriptors are then used to find a residual that captures the spatio-temporal and shape similarities between current and previous frame detections, as well as between detections and anchors. The residual is then used to predict the affinity matrix for probabilistic data associations, track lifecycle management, FP elimination, FN propagation, and track confidence refinement.

detections to terminate it as a DT. If both of these conditions are not met, then the retention frame limit is used from the greedy algorithm before terminating the unmatched track [7].

### B. Learning to Predict Affinity Matrices

ShaSTA uses off-the-shelf 3D detections, as well as the pre-trained LiDAR backbone used to generate the detections. ShaSTA first learns bounding box and shape representations for the anchors so we can later use this information to learn a residual that encodes spatio-temporal and shape similarities not only between current frame detections and previous frame tracks, but also between current frame detections and FP and NB anchors or between previous frame tracks and FN and DT anchors. The residual is then used to predict affinity matrices that assign probabilities for matching detections to tracks, eliminating FPs, propagating FNs, initializing NBs, and terminating DTs.

#### 1) Learning Bounding Box Representations for Anchors:

Our goal is to learn bounding box representations for each of the FP, FN, NB, and DT anchors in a given frame pair. The aim of the bounding box representation of each anchor is to capture the commonalities between detection bounding boxes that fall under each of these four categories.

We use off-the-shelf 3D detections, which include each 3D bounding box's center coordinate, dimensions, and yaw rotation angle:  $b := (x, y, z, w, l, h, r_y)$ . We take the set of bounding boxes for the current frame  $B_t \in \mathbb{R}^{N_{max} \times 7}$  and previous frame  $B_{t-1} \in \mathbb{R}^{N_{max} \times 7}$ . We fix the maximum number of detections per frame to be  $N_{max}$ ; the matrix is zero padded if there are fewer than  $N_{max}$  detections, and we sample the top  $N_{max}$  detections if the threshold is exceeded.

Then, we create a learned bounding box representation at time step  $t$  for FP and NB anchors using the current frame detections  $B_t$  as follows, where each  $\sigma$  represents an MLP:  $b_{fp} = \sigma_{fp}^b(B_t)$  and  $b_{nb} = \sigma_{nb}^b(B_t)$ .

Similarly, we find learned bounding box representations for FN and DT anchors using previous frame detections  $B_{t-1}$ :  $b_{fn} = \sigma_{fn}^b(B_{t-1})$  and  $b_{dt} = \sigma_{dt}^b(B_{t-1})$ .

For all four cases, we apply the absolute value to the MLP outputs corresponding to  $(w, l, h)$ , since dimensions need to be non-negative values. We then concatenate  $b_{fp} \in \mathbb{R}^7$  and  $b_{nb} \in \mathbb{R}^7$  to  $B_{t-1}$  to get  $\hat{B}_{t-1} \in \mathbb{R}^{(N_{max}+2) \times 7}$ , as well as  $b_{fn} \in \mathbb{R}^7$  and  $b_{dt} \in \mathbb{R}^7$  to  $B_t$  to get  $\hat{B}_t \in \mathbb{R}^{(N_{max}+2) \times 7}$ . The reason we append FP and NB anchors to  $B_{t-1}$  is so that current frame detections can match to them, and the same logic applies for appending FN and DT anchors to  $B_t$ .

2) Learning Shape Representations for Anchors: In this section, we extract shape descriptors for each 3D detection using the pre-trained LiDAR backbone of our off-the-shelf detector to leverage spatio-temporal and shape information from the raw LiDAR data. In similar fashion to the previous section, our goal is to learn shape representations for the FP, FN, NB, and DT anchors using the shape descriptors from existing detections. These shape descriptors will be used to match detections to each other or one of the anchors if they share similar shape information.

We pass the current frame's 4D LiDAR point cloud with an added temporal dimension into the frozen pre-trained VoxelNet [13] LiDAR backbone used in the CenterPoint [7] detector. This outputs a bird's-eye-view (BEV) map for the current frame, where each voxelized region encodes a high-dimensional volumetric representation of the shape information in that region. Then, for each current frame detection bounding box, we extract a shape descriptor using bilinear interpolation from the BEV map as in [7] using the bounding box center, left face center, right face center, front face center, and back face center. We concatenate these 5 shape features to create the overall shape descriptor for each bounding box. Note that in BEV, the box center, bottom face center, and top face center all project to the same point, so we forgo extracting the latter two centers. We accumulate all of the shape features ex-

tracted for the current frame detections and call this cumulative shape descriptor  $S_t \in \mathbb{R}^{N_{max} \times F}$ . We repeat the same procedure for the previous frame's LiDAR point cloud and bounding boxes to get the overall shape descriptor  $S_{t-1} \in \mathbb{R}^{N_{max} \times F}$ .

Using the extracted shape features for the current frame  $S_t$ , we obtain our learned shape descriptors for the FP and NB anchors with MLPs  $\sigma_{fp}$  and  $\sigma_{nb}$ , respectively:  $s_{fp} = \sigma_{fp}^s(S_t)$  and  $s_{nb} = \sigma_{nb}^s(S_t)$ .

Likewise, we use our previous frame's shape features  $S_{t-1}$  to learn shape descriptors for FN and DT anchors as such:  $s_{fn} = \sigma_{fn}^s(S_{t-1})$  and  $s_{dt} = \sigma_{dt}^s(S_{t-1})$ .

Just as we did in the previous section, we concatenate  $s_{fp} \in \mathbb{R}^F$  and  $s_{nb} \in \mathbb{R}^F$  to  $S_{t-1}$  to get the augmented shape descriptor  $\hat{S}_{t-1} \in \mathbb{R}^{(N_{max}+2) \times F}$ , as well as  $s_{fn} \in \mathbb{R}^F$  and  $s_{dt} \in \mathbb{R}^F$  to  $S_t$  to get  $\hat{S}_t \in \mathbb{R}^{(N_{max}+2) \times F}$ .

3) *Residual Prediction*: Using the augmented 3D bounding boxes and shape descriptors, we aim to find three residuals that measure the similarities between current and previous frames' bounding box and shape representations. Since the boxes are low-dimensional abstractions, we aim to maximize the amount of spatio-temporal information we extract from them. Thus, we obtain two residuals for the augmented 3D bounding boxes called the VoxelNet and bounding box residuals,  $R_v$  and  $R_b$  respectively. We also learn one shape residual  $R_s$  between the augmented shape descriptors. We obtain our overall residual  $R$  that captures the spatio-temporal and shape similarities by taking a weighted sum of the three residuals, where the weights are also learned.

**VoxelNet Residual.** Our first residual  $R_v \in \mathbb{R}^{(N_{max}+2) \times (N_{max}+2)}$  is a variation of the VoxelNet [13] bounding box residual. Although the VoxelNet residual from [13] is used for supervised learning on the bounding box predictions, we adapt the VoxelNet residual for the end-goal of data association between  $\hat{B}_{t-1}$  and  $\hat{B}_t$  to capture their similarities based on 3D box center, dimension, and rotation.

**Bounding Box Residual.** The learned bounding box residual  $R_b \in \mathbb{R}^{(N_{max}+2) \times (N_{max}+2)}$  is found by first expanding  $\hat{B}_{t-1}$  and  $\hat{B}_t$  and concatenating them to get a matrix  $\hat{B} \in \mathbb{R}^{(N_{max}+2) \times (N_{max}+2) \times 6}$ . Note that for this step we only use the box centers from  $\hat{B}_{t-1}$  and  $\hat{B}_t$ . Then, we obtain our residual  $R_b \in \mathbb{R}^{(N_{max}+2) \times (N_{max}+2)}$  with an MLP  $\sigma_r^b$  as such:  $R_b = \sigma_r^b(\hat{B})$ .

**Shape Residual.** We use  $\hat{S}_t$  and  $\hat{S}_{t-1}$  to obtain the learned shape residual  $R_s$  between the two frames. We expand  $\hat{S}_{t-1}$  and  $\hat{S}_t$  and concatenate them to get a matrix  $\hat{S} \in \mathbb{R}^{(N_{max}+2) \times (N_{max}+2) \times 2F}$ . Then, we obtain our residual  $R_s \in \mathbb{R}^{(N_{max}+2) \times (N_{max}+2)}$  with an MLP:  $R_s = \sigma_r^s(\hat{S})$ .

**Overall Residual.** The overall residual  $R$  is a weighted sum of the three residuals we previously obtained:  $R_v$ ,  $R_b$ , and  $R_s$ . We concatenate  $\hat{B}$  and  $\hat{S}$  that we previously obtained in the early steps to create our input  $\hat{W} \in \mathbb{R}^{(N_{max}+2) \times (N_{max}+2) \times (2F+6)}$ . Then, we pass this through an MLP  $\sigma_\alpha$  to get the learned weights  $\alpha \in \mathbb{R}^{(N_{max}+2) \times (N_{max}+2) \times 3}$  as follows:  $\alpha = \sigma_\alpha(\hat{W})$ . We split  $\alpha$  into matrices  $\alpha_v, \alpha_b, \alpha_s \in \mathbb{R}^{(N_{max}+2) \times (N_{max}+2)}$ , and obtain our overall residual  $R \in \mathbb{R}^{(N_{max}+2) \times (N_{max}+2)}$ :

$$R = \alpha_v \odot R_v + \alpha_b \odot R_b + \alpha_s \odot R_s. \quad (3)$$

Note that  $\odot$  is the Hadamard product.

4) *Affinity Matrix Estimation*: Given our overall residual, we have information about the pairwise similarities between current frame detections and previous frame tracks, as well as each of the four augmented anchors. We want to use these spatio-temporal and shape relationships to learn probabilities for matching the detections and tracks to each other or the anchors. This will ultimately allow us to leverage global relationships for learning our matching probabilities, as opposed to using greedy local searches on residuals for matching [2, 7]. We apply an MLP to get our overall affinity matrix  $A$ :

$$A = \sigma_{aff}(R). \quad (4)$$

### C. Log Affinity Loss

Inspired by [28–30], we train our model with the log affinity loss  $\mathcal{L}_{la}$ . We define the ground-truth affinity matrix  $A_{gt}$ , the estimated affinity matrix  $A$ , and the Hadamard product  $\odot$  to get:

$$\mathcal{L}_{la} = \frac{\sum_i \sum_j (A_{gt} \odot - \log(A))}{\sum_i \sum_j A_{gt}}. \quad (5)$$

The affinity matrix prediction has corresponding ground-truth matrices  $A_{gt, fm}$  and  $A_{gt, bm}$  for forward matching and backward matching, respectively. In total, our overall loss function  $\mathcal{L}$  is defined as follows:

$$\mathcal{L}_{fm} = \frac{\sum_i \sum_j (A_{gt, fm} \odot - \log(A_{fm}))}{\sum_i \sum_j A_{gt, fm}} \quad (6)$$

$$\mathcal{L}_{bm} = \frac{\sum_i \sum_j (A_{gt, bm} \odot - \log(A_{bm}))}{\sum_i \sum_j A_{gt, bm}} \quad (7)$$

$$\mathcal{L} = \frac{1}{2} (\mathcal{L}_{fm} + \mathcal{L}_{bm}). \quad (8)$$

### D. Sequential Track Confidence Refinement

We propose a sequential track confidence refinement technique that is used to assess the quality of tracks in real-time. We first offer an intuition behind our sequential track confidence refinement technique before formalizing our approach in Equation 9.

As stated earlier, track confidence should accurately reflect a track's quality relative to the quality of tracks for objects of the same class in a given time step. Thus, to obtain more accurate confidence scores to characterize our tracks, we want to leverage our affinity matrix estimation as follows: (1) If the affinity matrix indicates that the detection matched to our track has a high probability of being true-positive, we want to take a weighted average between the confidence of our matched detection and the existing confidence of our track. (2) If the affinity matrix indicates that the detection matched to our track almost got eliminated as an FP – only marginally missing the elimination threshold  $\tau_{FP}$  – then we want to downscale the existing confidence for our track, because it likely should not be kept alive with this matched detection.

This approach is captured in Equation 9.

$$\begin{aligned} c_{trk,i}^{(t)} &\leftarrow \mathbb{1}[P_{FP,i}^{(t)} < \beta_1] \beta_2 c_{det,i}^{(t)} + (1 - \beta_2) c_{trk,i}^{(t-1)} \quad (9) \\ \text{s.t. } 0 &\leq c_{det,i}^{(t)} \leq 1 \quad \forall i, t \\ 0 &\leq c_{trk,i}^{(t)} \leq 1 \quad \forall i, t \\ 0 &\leq \beta_1 \leq 1 \\ 0 &\leq \beta_2 \leq 1 \end{aligned}$$

According to Equation 9, for track ID  $i$  at time step  $t$ , we have a confidence of  $c_{trk,i}^{(t)}$ . The affinity matrix gets leveraged in the indicator function with  $P_{FP,i}^{(t)}$ , which is the probability that the detection got matched with the false-positive anchor at time step  $t$ . We fix  $\beta_1$  to be a value slightly less than  $\tau_{FP}$  to account for detections that are very close to getting eliminated as FP. In other words, we view detections with  $P_{FP,i}^{(t)} \in [\beta_1, \tau_{FP}]$  to be very uncertain detections, and we take this into account by reducing the overall track confidence for tracks matching to such detections. However, if the matched detection has a very high probability of being true-positive, then the track confidence  $c_{trk,i}^{(t)}$  becomes a weighted average between  $c_{det,i}^{(t)}$  and  $c_{trk,i}^{(t-1)}$ . Based on cross validation, we found that the values for  $\beta_1$  and  $\beta_2$  are not sensitive. We set  $\beta_1 = 0.5$  for all object classes, and  $\beta_2 = 0.5$  for all object classes except for bicycle ( $\beta_2 = 0.4$ ), bus ( $\beta_2 = 0.7$ ), and trailer ( $\beta_2 = 0.4$ ). For bicycle and trailer, we choose a slightly lower  $\beta_2$  to give more weight to the tracking history since 3D detectors tend to skew towards lower detection confidence scores for these more difficult classes. For the converse reason, we give bus a higher  $\beta_2$ .

In the special case of a newborn track, the track confidence is only the first term of Equation 9, i.e.

$$c_{trk,i}^{(t)} \leftarrow \mathbb{1}[P_{FP,i}^{(t)} < \beta_1] \beta_2 c_{det,i}^{(t)}.$$

This works well, because the detection confidences are also created in relative terms. Thus, if all newly initialized track confidences are the detection confidences scaled with  $\beta_2$ , their relative confidence rankings will be preserved.

#### IV. DATA PREPARATION

We evaluate our tracking performance on nuScenes [15], which is the primary benchmark in the 3D MOT domain mainly due to its large dataset, diverse range of driving scenarios from North America and Asia, and fine-grained evaluation. Specifically, nuScenes contains 1000 20-second long scenes generated with a 2Hz frame-rate. For the tracking task, nuScenes evaluates 7 classes: bicycle, bus, car, motorcycle, pedestrian, trailer, and truck.

We provide an overview of our LiDAR preprocessing and ground-truth affinity matrix formation used to deploy ShaSTA for the nuScenes dataset.

##### A. LiDAR Point Cloud Preprocessing

Since CenterPoint [7] provides detections at 2Hz, but the LiDAR has a sampling rate of 20Hz, we use the nuScenes [15] feature for accumulating multiple LiDAR sweeps with motion compensation. This provides us with a denser 4D point cloud

with an added temporal dimension, and makes the LiDAR input match the detection sampling rate. We use 10 LiDAR sweeps to generate the point cloud.

##### B. Ground-Truth Affinity Matrix.

We find the true-positive, false-positive, and false-negative detections using the matching algorithm from the nuScenes [15] dev-kit. For each frame, we define a ground-truth affinity matrix  $A_t \in \mathbb{R}^{(N_{max}+2) \times (N_{max}+2)}$  between the current frame and previous frame off-the-shelf detections. For all true-positive previous and current frame detections,  $B_{t-1}^{tp}$  and  $B_t^{tp}$ , we say they are matched if the ground-truth boxes they each correspond to have the same ground-truth tracking IDs. If there is a match between true-positive boxes  $b_{t-1,i}^{tp} \in B_{t-1}^{tp}$  and  $b_{t,j}^{tp} \in B_t^{tp}$ , then  $A_t(i, j) = 1$ . For all false-positive previous frame detections and true-positive previous frame detections whose tracking IDs do not appear in any of the current frame ground-truth tracking IDs, we call them dead tracks  $B_{t-1}^{dt}$ . For all  $b_{t-1,i}^{dt} \in B_{t-1}^{dt}$ , we set  $A_t(i, N_{max} + 1) = 1$ . For all true-positive previous frame detections that do not have a match in the current frame but whose tracking IDs do appear in the current frame ground-truth tracking IDs, we call them false-negatives  $B_{t-1}^{fn}$ . For all  $b_{t-1,i}^{fn} \in B_{t-1}^{fn}$ , we set  $A_t(i, N_{max} + 2) = 1$ . Additionally, for all true-positive current frame detections whose tracking IDs do not appear in the previous frame ground-truth tracking IDs, we call them newborn tracks  $B_t^{nb}$ , and we set  $A_t(N_{max} + 1, j) = 1$  for all  $b_{t,j}^{nb} \in B_t^{nb}$ . Finally, we define the set of false-positive current frame detections as  $B_t^{fp}$ , and we set  $A_t(N_{max} + 2, j) = 1$  for all  $b_{t,j}^{fp} \in B_t^{fp}$ . Unmatched  $(i, j)$  entries are set to 0.

#### V. EXPERIMENTAL RESULTS

This section provides an overview of evaluation metrics, training details, comparisons against state-of-the-art 3D multi-object tracking techniques, and ablation studies to analyze our technique.

##### A. Evaluation Metrics

The primary metrics for evaluating nuScenes are Average Multi-Object Tracking Accuracy (AMOTA) and Average Multi-Object Tracking Precision (AMOTP) [1]. AMOTA measures the ability to track objects correctly, while AMOTP measures the quality of the estimated tracks. We also present secondary metrics, including the total number of true-positives (TP), FPs, and FNs to better analyze tracking effectiveness for downstream decision-making tasks. We direct readers to [1] for a more thorough definition of the metrics.

##### B. Training Specifications

We train a different network for each object category following [3, 4]. The value of  $N_{max}$  depends on the object type since some classes are more common than others. We choose the value for  $N_{max}$  based on the per-object detection frequency we gather in the training set. The range is  $N_{max} \in [20, 90]$ . We use the Adam optimizer with a constant learning rate of  $10^{-4}$  and L2 weight regularization of  $10^{-2}$ .

**IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.**

TABLE I: nuScenes Test Results with evaluation in terms of overall and individual AMOTA and AMOTP. Our method achieves the highest AMOTA, while effectively balancing AMOTP to get the second lowest AMOTP only 0.005m shy of the first spot. We only compare on LiDAR-only methods that use CenterPoint detections to ensure a fair comparison. The blue entry is our LiDAR-only method ShaSTA.

Metric	Method	Detector	Input	Overall	Bicycle	Bus	Car	Motorcycle	Pedestrian	Trailer	Truck
AMOTA $\uparrow$	ShaSTA (Ours)	CenterPoint	LiDAR	<b>69.6</b>	41.0	<b>73.3</b>	<b>83.8</b>	<b>72.7</b>	<b>81.0</b>	<b>70.4</b>	<b>65.0</b>
	NEBP [6]	CenterPoint	LiDAR	68.3	<b>44.7</b>	70.8	83.5	69.8	80.2	69.0	59.8
	OGR3MOT [8]	CenterPoint	LiDAR	65.6	38.0	71.1	81.6	64.0	78.7	67.1	59.0
	CenterPoint [7]	CenterPoint	LiDAR	65.0	33.1	71.5	81.8	58.7	78.0	<b>69.3</b>	62.5
AMOTP $\downarrow$	ShaSTA (Ours)	CenterPoint	LiDAR	0.540	0.674	<b>0.629</b>	<b>0.383</b>	<b>0.504</b>	<b>0.369</b>	0.742	0.650
	NEBP [6]	CenterPoint	LiDAR	0.624	1.026	0.677	0.391	0.557	0.393	0.781	0.541
	OGR3MOT [8]	CenterPoint	LiDAR	0.620	0.899	0.675	0.395	0.615	0.383	0.790	0.585
	CenterPoint [7]	CenterPoint	LiDAR	<b>0.535</b>	<b>0.561</b>	0.636	0.391	0.519	0.409	<b>0.729</b>	<b>0.5</b>

TABLE II: nuScenes Test Results with evaluation in terms of overall TP, FP, and FN. We only compare on LiDAR-only methods that use CenterPoint detections to ensure a fair comparison. The blue entry is our LiDAR-only method ShaSTA.

Method	Detector	Input	TP $\uparrow$	FP $\downarrow$	FN $\downarrow$
ShaSTA (Ours)	CenterPoint	LiDAR	<b>97,799</b>	<b>16,746</b>	<b>21,293</b>
NEBP [6]	CenterPoint	LiDAR	97,367	16,773	21,971
OGR3MOT [8]	CenterPoint	LiDAR	95,264	17,877	24,013
CenterPoint [7]	CenterPoint	LiDAR	94,324	17,355	24,557

We use the pre-trained VoxelNet weights from CenterPoint [7] and freeze them for our LiDAR backbone so that the shape information correlates with our detection bounding boxes from [7]. Additionally, we downsample the number of FP detections during training to manage the class imbalance between FPs and TPs. We do so in order to prevent the model from overfitting to FPs by matching all the detections to the FP anchor to minimize the training loss. Since we have the labels for FP and TP detections for the training set, we can randomly remove the FPs from each frame to achieve the chosen per-frame ratio. Thus, we choose an FP to TP ratio during training, i.e. for every 1 TP keep 3 FPs per frame. Finally, we fix the thresholds  $\tau_{fp}$  for FP elimination,  $\tau_{fn}$  for FN propagation,  $\tau_{nb}$  for NB initialization, and  $\tau_{dt}$  for DT termination with cross-validation. We found that  $\tau_{fp} = 0.7$ ,  $\tau_{fn} = 0.5$ ,  $\tau_{nb} = 0.5$ , and  $\tau_{dt} = 0.5$  works best for all classes. Only the FP threshold is sensitive to tuning, since eliminating TP detections can contribute to catastrophic missing detections.

### C. Comparison with State-of-the-Art LiDAR Tracking

Our nuScenes test results can be seen in Tables I and II. Because previous works have shown that the standard AMOTA tracking metric heavily depends on the upstream 3D detections [3, 17, 18], all tracking methods reported use CenterPoint detections to ensure a fair comparison<sup>1</sup>.

The CenterPoint tracking algorithm [7] is our baseline, and we include recent, peer-reviewed works cited in Section II. Most notably, before ShaSTA was developed, NEBP [6] was the #1 LiDAR-only tracker using CenterPoint detections. The test results achieve a running speed of 10Hz (FPS in Hz) with 1 NVIDIA A100 GPU and batch size of 1.

In Table I, we can see that ShaSTA balances correct tracking (AMOTA) with precise, high-quality tracking (AMOTP). ShaSTA not only achieves the highest overall AMOTA, but also does so for 6 out of 7 object classes. Though

<sup>1</sup>CenterPoint detections currently have the highest mAP among publicly-available 3D detections. While the nuScenes leaderboard showcases tracking algorithms employing higher-scoring detections, these particular detections are not publicly available. It is crucial to note that comparing trackers utilizing different detectors directly is challenging due to the significant impact of detection quality on tracking performance. Only four LiDAR trackers rank higher than ShaSTA, and none of them utilize CenterPoint detections or make their detections publicly available. Additionally, some of these trackers lack code, research papers, or both.

NEBP achieves higher AMOTA for bicycles, our precision is much better. NEBP averages over 1m in distance from the ground-truth tracks on bicycles, which can be perilous in real-world traffic scenes. For the overall AMOTP, we only trail CenterPoint by 0.005m, which is a favorable trade-off when compared to our significant boost in overall AMOTA.

Table II shows that ShaSTA is effective in de-cluttering environments suffering from FPs and compensating for missed detections that lead to missing TPs and rising FNs. According to [1], AMOTA penalizes FPs and FNs. Concretely, even for a similar total number of TPs, FPs, and FNs, if ShaSTA predicts a greater confidence for TPs and a smaller confidence for FPs, the AMOTA metric will reward us more than others who have worse confidence estimates. Thus, our significant boost in AMOTA can be attributed to two factors: (1) we decrease the total number of FPs and FNs, and (2) we assign more accurate track confidence scores with our novel track confidence refinement technique.

### D. Ablation Studies

**Sequential Track Confidence Refinement.** Table III reports a significant boost in the overall AMOTA when using the proposed confidence refinement method. Smaller objects like bicycles and motorcycles that tend to suffer the most from LiDAR point cloud sparsity benefit the most from this technique. The pure detection confidences can diminish when there are scenarios such as occlusion, but our sequential refinement takes into account the existing confidence that has accumulated for a track before such scenarios occur. Of the large vehicles, trailer benefits the most since it is an underrepresented class with low-fidelity detection confidence scores.

**Shape Descriptor Extraction.** Table IV demonstrates that as we add more bounding box face centers for extracting shape descriptors, the accuracy of our model monotonically increases. This result supports our argument that leveraging raw LiDAR point clouds to get shape information is integral to the success of our tracking framework.

**Residuals.** Table V shows that even though we can achieve competitive tracking accuracy solely with spatio-temporal (VoxelNet or Bounding Box only) or shape (Shape only) information, we achieve the most optimal results when leveraging both types of information.

**Learned Affinity Matrix Augmentations.** According to Table VI FP elimination and NB initialization are the most impactful augmentations. Eliminating FPs is expected to boost our accuracy since off-the-shelf detections suffer from such high rates of FP detections. NB initialization can be seen as a reinforcement on eliminating FP detections because

## IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2024, Yokohama, Japan. Cite as RA-L paper.

TABLE III: Ablation on Confidence Refinement. We assess the effect of confidence refinement on tracking performance for each object. The version of ShaSTA without refinement indicates that we apply the CenterPoint detection confidence scores to our tracks. Evaluation is on nuScenes validation set in terms of overall AMOTA. The blue entry is our full method.

Metric	Method	Overall	Bicycle	Bus	Car	Motorcycle	Pedestrian	Trailer	Truck
AMOTA $\uparrow$	Without Refinement	70.2	49.4	85.2	85.3	70.1	80.9	50.1	70.1
	With Refinement	<b>72.8</b>	<b>58.8</b>	<b>85.6</b>	<b>85.6</b>	<b>74.5</b>	<b>81.4</b>	<b>53.1</b>	<b>70.3</b>
	<i>Difference</i>	+2.6	+9.4	+0.4	+0.3	+4.4	+0.5	+3.0	+0.2

TABLE IV: Ablation on Shape Descriptors. We assess the effect of the number of 3D bounding box points used to extract shape descriptors. Evaluation is on nuScenes validation set in terms of overall AMOTA. The blue entry is our full method.

Box Center Used	Face Centers Used	AMOTA $\uparrow$
$\checkmark$	$\times$	68.8
$\times$	$\checkmark$	70.2
$\checkmark$	$\checkmark$	<b>72.8</b>

TABLE V: Ablation on Residuals. We assess the effect of isolating each residual on the tracking algorithm's performance. Evaluation is on nuScenes validation set in terms of overall AMOTA. The blue entry is our full method.

Residual Used	AMOTA $\uparrow$
VoxelNet Only	68.9
Bounding Box Only	69.1
Shape Only	69.5
All (ShaSTA)	<b>72.8</b>

TABLE VI: Ablation on Learned Affinity Matrix. We assess the effect of each augmentation from our learned affinity matrix on track formation. Evaluation is on nuScenes validation set in terms of overall AMOTA. The blue entry is our full method.

Affinity Augmentation Method	AMOTA $\uparrow$
False-Positive Elimination Only	72.7
False-Negative Propagation Only	70.4
Newborn Initialization Only	72.6
Dead Track Termination Only	70.4
All (ShaSTA)	<b>72.8</b>

unmatched detections that do not have a probability of at least  $\tau_{nb}$  for being NBs are discarded.

## VI. CONCLUSION

We present ShaSTA, a 3D multi-object tracking framework that leverages shape and spatio-temporal information from LiDAR sensors to learn affinities for robust data association, track lifecycle management, FP elimination, FN propagation, and track confidence refinement. Our approach effectively addresses false-positives and missing detections against cluttered scenes and occlusion, yielding state-of-the-art performance. Since ShaSTA offers a flexible framework, we can extend it to fuse various sensor modalities or use our learned affinities in prediction and planning.

## REFERENCES

- X. Weng, J. Wang, D. Held, and K. Kitani, "3D Multi-Object Tracking: A Baseline and New Evaluation Metrics," *IROS*, 2020.
- H.-k. Chiu, A. Prioletti, J. Li, and J. Bohg, "Probabilistic 3d multi-object tracking for autonomous driving," *arXiv preprint arXiv:2001.05673*, 2020.
- H.-k. Chiu, J. Li, R. Ambruş, and J. Bohg, "Probabilistic 3d multi-modal, multi-object tracking for autonomous driving," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 14 227–14 233.
- C. Stearns, D. Rempe, J. Li, R. Ambruş, S. Zakhharov, V. Guizilini, Y. Yang, and L. J. Guibas, "Spot: Spatiotemporal modeling for 3d object tracking," in *European Conference on Computer Vision*. Springer, 2022, pp. 639–656.
- F. Meyer, T. Kropfreiter, J. L. Williams, R. Lau, F. Hlawatsch, P. Braca, and M. Z. Win, "Message passing algorithms for scalable multitarget tracking," *Proceedings of the IEEE*, vol. 106, no. 2, pp. 221–259, 2018.
- M. Liang and F. Meyer, "Neural enhanced belief propagation for data association in multiobject tracking," in *2022 25th International Conference on Information Fusion (FUSION)*. IEEE, 2022, pp. 1–7.
- T. Yin, X. Zhou, and P. Krahenbuhl, "Center-based 3d object detection and tracking," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 11 784–11 793.
- J.-N. Zaech, A. Liniger, D. Dai, M. Danelljan, and L. Van Gool, "Learnable online graph representations for 3d multi-object tracking," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 5103–5110, 2022.
- A. Kim, A. Oşep, and L. Leal-Taixé, "Eagermot: 3d multi-object tracking via sensor fusion," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 11 315–11 321.
- X. Wang, C. Fu, Z. Li, Y. Lai, and J. He, "Deepfusionmot: A 3d multi-object tracking framework based on camera-lidar fusion with deep association," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 8260–8267, 2022.
- X. Weng, Y. Wang, Y. Man, and K. Kitani, "Gnn3dmot: Graph neural network for 3d multi-object tracking with multi-feature learning," in *CVPR*, 2020.
- M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman, "The pascal visual object classes (voc) challenge," *International journal of computer vision*, vol. 88, no. 2, pp. 303–338, 2010.
- Y. Zhou and O. Tuzel, "Voxelnet: End-to-end learning for point cloud based 3d object detection," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 4490–4499.
- S. Chen, X. Wang, T. Cheng, Q. Zhang, C. Huang, and W. Liu, "Polar parametrization for vision-based surround-view 3d detection," *arXiv preprint arXiv:2206.10965*, 2022.
- H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nuscenes: A multimodal dataset for autonomous driving," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 11 621–11 631.
- M. Liang, B. Yang, W. Zeng, Y. Chen, R. Hu, S. Casas, and R. Urtasun, "Pnpnet: End-to-end perception and prediction with tracking in the loop," in *CVPR*, 2020.
- J. Luiten, A. Osep, P. Dendorfer, P. Torr, A. Geiger, L. Leal-Taixé, and B. Leibe, "Hota: A higher order metric for evaluating multi-object tracking," *International journal of computer vision*, vol. 129, pp. 548–578, 2021.
- L. Wang, X. Zhang, W. Qin, X. Li, J. Gao, L. Yang, Z. Li, J. Li, L. Zhu, H. Wang, et al., "Camo-mot: Combined appearance-motion optimization for 3d multi-object tracking with camera-lidar fusion," *IEEE Transactions on Intelligent Transportation Systems*, 2023.
- R. Girshick, "Fast r-cnn," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 1440–1448.
- X. Zhou, D. Wang, and P. Krähenbühl, "Objects as points," in *arXiv preprint arXiv:1904.07850*, 2019.
- C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "Pointnet: Deep learning on point sets for 3d classification and segmentation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 652–660.
- M. Engelcke, D. Rao, D. Z. Wang, C. H. Tong, and I. Posner, "Vote3deep: Fast object detection in 3d point clouds using efficient convolutional neural networks," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 1355–1361.
- B. Yang, W. Luo, and R. Urtasun, "Pixor: Real-time 3d object detection from point clouds," in *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2018, pp. 7652–7660.
- A. H. Lang, S. Vora, H. Caesar, L. Zhou, J. Yang, and O. Beijbom, "Pointpillars: Fast encoders for object detection from point clouds," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 12 697–12 705.
- H. Choi, M. Kang, Y. Kwon, and S.-e. Yoon, "An objectness score for accurate and fast detection during navigation," *arXiv preprint arXiv:1909.05626*, 2019.
- C. R. Qi, O. Litany, K. He, and L. J. Guibas, "Deep hough voting for 3d object detection in point clouds," in *proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 9277–9286.
- A. Simonelli, S. R. Buló, L. Porzi, M. López-Antequera, and P. Kotschieder, "Disentangling monocular 3d object detection," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 1991–1999.
- X. Guo and Y. Yuan, "Joint class-affinity loss correction for robust medical image segmentation with noisy labels," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2022, pp. 588–598.
- M. Hayat, S. Khan, S. W. Zamir, J. Shen, and L. Shao, "Gaussian affinity for max-margin class imbalanced learning," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 6469–6479.
- S. Sun, N. Akhtar, H. Song, A. Mian, and M. Shah, "Deep affinity network for multiple object tracking," *IEEE transactions on pattern analysis and machine intelligence*, vol. 43, no. 1, pp. 104–119, 2019.