

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

MUTR3D: A Multi-camera Tracking Framework via 3D-to-2D Queries

Tianyuan Zhang Carnegie Mellon University tianyuaz@andrew.cmu.edu

yuewang@csail.mit.edu

Yue Wang Massachusetts Institute of Technology

Yilun Wang Li Auto yilunw@cs.stanford.edu

Xuanyao Chen

Fudan University

xuanyaochen19@fudan.edu.cn

Hang Zhao ¶ Tsinghua University hangzhao@mail.tsinghua.edu.cn

Abstract

Accurate and consistent 3D tracking from multiple cameras is a key component in a vision-based autonomous driving system. It involves modeling 3D dynamic objects in complex scenes across multiple cameras. This problem is inherently challenging due to depth estimation, visual occlusions, appearance ambiguity, etc. Moreover, objects are not consistently associated across time and cameras. To address that, we propose an end-to-end MUlti-camera **TR**acking framework called MUTR3D. In contrast to prior works, MUTR3D does not explicitly rely on the spatial and appearance similarity of objects. Instead, our method introduces 3D track query to model spatial and appearance coherent track for each object that appears in multiple cameras and multiple frames. We use camera transformations to link 3D trackers with their observations in 2D images. Each tracker is further refined according to the features that are obtained from camera images. MUTR3D uses a set-to-set loss to measure the difference between the predicted tracking results and the ground truths. Therefore, it does not require any post-processing such as non-maximum suppression and/or bounding box association. MUTR3D outperforms state-of-the-art methods by 5.3 AMOTA on the nuScenes dataset. Code will be released.

1. Introduction

3D tracking is crucial in various perception systems, such as autonomous driving, robotics, and virtual reality. In its most basic incarnation, 3D tracking involves predicting per-frame objects and finding the correspondences between them temporally. Given per-frame object detection results, this problem boils down to associating objects across frames in a coherent fashion according to object similarity. On the other hand, tracking improves detection stability and enforces consistency of detection predictions across frames. However, this induces a complicated iterative optimization problem.

More challenges arise when detailing multi-camera cases. First, accurate 3D detection is necessary for accurate tracking. However, camera-based 3D object detection remains an unsolved problem. Second, vision trackers are fragile regarding occlusion and appearance ambiguity in complex scenes. For example, a person of interest may walk behind a car and re-appear after a couple of seconds in a different pose. Third, trackers often lose objects moving across camera view boundaries. Therefore, beyond temporal association, we need to perform cross-camera association when objects span or cross different cameras to make spatially consistent predictions. These challenges hamper the practical use of 3D vision trackers.

There are only a handful of works on vision-based 3D object tracking. Classical Kalman filtering-based methods [38] take detection results from any detectors as input and further make object state estimation and associations across time. More recent learning-based methods also follow a detect-to-track paradigm, where they first perform object proposals for each frame and then associate them in the feature space with a deep neural network [8, 11, 40].

In this work, we propose MUTR3D, an online multicamera 3D multi-object tracking framework that associate



Figure 1. We propose an end-to-end Multi-camera 3D tracking framework, named MUTR3D. Our algorithm works with arbitary camera rigs with known parameters. It handles multi-camera 3D detection, and cross-camera, cross-frame objects association end-to-end fashion.

objects into 3D tracks using spatial and appearance similarities in an end-to-end manner. More concretely, we introduce 3D track query, which directly models the 3D states and appearance features of an object track over time and across cameras. At each frame, a 3D track query sample features from all visible cameras, and learn to create/track/end a track. In contrast to previous works, MUTR3D performs detection and tracking simultaneously in a unified and end-to-end framework. Objects decoded from the same queries across frames are inherently associated.

In summary, our contributions are three-fold:

- To the best of our knowledge, MUTR3D is the first fully end-to-end multi-camera 3D tracking framework. Unlike existing detect-to-track methods that use explicit tracking heuristics, our method implicitly models the position and appearance variances of object tracks. Furthermore, we simplify the 3D tracking pipeline by eliminating commonly used post-processing steps such as non-maximum suppression, bounding box association, and object re-identification (Re-ID).
- We introduce a *3D track query* which models the 3D states of the entire track of an object. 3D track query samples feature from all visible cameras and update the track frame-by-frame end-to-end.
- Our end-to-end 3D tracking method achieves stateof-the-art performance on NuScenes vision-only 3D tracking dataset with 27.0% AMOTA. More specifically, MUTR3D performs much better than previous

SOTA methods in the multi-camera setting with 12% less ID switch.

• We propose two metrics to evaluate motion models in the current 3D tracker: Average Tracking Velocity Error (ATVE) and Tracking Velocity Error (TVE). They measure the error in the estimated motion of tracked objects.

2. Related Work

2.1. 3D MOT in Autonomous driving

For autonomous cars, it is critical to track surrounding objects while estimating their position, orientation, size, and velocity. Due to recent advances on 3D detection [14, 28, 41, 43, 49], modern 3D MOT follows tracking-bydetection paradigm. These methods detect objects in the current frame and then associate them with previous tracklets. Weng et al. [38] benchmark a simple yet effective association methods. They predict the location of previous tracklets through Kalman filtering, then associate current detections using 3D IoU. Beyond IoU, several works used L2 distance [43] and generalized 3D IoU [21] to associate 3D box with pure location cues. Many works use more advanced association by adding learned motion and appearance features [2, 8, 9] or using graph neural networks [5, 39, 45]. Several works study how to improve life cycle management [3,21] by utilizing cues from detection scores. QD3DT current SOTA (State-of-The-Art) camera-based tracking algorithms learn an appearance matching feature through dense contrastive learning They use an LSTM-

based motion model to learn motion features and predict current locations. Finally, it combines visual features, motion cues, and depth-ordering for the association. Though with strong RGB appearance cues, performance of camerabased 3D MOT [2, 8, 11, 27, 32, 47] has been lagged behind compared to LiDAR-based. On the nuScenes 3D MOT challenge's public leaderboard, STOA camera-based methods achieve 21.7% AMOTA while STOA LiDAR-based methods reach 67.9% AMOTA. The problem of tracking through multiple distinct viewpoints also draws attention [30].

2.2. Camera-based 3D Detection

3D object detection have seen great advances in recent years. A stream of algorithms build upon 2D detection framework [28, 34, 42, 48]. To resolve the fundamental ambiguity of instance depth and scales, categorical canonical shapes [1,19], geometric relation graphs [33] and pretrained monocular depth [18,22] are used. Another stream of methods works with representations on 3D space or Birds-Eye-View. Pseudolidar [35,44] use pre-trained monocular depth models to lift pixels to 3D point clouds, then perform 3D detection using a LiDAR-based detector. Lift-Splat-Shot [23] makes the lifting process fully differentiable and joint trains the lifting modules with downstream tasks. Later CaDDN [24] and BEVDet [12] used similar represents for 3D detection. DETR3D [36] adopt an inverse projecting process and build query-based multi-camera 3D detectors. Compared to working on perspective image planes directly, one major advantage of working in 3D space is the ease of adopting arbitrary camera rigs and fusing multiple sensor features. Currently, there is no clear advantages on performance [22]. More comparisons are still yet under-explored.

2.3. Query based detection and tracking

A dominant type of modern detection and tracking approach is to reduce the task of detection to pixel-wise regression and classifications [13, 17, 25, 26, 31, 48], then perform tracking by associating detection boxes. Recently, DETR [7] successfully used query-based set prediction to achieve state-of-the-art detection results. Later Track-Former [20], MOTR and TransTrack [29, 46] extends this idea to online 2D MOT. Our work builds upon the framework of query-based tracking. We extend the framework to multi-camera 3D MOT with a motion model.

3. Methods

3.1. Query based Object tracking

We adopt query-based tracking for our algorithms. Query-based tracking is extended from query-based detection [7], where detect queries, a fixed-size set of embedding, are used to represent 2D object candidates. Track query extends the concept of the detect query to multi-frames, i.e., representing a whole tracklet across frames [20,37,46]. Specifically, we initialize a set of newborn queries at the beginning of each frame, then queries update themselves frame-by-frame in an auto-regressive way. A decoder head predicts one object candidate from each track query in each frame, and boxes decoded in different frames from the same track query are directly associated. With proper query life cycle management, query-based tracking can perform joint detection and track in an online fashion.

There are three key ingredients in our query-based multicamera 3D tracker. (1) A query-based object tracking loss assigns different regression targets for two different types of queries, newborn queries, and old queries. (2) A multicamera sparse attention uses 3D reference points to sample image features for each query. (3) A motion model estimates object dynamics and updates the query's reference point across frames. We illustrate the flow of our trackers in Figure 2.

3.2. End-to-end object tracking loss

We first explain the concept of *label assignment* in the context of query-based tracking. Our algorithms maintain a changing set of track queries across frames. At the current frame, we decode one object candidate from each query. Ideally, The decoded object candidates from the same query should represent the same object across frames, thus forming a whole tracklet. To train the query-based tracker, we need to assign one target ground truth object for each query in each frame, and the assigned ground-truth object acts as the regression target for the query. Specifically, label assignment is a mapping function between ground-truth objects and track queries. We typically pad the set of ground truth objects with \emptyset (no object) to the number of predicted object candidates to ensure the mapping is a one-to-one mapping. Suppose we have N decoded object candidates $\{\hat{y}_1, \dots, \hat{y}_N\}$ in current frame, label assignment can be denoted as a mapping $\pi \in \{1, 2, \dots, N\} \mapsto \{1, 2, \dots, N\}$. Then the training loss can be expressed as a sum of paired box loss:

$$\mathcal{L} = \sum_{i=1}^{N} \mathcal{L}_{\text{box}}(y_{\pi(i)}, \hat{y}_i), \qquad (1)$$

where $y_{\pi(i)}$ denotes the assigned target ground-truth object, and \mathcal{L}_{box} could be any bounding box loss. There are two types of queries for each frame, and they have different label assignment strategies. *Newborn queries* are a set of learned queries. They are input-agnostic and will be added to the set of queries at the beginning of each frame. Newborn quires are responsible for detecting newly appeared objects in the current frame. So we perform bipartite matching between object candidates from newborn queries with newly



Figure 2. Pipeline overview of our online multi-camera tracker. All small colored squares in the black dashed box represent track queries. Blue boxes represent newborn queries, a fixed-set of learnable queries added to the set of track queries at the beginning of each frame. Orange boxes denote old queries, which are active queries from previous frames. Track queries attend with multi-camera features to decode object candidates in the current frame. Then we filter out inactive queries. We also update reference points of active queries to compensate for object motions and ego-motion. Finally, the updated queries went to the following frames to track the same objects.

appeared ground truth objects as DETR [7]. Old queries are active queries from previous frames which successfully detected or tracked objects. Old queries are responsible for tracking previously appeared objects in the current frame. The assignment for old queries is fixed after the first time it successfully detected a ground truth object. It is assigned to track the same object if they are in the current frame; otherwise, \emptyset (no object).

The 3D box loss \mathcal{L}_{box} in equation 1 is defined as:

$$\mathcal{L}_{\text{box}}(y_{\pi(i)}, \hat{y}_j) = \begin{cases} L_{\text{cls}}(c_{\pi(i)}, \hat{c}_j) + \lambda L_{\text{reg}}(b_{\pi(i)}, \hat{b}_j) & y_{\pi(i)} \neq Q \\ L_{\text{cls}}(c_{\pi(i)}, \hat{c}_j) & y_{\pi(i)} = Q \end{cases}$$
(2)

We use L_1 loss for L_{reg} , and L_{cls} is the focal loss [17], and the 3D object $y_{\pi(i)}$ is parameterized using the class label $c_{\pi(i)}$, and bounding box parameters $b_{\pi(i)}$, details for the parameterization is in equation 6.

3.3. Multi-camera Track query decoding

Our transformer decoder head takes track queries and attends them with multi-camera image features, and the extracted query featurees would be used to decode object candidates. Our decoder has two types of attention modules: self-attention between queries and cross attention between queries and image features. For memory efficiency, we adopt a reference-point based attention from DETR3D [36] for cross attentions. For notation in this section, only 3D coordinates or their 2D projections are in bold, e.g., 3D coordinates of reference points, c_i , estimated velocities v_i .

Query initialization. We assign a 3D reference point c_i to each query when it is initialized, i.e., when it is introduced as a newborn query at a certain frame. The 3D reference point is decoded from its learnable embedding using a shared MLP (multi-layer perceptrons):

$$\boldsymbol{c}_i = \Phi^{\text{ret}}(q_i),\tag{3}$$

where q_i denotes the learnable query embedding, the 3D reference points would be updated auto-regressively through layers of transformer decoders and across frames. It aims to approximate the 3D location of an object candidate.

Query feature extraction. The cross attention works by projecting the reference point of each query to all the cameras and sampling point features. Suppose we have synchronized images from M cameras for each frame. We

extract pyramidal features for each image independently. We denote the set of pyramidal features as: $\mathcal{F}_1, \mathcal{F}_2, \mathcal{F}_3, \mathcal{F}_4$. Each item $\mathcal{F}_k = \{F_{k1}, \ldots, F_{kM}\}, F_{ki} \in \mathbb{R}^{H \times W \times C}$ corresponds to a level of features of the M images. We denote the provided camera projection matrices as $\mathcal{T} = \{T_1, \ldots, T_M\}, T_i \in \mathbb{R}^{3 \times 4}$. Specifically, the sampled point feature f_{c_i} is :

$$\boldsymbol{c}_{mi} = T_m(\boldsymbol{c}_i \oplus 1), w_i = \text{MLP}(q_i),$$

$$\boldsymbol{f}_{\boldsymbol{c}_i} = \sum_k^4 \sum_m^M F_{km}(\boldsymbol{c}_{mi}) \cdot \sigma(w_{kmi}),$$

(4)

where c_{mi} denotes the projected 2D coordinates on the image plane of camera m, $F_{km}(c_{mi})$ represents bilinear sampling from image features, and $\sigma()$ denotes sigmoid function, which is used to normalize the weighting factor.

Then we use the extracted feature to update the query and its reference point

$$q_i \leftarrow q_i + \text{MLP}(f_{\boldsymbol{c}_i} + \text{PE}(q_i)), \\ \boldsymbol{c}_i \leftarrow \boldsymbol{c}_i + \text{MLP}(f_{\boldsymbol{c}_i}),$$
(5)

where PE is learnable positional encoding, it is initialized with each query. After layers of transformer decoder, we use the final query feature to decode object candidate in the current frame.

3D Object Parametrization. We use two small FFNs to decode 3D box parameters and categorical labels. We parameterize the 3D box by additional ten dimensional parameters: coordinates of the box center in ego frame, $x_i \in \mathbb{R}^3$, size of the 3D box $s_i = (w_i, l_i, h_i) \in \mathbb{R}^3$, 2D velocity in ego frame $v_i = (v_i^x, v_i^y) \in \mathbb{R}^2$ and orientation $(\sin \theta_i, \cos \theta_i)$, where θ_i is the yaw-angle in ego frame. The coordinates of the box center is predicted by adding a residual to the reference point:

$$\boldsymbol{x}_i = \boldsymbol{c}_i + \mathrm{MLP}(q_i). \tag{6}$$

3.4. Query Life Management

To deal with disappearing objects in an online fashion, we need to remove inactive queries after each frame. We define the confidence score of each query as the classification score of their predicted box. We use two threshold parameters τ_{new} and τ_{old} for box scores and a time length, T to control the life management.

During inference, for newborn queries in each frame, if the score is lower than τ_{new} , we remove it. For old queries, if their scores have been lower than τ_{old} for successive Tframes, we remove it. We select $\tau_{\text{new}} = 0.4$, and $\tau_{\text{old}} = 0.35$ and T = 5 for nuScenes dataset after several trails.

During training, we view queries matched to \emptyset as inactive. For newborn queries in the current frame, if it is matched to \emptyset , we remove it. For old queries, we remove it if it has been matched to \emptyset for successive T times. Note that old queries that have been matched to \emptyset but have not been removed continue to update themselves through the transformer decoder.

3.5. Query Update and Motion model

After filtering out outdated (dead) queries, we update track queries, both their features and 3D reference points. The purpose of updating the 3D reference point is to model object dynamics and compensate for ego-motion. There are two commonly used motion models in 3D tracking, Kalman Filter, e.g., [21, 38], which uses observed position across frames to estimate unknown velocity, and predicted velocity from detectors, e.g., CenterTrack [43, 47]. We use velocity predicted from queries, which updates through frames and can aggregate multi-frame features. We use a small FFN to predict ego frame velocity. The predicted velocity is supervised with ground truth. Denote the ego pose of current frame and next frame as $R_t, R_{t+1} \in \mathbb{R}^{3 \times 3}, T_t, T_{t+1} \in \mathbb{R}^3$. Denote the time gap between these two frames as Δt . We update the reference point c_i of the *i*-th query using the predicted box velocity $\boldsymbol{v}_i = (v_i^x, v_i^y, 0) \in \mathbb{R}^3$:

$$\boldsymbol{c}_i \leftarrow R_{t+1}^{-1}(R_t(\boldsymbol{c}_i + \boldsymbol{v}_i \times \Delta t) + T_t - T_{t+1}).$$
(7)

To implicitly model multi-frame appearance variations, we update the track query using features from previous frames. Following MOTR [46], we maintain a fixed-size first-in-first-out queue for each of the active queries, named *memory bank*. After each frame, we apply an attention module for each query and its memory bank. The track query acts as the query for the attention module, and the corresponding memory bank act as a set of keys and values.

4. Experiments

4.1. Datasets

We use nuScenes [6] dataset for all of our experiments. It consists of 1000 real-world sequences, 700 sequences for training, 150 for validation, and 150 for the test. Each sequence has roughly 40 annotated keyframes. Keyframes are synchronized frames for each sensor with a sampling rate of 2 FPS. Each frame includes images from six cameras with a full 360-degree field of view. It provides 3D tracking annotations for 7 Object categories.

4.2. Evaluation Metrics

Average multi-object tracking accuracy (AMOTA) and average multi-object tracking precision (AMOTP) are the major metrics for nuScenes 3D tracking benchmark. AMOTA and AMOTP are computed by integrating MOTA(multi-object tracking accuracy) and MOTP(multiobject tracking precision) values over all recalls:

$$AMOTA = \frac{1}{L} \sum_{r \in \{\frac{1}{L}, \frac{2}{L}, \dots, 1\}} MOTA_r, \qquad (8)$$

$$MOTA_r = \max\left(0, 1 - \frac{FP_r + FN_r + IDS_r - (1 - r)GT}{rGT}\right)$$
(9)

where FP_r , FN_r and IDS_r represents the number of false positives, false negatives, and identity switches computed at the corresponding recall r. GT is the number of ground truth bounding boxes. AMOTA can be formulated as:

AMOTP =
$$\frac{1}{L} \sum_{r \in \{\frac{1}{L}, \frac{2}{L}, ..., 1\}} \frac{\sum_{i,t} d_{i,t}}{\text{TP}_r}$$
, (10)

where $d_{i,t}$ denotes the 2D birds-eye-view position error of matched track *i* at time *t*, and TP_r indicates the number of matches computed at the corresponding recall *r*.

We also report tracking metrics from CLEAR [4] and Li et al. [15] such as MOTA, MOTP, IDS. The confidence threshold for these metrics is selected by independently picking the threshold with the highest MOTA for each category.

4.3. Implementation Details

Feature extractor Following prior works [34] [36] ResNet-101 with deformable convolutions [10] and FPN [16] are used for image feature extractors. For ablation study, we replace the ResNet-101 with ResNet-50 for memory efficiency.

Training details We use 3D detection pre-trained models from DETR3D [36]. Then we replace the head and train our tracker with three frames video clips for 72 epochs.

Kalman filter baselines Kalman filter-based methods have been state-of-the-art trackers on LiDAR-based 3D tracking across datsets [21]. However, camera-based SOTA methods typically use learned appearance and motion features for matching. To further understand the field of camera-based 3D MOT, we provide two Kalman filter baselines with DETR3D [36] detector. (1) A basic version with no advanced design. The basic version improves over the public implementation of AB3DMOT [38]. To handle the failure of IoU(Intersection over Union) during association with low frame rate data, we enlarge the prediction boxes by 20% when computing 3D IoU. (2) We also provide an advanced version of Kalman filter baselines from SimpleTrack [21], which used 3D generalized IoU and two-stage associations. SimpleTrack obtained SOTA result on LiDAR-based MOT.

4.4. Compare with State-of-the-art

We compare our method with SOTA methods in Table 1. We outperform current SOTA methods for the camerabased tracker by a large margin. The gain in AMOTA from the current SOTA method QD3DT [11] is over 5.2 points on the validation set and 5.3 points on the test set. Our tracker operates in an end-to-end fashion, with no NMS and no association stages as in QD3DT [11].

We put the comparisons of two of our Kalman filter baselines in Table 2. We outperform the basic version of the Kalman filter. However, when compared with more tailored baselines from SimpleTrack [21], we only have slight gains on metrics like AMOTA, MOTA, MOTP.

4.5. Evaluating Motion Models

The motion model provides one of the primary cues for 3D Multi-object Tracking. The motion model aims to describe the moving patterns of tracklets. To evaluate the motion models of different tracking algorithms, we develop two metrics, Average Tracking Velocity Error (ATVE) and Tracking Velocity Error (TVE), following the idea of AMOTP and MOTP. ATVE can be computed as:

$$ATVE = \frac{1}{L} \sum_{r \in \{\frac{1}{L}, \frac{2}{L}, \dots, 1\}} \frac{\sum_{i,t} ||v_i - v_t||_2}{TP_r}, \quad (11)$$

where we traverse over all pairs of matched tracking predictions and ground truth and compute the L_2 error between the predicted velocity v_i and the ground-truth velocity v_t . Average Tracking Velocity Error is computed by averaging over all recalls r, and TP_r represents the number of matches in corresponding recall r. Like MOTP, Tracking Velocity Error is the average velocity error computed at the recall with the highest MOTA. We evaluate the evaluation of motion models in Table 3. Compared to the previous state-of-the-art camera tracker QD3DT [11], our velocity is more accurate. Compared to Kalman filtering-based motion models, our algorithm achieves better Tracking velocity Error.

4.6. Ablation study

We study two factors in the ablation study. First, we study the effect of dropping our motion model, i.e., do not update the 3D reference points at the end of each frame. We show the ablation results in Table 4. Removing our motion model degrades the performance in all metrics.

Second, we study the effect of the number of training frames. Our methods track objects in an auto-regressive way, and no teacher-forcing is applied. During training, gradients computed in latter frames will still propagate to compute graphs in previous frames. In the ablation study, we perform all the experiments using ResNet-50 backbones. We report the performance of training with 3,4,5 frames in

Table 1. Comparison with state-of-the-art methods on nuScenes dataset. For public camera-based 3D tracking, our algorithm achieves state-of-the-art results, outperforming QD3DT [11] by 0.052 in AMOTA on validation set and 0.053 on test split.

	Modality	AMOTA ↑	AMOTP \downarrow	RECALL \uparrow	MOTA \uparrow	$\mathrm{IDS}\downarrow$	#params
Validation Split							
CenterPoint [43]	LiDAR	0.665	0.567	69.9%	0.562	562	9M
SimpleTrack [21]	LiDAR	0.687	0.573	72.5%	0.592	519	9M
DEFT [8]	Camera	0.201	N/A	N/A	0.171	N/A	22M
QD3DT [11]	Camera	0.242	1.518	39.9%	0.218	5646	91M
Ours	Camera	0.294	1.498	42.7%	0.267	3822	56M
Test Split							
CenterTrack [47]	Camera	0.046	1.543	23.3%	0.043	3807	20M
DEFT [8]	Camera	0.177	1.564	33.8%	0.156	6901	22M
QD3DT [11]	Camera	0.217	1.550	37.5%	0.198	6856	91M
Ours	Camera	0.270	1.494	41.1%	0.245	6018	56M

Table 2. Comparison with Kalman Filter based methods on nuScenes validation split. We construct two kalman filter baselines using our pretrained detector DETR3D [36]. We compare them with out tracker.

	AMOTA ↑	AMOTP \downarrow	RECALL \uparrow	MOTA \uparrow	$\text{MOTP} \downarrow$	$\mathrm{IDS}\downarrow$
DETR3D [36] + KF	0.263	1.569	39.7%	0.260	0.952	4698
DETR3D + SimpleTrack [21]	0.293	1.307	41.8%	0.263	0.84	1695
Ours	0.294	1.498	42.7%	0.267	0.799	3822

Table 3. **Evaluate velocity estimation.** We report ATVE (Average Tracking Velocity error) and TVE(Tracking Velocity error) for on nuScenes validation split. Compared with kalman filter based motion models, our methods obtain better TVE.

	Modality	ATVE \downarrow	$TVE\downarrow$
CenterPoint [43]	LiDAR	0.572	0.298
QD3DT [11]	Camera	1.876	1.373
DETR3D + SimpleTrack	Camera	1.344	0.836
Ours	Camera	1.548	0.768

Table 4. **Ablation on motion models.** When removing motion models, the performance of our algorithm drops in all metrics.

	AMOTA	AMOTP	RECALL	MOTA	IDS
w/o Motion	0.215	1.598	35.8%	0.198	4100
w/ Motion	0.234	1.585	38.7%	0.22	3775

Table 5. Results showed increasing the number of training frames gradually improves the performance.

Table 5. Ablation on the number of training frames. Training our models with longer video clips is beneficial.

#frames	AMOTA	AMOTP	RECALL	IDS	ATVE
3	0.234	1.585	38.7%	3775	1.606
4	0.242	1.580	39.7%	4623	1.545
5	0.251	1.573	39.9%	3873	1.565

4.7. Qualitative results

We provide visualizations of our tracking algorithms in both BEV and camera views for an 8 seconds clip in Figure 3. Near-filed objects on the left/right side of the car are usually truncated by several cameras, which is a substantial challenge for multi-camera 3D tracking. See, the gray and black cars are truncated in the Front-Left camera and Back-Left camera(3-rd/7-th and 4-th/8-th row), and our algorithm handles them correctly.

5. Conclusion

We design an end-to-end multi-camera 3D MOT framework. Our framework can perform 3D detection, compensate for ego-motion and object motions, and perform crosscamera and cross-frame object association end-to-end. In the nuScenes test dataset, our tracker outperforms the current state-of-the-art camera-based 3D tracker QD3DT [11] by 5.3 AMOTA and 4.7 MOTA. We also study the quality of the motion models in current 3D trackers by evaluating two new metrics: Average Tracking Velocity Error (ATVE) and Tracking Velocity Error (TVE). Compared to hand-designed associating methods, we believe our end-toend learnable tracker can enjoy the abundant amount of data in autonomous driving fields in the future.



Figure 3. Visualization on 8 consecutive frames with FPS as 1. We plot the results for 4 views, from top to bottom: Birds-Eye-View, Front camera, Front-left camera, and Back-left camera. Objects with the same identity are painted with the same color. We plot the estimated velocity using arrows, and longer arrows represent larger velocity. The example we showed contains multiple frames with truncated objects across cameras. Our algorithm are designed to fuse multi-camera features automatically, and handle the truncation correctly.

References

- Ivan Barabanau, Alexey Artemov, Evgeny Burnaev, and Vyacheslav Murashkin. Monocular 3D Object Detection via Geometric Reasoning on Keypoint. arXiv preprint arXiv:1905.05618, 2019. 3
- [2] Erkan Baser, Venkateshwaran Balasubramanian, Prarthana Bhattacharyya, and Krzysztof Czarnecki. FANTrack: 3D Multi-Object Tracking with Feature Association Network. In *IEEE Intelligent Vehicles Symposium (IV)*, 2019. 2, 3
- [3] Nuri Benbarka, Jona Schröder, and Andreas Zell. Score refinement for confidence-based 3D multi-object tracking. In *IROS*, 2021. 2
- [4] Keni Bernardin, Alexander Elbs, and Rainer Stiefelhagen. Multiple Object Tracking Performance Metrics and Evaluation in a Smart Room Environment. In *ECCV Workshops*, 2006. 6
- [5] Guillem Brasó and Laura Leal-Taixé. Learning a Neural Solver for Multiple Object Tracking. In CVPR, 2020. 2
- [6] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuScenes: A multimodal dataset for autonomous driving. In CVPR, 2020. 5
- [7] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-toend Object Detection with Transformers. In *ECCV*, 2020. 3, 4
- [8] Mohamed Chaabane, Peter Zhang, J Ross Beveridge, and Stephen O'Hara. DEFT: Detection Embeddings for Tracking. *CVPR Workshops*, 2021. 1, 2, 3, 7
- [9] Hsu-kuang Chiu, Jie Li, Rareş Ambruş, and Jeannette Bohg. Probabilistic 3D Multi-Modal, Multi-Object Tracking for Autonomous Driving. In *ICRA*, 2021. 2
- [10] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable Convolutional Networks. In *ICCV*, 2017. 6
- [11] Hou-Ning Hu, Yung-Hsu Yang, Tobias Fischer, Trevor Darrell, Fisher Yu, and Min Sun. Monocular Quasi-Dense 3D Object Tracking. *arXiv preprint arXiv:2103.07351*, 2021. 1, 3, 6, 7
- [12] Junjie Huang, Guan Huang, Zheng Zhu, and Dalong Du. BEVDet: High-performance Multi-camera 3D Object Detection in Bird-Eye-View. arXiv preprint arXiv:2112.11790, 2021. 3
- [13] Lichao Huang, Yi Yang, Yafeng Deng, and Yinan Yu. Dense-Box: Unifying Landmark Localization with End to End Object Detection. arXiv preprint arXiv:1509.04874, 2015. 3
- [14] Alex H Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. PointPillars: Fast Encoders for Object Detection from Point Clouds. In CVPR, 2019. 2
- [15] Yuan Li, Chang Huang, and Ram Nevatia. Learning to associate: HybridBoosted multi-target tracker for crowded scene. In *CVPR*, 2009. 6
- [16] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature Pyramid Networks for Object Detection. In *CVPR*, 2017. 6

- [17] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal Loss for Dense Object Detection. In *ICCV*, 2017. 3, 4
- [18] Xinzhu Ma, Shinan Liu, Zhiyi Xia, Hongwen Zhang, Xingyu Zeng, and Wanli Ouyang. Rethinking Pseudo-LiDAR Representation. In *ECCV*, 2020. 3
- [19] Fabian Manhardt, Wadim Kehl, and Adrien Gaidon. ROI-10D: Monocular Lifting of 2D Detection to 6D Pose and Metric Shape. In *CVPR*, 2019. 3
- [20] Tim Meinhardt, Alexander Kirillov, Laura Leal-Taixe, and Christoph Feichtenhofer. TrackFormer: Multi-Object Tracking with Transformers. arXiv preprint arXiv:2101.02702, 2021. 3
- [21] Ziqi Pang, Zhichao Li, and Naiyan Wang. SimpleTrack: Understanding and Rethinking 3D Multi-object Tracking. arXiv preprint arXiv:2111.09621, 2021. 2, 5, 6, 7
- [22] Dennis Park, Rares Ambrus, Vitor Guizilini, Jie Li, and Adrien Gaidon. Is Pseudo-Lidar needed for Monocular 3D Object detection? In *ICCV*, 2021. 3
- [23] Jonah Philion and Sanja Fidler. Lift, Splat, Shoot: Encoding Images From Arbitrary Camera Rigs by Implicitly Unprojecting to 3D. In ECCV, 2020. 3
- [24] Cody Reading, Ali Harakeh, Julia Chae, and Steven L Waslander. Categorical Depth Distribution Network for Monocular 3D Object Detection. In CVPR, 2021. 3
- [25] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You Only Look Once: Unified, Real-Time Object Detection. In *CVPR*, 2016. 3
- [26] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. *NeurIPS*, 2015. 3
- [27] Samuel Scheidegger, Joachim Benjaminsson, Emil Rosenberg, Amrit Krishnan, and Karl Granström. Mono-Camera 3D Multi-Object Tracking Using Deep Learning Detections and PMBM Filtering. In *IEEE Intelligent Vehicles Symposium (IV)*, 2018. 3
- [28] Andrea Simonelli, Samuel Rota Bulo, Lorenzo Porzi, Manuel López-Antequera, and Peter Kontschieder. Disentangling monocular 3d object detection. In *ICCV*, 2019. 2, 3
- [29] Peize Sun, Jinkun Cao, Yi Jiang, Rufeng Zhang, Enze Xie, Zehuan Yuan, Changhu Wang, and Ping Luo. Transtrack: Multiple object tracking with transformer. *arXiv preprint* arXiv:2012.15460, 2020. 3
- [30] Xiao Tan, Zhigang Wang, Minyue Jiang, Xipeng Yang, Jian Wang, Yuan Gao, Xiangbo Su, Xiaoqing Ye, Yuchen Yuan, Dongliang He, et al. Multi-camera vehicle tracking and re-identification based on visual and spatial-temporal features. In *CVPR Workshops*, 2019. 3
- [31] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. FCOS: Fully Convolutional One-Stage Object Detection. In *ICCV*, 2019. 3
- [32] Pavel Tokmakov, Jie Li, Wolfram Burgard, and Adrien Gaidon. Learning to Track with Object Permanence. In *ICCV*, 2021. 3
- [33] Tai Wang, ZHU Xinge, Jiangmiao Pang, and Dahua Lin. Probabilistic and Geometric Depth: Detecting Objects in Perspective. In *CoRL*, 2022. 3

- [34] Tai Wang, Xinge Zhu, Jiangmiao Pang, and Dahua Lin. FCOS3D: Fully Convolutional One-Stage Monocular 3D Object Detection. *ICCV Workshops*, 2021. 3, 6
- [35] Yan Wang, Wei-Lun Chao, Divyansh Garg, Bharath Hariharan, Mark Campbell, and Kilian Q Weinberger. Pseudo-LiDAR from Visual Depth Estimation: Bridging the Gap in 3D Object Detection for Autonomous Driving. In CVPR, 2019. 3
- [36] Yue Wang, Vitor Campagnolo Guizilini, Tianyuan Zhang, Yilun Wang, Hang Zhao, and Justin Solomon. DETR3D: 3D Object Detection from Multi-view Images via 3D-to-2D Queries. In *CoRL*, 2021. 3, 4, 6, 7
- [37] Yuqing Wang, Zhaoliang Xu, Xinlong Wang, Chunhua Shen, Baoshan Cheng, Hao Shen, and Huaxia Xia. End-to-End Video Instance Segmentation with Transformers. In *CVPR*, 2021. 3
- [38] Xinshuo Weng, Jianren Wang, David Held, and Kris Kitani. 3D Multi-Object Tracking: A Baseline and New Evaluation Metrics. *IROS*, 2020. 1, 2, 5, 6
- [39] Xinshuo Weng, Yongxin Wang, Yunze Man, and Kris M Kitani. GNN3DMOT: Graph Neural Network for 3D Multi-Object Tracking with Multi-Feature Learning. In CVPR, 2020. 2
- [40] Nicolai Wojke, Alex Bewley, and Dietrich Paulus. Simple Online and Realtime Tracking with a Deep Association Metric. In *ICIP*, 2017. 1
- [41] Yan Yan, Yuxing Mao, and Bo Li. SECOND: Sparsely Embedded Convolutional Detection. *Sensors*, 18(10), 2018. 2
- [42] Zetong Yang, Yanan Sun, Shu Liu, and Jiaya Jia. 3DSSD: Point-based 3D Single Stage Object Detector. In *CVPR*, 2020. 3
- [43] Tianwei Yin, Xingyi Zhou, and Philipp Krähenbühl. Centerbased 3D Object Detection and Tracking. arXiv preprint arXiv:2006.11275, 2020. 2, 5, 7
- [44] Yurong You, Yan Wang, Wei-Lun Chao, Divyansh Garg, Geoff Pleiss, Bharath Hariharan, Mark Campbell, and Kilian Q Weinberger. Pseudo-LiDAR++: Accurate Depth for 3D Object Detection in Autonomous Driving. arXiv preprint arXiv:1906.06310, 2019. 3
- [45] Jan-Nico Zaech, Alexander Liniger, Dengxin Dai, Martin Danelljan, and Luc Van Gool. Learnable Online Graph Representations for 3D Multi-Object Tracking. *IEEE Robotics* and Automation Letters, 2022. 2
- [46] Fangao Zeng, Bin Dong, Tiancai Wang, Cheng Chen, Xiangyu Zhang, and Yichen Wei. MOTR: End-to-End Multiple-Object Tracking with TRansformer. arXiv preprint arXiv:2105.03247, 2021. 3, 5
- [47] Xingyi Zhou, Vladlen Koltun, and Philipp Krähenbühl. Tracking Objects as Points. In ECCV, 2020. 3, 5, 7
- [48] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as Points. arXiv preprint arXiv:1904.07850, 2019. 3
- [49] Yin Zhou and Oncel Tuzel. VoxelNet: End-to-End Learning for Point Cloud Based 3D Object Detection. In *CVPR*, 2018.2