

3DMODT: Attention-Guided Affinities for Joint Detection & Tracking in 3D Point Clouds

Jyoti Kini¹, Ajmal Mian² and Mubarak Shah¹

Abstract—We propose a method for joint detection and tracking of multiple objects in 3D point clouds, a task conventionally treated as a two-step process comprising object detection followed by data association. Our method embeds both steps into a single end-to-end trainable network eliminating the dependency on external object detectors. Our model exploits temporal information employing multiple frames to detect objects and track them in a single network, thereby making it a utilitarian formulation for real-world scenarios. Computing affinity matrix by employing features similarity across consecutive point cloud scans forms an integral part of visual tracking. We propose an attention-based refinement module to refine the affinity matrix by suppressing erroneous correspondences. The module is designed to capture the global context in affinity matrix by employing self-attention within each affinity matrix and cross-attention across a pair of affinity matrices. Unlike competing approaches, our network does not require complex post-processing algorithms, and directly processes raw LiDAR frames to output tracking results. We demonstrate the effectiveness of our method on three tracking benchmarks: JRDB, Waymo, and KITTI. Experimental evaluations indicate the ability of our model to generalize well across datasets.

I. INTRODUCTION

Rapid advancements in 3D sensing technology and the recent commercial interest in autonomous driving [1], [2] and assistive robots [3], [4], have sparked great research interest in multiple object tracking (MOT) in 3D point clouds obtained from LiDAR (Light Detection and Ranging) sensors. MOT aims to predict the trajectories of distinct objects in a given sequence. In 2D tracking, two major paradigms: (a) tracking-by-detection [5], [6], [7], and (b) joint detection and tracking [8], [9], [10] have commonly been used to advance research in this direction. Although numerous 2D MOT approaches [11], [10], [9], [12], [13], [14], [15] have leveraged the benefits of jointly performing detection and tracking in a single unified framework, several open challenges have afflicted this long-standing research problem of MOT in 3D point clouds.

3D point cloud-based tracking provides robust depth perception, motion prediction, and path planning for critical robot navigation tasks. However, LiDAR point clouds are characterized by challenges such as data sparsity, uneven

*This work was conducted at UCF and supported by Lockheed Martin Corporate Engineering, Technology and Operations (CETO) University Engagement (UE) - Research. Professor Ajmal Mian is the recipient of an Australian Research Council Future Fellowship Award (project number FT210100268) funded by the Australian Government.

¹Center for Research in Computer Vision, University of Central Florida, USA jyoti.kini@knights.ucf.edu, shah@crvc.ucf.edu

²University of Western Australia ajmal.mian@uwa.edu.au

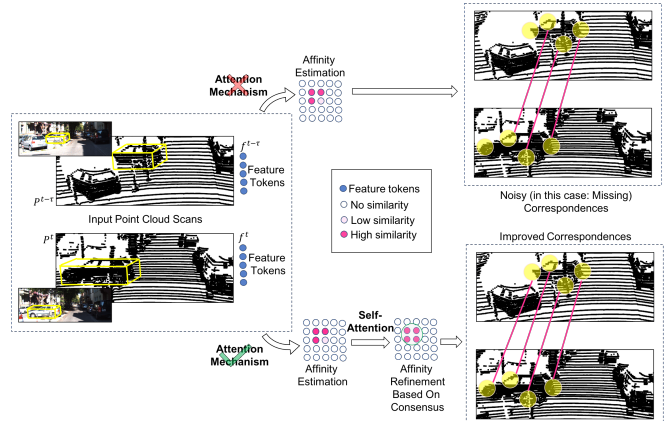


Fig. 1: Overview of our attention-guided affinity refinement. Feature tokens of two consecutive point cloud scans are used to construct an affinity matrix that stores dense similarities between the tokens. We use attention (bottom) to refine the affinities. Initially, the *light pink* node is linked incorrectly. However, self-attention on affinities seeks consensus from the other nodes (i.e. *dark pink* nodes). Supporting cues provided by the other nodes boost the correct correspondences and weaken the noisy correspondence. The input to our network is point cloud scans. RGB images and bounding boxes are shown for visualization purposes only.

sampling density, lack of texture, etc. To accomplish tracking across point cloud scans, typically object features are used to compute robust affinities for comprehensive data association (correspondences between the detected objects across multiple scans). Popular 3D MOT approaches [16], [17], [18] rely heavily on external off-the-shelf detectors to localize objects of interest. This not only adds an overhead to the tracking, but sub-par detections also inherently affect the overall performance of the tracker. Moreover, given the sparse and non-uniform nature of the point clouds, simple affinities estimation in current 3D joint detection and tracking approach [19] leads to noise in correspondences, which are crucial for robust tracking.

We propose an end-to-end learning framework for joint detection and tracking in 3D point clouds. As opposed to the conventional use of attention on features, we take a unique approach and propose self and cross-attention on the affinity matrices. Three consecutive point cloud frames are passed through a single end-to-end network to extract features employing transformer encoder, followed by the computation of pairwise affinity matrices. The affinity matrices are

refined (see Figure 1) using an attention mechanism and intermediate supervision. Thereafter, 3D tracking offsets are computed. The final output is supervised using ground truth object detections and is composed of regressed values of object centers and bounding boxes in 3D. The output object detection information and tracking offsets are collectively used to generate tracks. Experiments on the JRDB, Waymo and KITTI 3D tracking benchmarks show that our end-to-end network achieves competitive results w.r.t. existing MOT methods, *without* the need for an external object detector during inference or extensive post-processing of the tracking results. To summarize our contributions:

- We propose a novel 3D point cloud-based joint detection and tracking method that eliminates the reliance on external object detector.
- We develop an end-to-end network that directly processes raw point clouds to produce multiple object tracks, without the need to pre-process the point clouds to eliminate detection uncertainties or post-process the output results using extensive association algorithms.
- We design a novel attention-based refinement module to capture global context of affinities, which is essential to refining the affinity matrices capturing correspondences across the point cloud scans and thereby improving tracking accuracy.
- We show competitive performance on the KITTI, JRDB, and Waymo tracking benchmarks as well as the generalization ability of our method across datasets.

II. RELATED WORK

2D Multi-Object Tracking: Given the continued advancements in object detection techniques [20], [21], [22], most 2D multi-object tracking methods use the tracking-by-detection approach. This pipeline involves first detecting the objects and then associating the detections across frames. The association step can be categorized into two groups: appearance-based and motion-based. Appearance-based methods [23], [24] establish object correspondences across frames by using re-identification features. On the other hand, motion-based approaches use temporal modeling to update the object detections using distance or intersection-over-union (IOU) metric. State estimation using the Kalman filter [25] is a common practice amongst these approaches. Motion modeling [9], [26] by predicting motion offsets is another typical strategy. Despite the popularity of the tracking-by-detection paradigm, joint detection and multi-object tracking has gained momentum due to the missing connection between detections and associations in the former case.

Joint detection and tracking methods co-optimize both tasks, thereby boosting performance gains. FairMOT [27] and JDE [28] introduce a re-identification module into the detector, while GSDT [14] incorporates Graph Neural Networks to accomplish the task. Inspired by the success of Transformers [29], TrackFormer [30] and TransTrack [31] employ query-key mechanisms to follow object locations.

3D Multi-Object Tracking: Although 3D information is

crucial for depth perception, the data is inherently challenging to work with. Therefore, the research trend in 3D space ensues the progress in 2D space. Similar to its 2D counterpart, tracking-by-detection is a more commonly used practice in 3D multi-object tracking. The availability of sophisticated detectors has facilitated the growth in this space. VoxelNet [32], SECOND [33], PointPillar [34] are LiDAR-only detectors. In the tracking-by-detection approach, detections are followed by association. AB3DMOT [17] uses the Kalman filter for state motion estimation. PnPNet [35] learns affinity matrix using 3D features and trajectories. CenterPoint [36] incorporates a regression head to estimate velocity. Chiu et al. [37] enhances association by introducing distance metric. However, in LiDAR point cloud scans, object associations rely solely on motion modeling and lack appearance cues. Therefore, most tracking approaches, additionally, employ complex association algorithms (e.g. Hungarian algorithm) for bipartite matching. Some other detectors, namely ContFuse [38], MMF [39], CLOCS [40], 3D-CVF [41], fuse camera images with LiDAR point clouds to introduce appearance information. However, fusing camera data introduces sensor calibration overhead.

Similar to 2D space, 3D joint detection and tracking attempts to leverage the power of optimizing the two tasks in a single end-to-end network. JRMOT [5] combines re-identification, detection, and tracking in a single framework. Similar to our work, SimTrack [19] works solely with LiDAR data for joint detection and tracking by introducing a motion estimation branch to eliminate the heuristic matching step. In our work, however, we focus on refining the affinity matrix, which maps the correspondences across point cloud scans, using the attention-based refinement module.

Attention Mechanism: Over the past few years, the attention mechanism has gained accolades due to its focus on non-local operations [42]. Its variants (self-attention and cross-attention) have been used in the spatial and temporal domains for dynamic information aggregation [43]. We use the concept of the global receptive field provided by the attention mechanism to refine the affinity matrix. We suppress the erroneous correspondences in the affinity matrix by seeking consensus among the affinities using attention.

III. APPROACH

Figure 2 provides detailed schematics of the proposed 3DMODT, which has four main building blocks: (1) transformer encoder for feature extraction, (2) affinity computation and attention-based refinement, (3) tracking offset and 3D detection prediction, and (4) data association and tracklet-generation. We extract feature tokens $f^t, f^{t-\tau}$ from consecutive point cloud scans $P^t, P^{t-\tau}$. Thereafter, we use these feature tokens to construct affinity matrix A^t that stores dense similarity matches of tokens associated with t and $t - \tau$. The affinity matrix is then passed through attention mechanism leveraging global receptive fields to get refined affinity matrix \hat{A}^t . Network heads are employed to compute: (a) tracking offsets $O_{i,j,k}^t$ that store spatio-temporal displacements for all points, denoted by 3D locations (i, j, k), from




Fig. 2: **Details of the 3DMODT network.** 3DMODT performs joint detection and tracking in 3D point cloud data. Input point cloud scans $\mathcal{P} = \{P^t, P^{t-1}, P^{t-2}\}$ are passed through a transformer encoder $\phi(\cdot)$. The feature tokens are used to construct affinity matrices $\mathcal{A} = \{A^t, A^{t-1}\}$ that store dense similarities between tokens in consecutive scans. Attention is then applied to the affinity matrices for refinement. The network heads $\varphi(\cdot)$ generate (a) tracking offsets that store spatio-temporal displacements of points between consecutive scans, and (b) 3D detections. The predictions from the network heads are collectively used to generate tracklets. L_a, L_c, L_b represent Forward Association Loss on Affinities, ℓ_1 Center Offset Loss and ℓ_1 Dimension Size Loss. \otimes, \oplus, \odot denote matrix multiplication, element-wise sum and Hadamard product.

time t to the corresponding points at $t - \tau$, and (b) regress 3D object information (center, bounding box, rotation). 3D bounding box information combined with tracking offset is used to generate tracklets. Our model processes three point cloud scans simultaneously in a single pass to leverage self-attention and cross-attention across affinity matrices, thereby improving the temporal context. The four building blocks of our model are explained below.

A. Feature Extraction

The input to our network is three point cloud scans $\mathcal{P} = \{P^t, P^{t-1}, P^{t-2}\} \in \mathbb{R}^{N \times 3}$ captured at time step $t, t-1, t-2$ and consisting N 3D points. Each point cloud is passed through a transformer encoder (feature extractor) $\phi(\cdot)$ [44] to generate feature tokens $\hat{\mathcal{P}} = \{f^t, f^{t-1}, f^{t-2}\} \in \mathbb{R}^{M \times dim}$. Here, $M \subset N$ and dim corresponds to 64 feature representation channels + 3 channels of positional information.

B. Attention-Guided Affinity Refinement

Affinity Estimation: We process the feature tokens to construct affinity matrices. These affinity matrices store similarity scores between tokens at location index d for a given scan and tokens at location index e for its successive scan, with indexes $d = 1, \dots, M$ and $e = 1, \dots, M$. Given the feature tokens f_d^t (or simply f^t) and $f_e^{t-\tau}$ (or simply $f_e^{t-\tau}$), cosine similarity is applied to compute the affinity matrix. The affinity matrix $A_{d,e}^t$ (or simply A^t) across time t and $t - \tau$ is given by:

$$A_{d,e}^t = \frac{f_d^t \cdot f_e^{t-\tau}}{\|f_d^t\| \|f_e^{t-\tau}\|} \quad (1)$$

Our network computes two affinity matrices $\mathcal{A} = \{A^t, A^{t-1}\}$ across time steps $(t, t-1)$, and $(t-1, t-2)$

respectively. The computed affinity matrices \mathcal{A} exhibit noisy affinities due to inter-object similarities and intra-object variations caused by viewpoint changes and occlusions. Therefore, we apply attention mechanism to suppress the erroneous similarities by capturing global context. Typically, attention is applied to improve input feature representation. In our case, however, we take a novel approach and introduce attention on affinity matrices, as opposed to the input feature embeddings.

Affinity Refinement: Attention-guided affinity refinement is a two-step process. First, we apply self-attention on the individual affinity matrix A^t and A^{t-1} to get intermediate affinity matrix S^t and S^{t-1} respectively. This is followed by cross-attention applied on the intermediate affinity matrices $\mathcal{S} = \{S^t, S^{t-1}\}$ to yield the refined affinity matrices $\hat{\mathcal{A}} = \{\hat{A}^t, \hat{A}^{t-1}\}$.

Self-attention, applied on individual affinity matrix, seeks consensus from the similarities within the affinity matrix and suppresses erroneous correspondences. We employ convolutions followed by projection layer [45] to transform the affinity matrices into query Q , and set of key-value (K, V) pair vectors. Thereafter, scalar product computed between Q and K is passed through a softmax activation, and the resultant output is used to re-weight V . The self-attention on individual affinity matrix $\zeta(\cdot)$ is formulated as:

$$\zeta(A^t) = \text{softmax}(\bar{Q}^t \bar{K}^{t\top}) \bar{V}^t, \quad (2)$$

$$\zeta(A^{t-1}) = \text{softmax}(\bar{Q}^{t-1} \bar{K}^{t-1\top}) \bar{V}^{t-1}, \quad (3)$$

where $\bar{Q} = QW_q$, $\bar{K} = KW_k$, $\bar{V} = VW_v$ represent various linear transformations [45], and W_q, W_k, W_v denote transformation weights for query, key and value.

Cross-attention, on the other hand, is applied across affinity matrices. It aims to suppress erroneous correspondences in a given affinity matrix, by seeking agreement between similarities of the given affinity matrix and its successive affinity matrix. Applying cross-attention on affinity matrices implicitly increases the temporal context. The cross-attention across affinity matrices $\xi(\cdot)$ is given by:

$$\xi(A^t) = \text{softmax}(\bar{Q}^{t-1} \bar{K}^{t\top}) \bar{V}^t, \quad (4)$$

$$\xi(A^{t-1}) = \text{softmax}(\bar{Q}^t \bar{K}^{t-1\top}) \bar{V}^{t-1}. \quad (5)$$

C. Tracking Offset and 3D Bounding Box Prediction

The initial feature tokens \hat{P} and refined affinity matrices \hat{A} are processed through network heads $\varphi(\cdot)_T$ and $\varphi(\cdot)_D$ to generate tracking offsets and 3D detections respectively. We generate tracking offset $O_{i,j,k}^t \in \mathbb{R}^{N \times 3}$ that estimates the spatio-temporal displacement for point (i, j, k) at time t to its corresponding point at time $t - \tau$. In addition, several light-weight convolutions are used to predict 3D bounding boxes $\mathcal{B} = \{B_k\}$. Each bounding box $B = (c_x, c_y, c_z, b_w, b_l, b_h, \alpha)$ comprises object center (c_x, c_y, c_z) , object size (b_w, b_l, b_h) and rotation in yaw α . The values are regressed in logarithmic scale.

D. Tracklet Generation

To generate tracklets, we connect the 3D detection bounding boxes \mathcal{B} across time. For time step t , we first look for bounding box associations within a limited space in $t - 1$, as prompted by the tracking offsets. In case of unmatched bounding boxes, we increase the search space in $t - 1$. Basically, given time step t , we use the bounding box center location (c_x, c_y, c_z) , add it to the tracking offset $O_{i,j,k}^t$ for that location and associate the closest matching bounding box within the area at time $t - 1$. If the bounding box does not find a match, we compute cosine similarity of its embedding with all unmatched embeddings at time $t - 1$, and associate based on highest similarity. In case the bounding box still fails to find a match, we initiate a new tracklet.

E. Proposed Loss Formulation

We use three different losses, inspired by [11], [19], to train 3DMODT. The first one, Forward Association Loss, is used to learn accurate correspondences in the affinity matrix computed across time t and $t - \tau$. It is given by

$$L_a = \frac{\sum(G \odot (-\log \hat{A}^t))}{\sum G}, \quad (6)$$

where G refers to the ground-truth affinities and \hat{A}^t denotes the refined affinity matrix prediction.

The second loss is an ℓ_1 Center Offset Loss for supervision of 3D object centers only i.e.,

$$L_c = \frac{1}{R} \sum_{r=1}^R |\hat{C}_r - C_r|, \quad (7)$$

where R is the number of 3D objects, $\hat{C} \in \mathbb{R}^3$ implies the predicted object center and C represents the ground truth

center. The third loss is an ℓ_1 Dimension Size Loss to regress the bounding boxes,

$$L_b = \frac{1}{R} \sum_{r=1}^R |\hat{\gamma}_r - \gamma_r|, \quad (8)$$

where γ refers to the width, length and height of the 3D bounding boxes. The combined loss L is the weighted sum of the three losses

$$L = L_a + \lambda_c L_c + \lambda_b L_b, \quad (9)$$

where λ_c and λ_b are the loss hyper-parameters.

IV. EXPERIMENTS

We performed experiments on the JRDB dataset [50], Waymo Open dataset [46] and the KITTI tracking benchmark [51]. These datasets provide 3D point clouds captured using distinct LiDAR sensors.

Datasets: *JRDB* dataset was collected in a university campus setting and contains 360° views of both indoor buildings and outdoor pedestrian spaces. It captures entities of pedestrian class in densely populated areas such as cafeterias, and classrooms, which makes it a challenging dataset. It provides 54 sequences divided into 27 training and 27 test sequences. We use the recommended 7 sequences from the training split as the validation set for our experiments. *Waymo* dataset was captured in multiple cities, namely San Francisco, Phoenix, and Mountain View, providing large geographical coverage compared to other autonomous driving datasets. It contains vehicle, pedestrian, cyclist, and road-sign information for tracking task. The dataset comprises 798 training, 202 validation, and 150 test sequences, each spanning 20 seconds in duration. To evaluate the generalizability of the models on previously unseen locations, the test set includes geographically held-out areas. The evaluation toolkit provides two levels of difficulty: LEVEL_1 for assessing objects with more than 5 LiDAR points, and LEVEL_2 for objects with at least one LiDAR point. *KITTI* dataset, captured in Germany, provides tracking information for cars and pedestrians. It comprises 21 training sequences and 29 test sequences. The training sequences are further split into a training set comprising 10 sequences (3975 frames) and a validation set comprising 11 sequences (3945 frames). The original KITTI benchmark provides test data only for 2D evaluations. Weng et al. [52] proposed a procedure that is used to conduct 3D MOT evaluation.

Evaluation Metrics: We use standard 3D MOT evaluation measures [52]: *Multiple-Object Tracking Accuracy (MOTA)*, *Average MOTA (AMOTA)*, *scaled AMOTA (sAMOTA)*, degree of overlap between actual and predicted objects *Multiple-Object Tracking Precision (MOTP)*, *Average MOTP (AMOTP)*, *False Negatives (FN)*, *False Positives (FP)*, *Identity Switches (IDS)*, percentage of objects tracked for at least 80% of their total trajectory lengths *Mostly Tracked Trajectories (MT)*, percentage of objects recovered for less than 20% of their life span *Mostly Lost Trajectories (ML)*, trajectory *Fragmentation (FRAG)*.

TABLE I: **Results on JRDB dataset.** Our method exhibits competitive results on the JRDB test set without using additional input modality, any external detector, or a complex post-processing algorithm. Here, L implies LiDAR data, and I refers to camera feed.

Method	Input Data	External Detector	Complex Post-Processing	$MOTA\uparrow$	$MOTP\uparrow$	$IDS\downarrow$
AB3DMOT [17]	L	✓	✓	19.34	42.01	6.179
JRMOT [5]	$L + I$	×	✓	20.15	42.44	4.216
PC-DAN [16]	L	✓	✓	22.56	6.01	26.022
3DMODT	L	×	×	21.74	47.61	3.805

TABLE II: **Generalization performance.** Demonstration of the ability of the network to generalize across datasets.

Training Set	Testing Set	$MOTA\uparrow$
JRDB	JRDB	21.74
Waymo	JRDB	17.83
KITTI	KITTI	92.33
Waymo	KITTI	83.48

TABLE III: **Results on Waymo Open dataset.** We report the best tracking results on the Waymo vehicle validation set and the values are in the format LEVEL_1 / LEVEL_2. Here, L implies LiDAR data, and I refers to camera feed.

Method	Input Data	External Detector	Complex Post-Processing	$MOTA\uparrow$	$MOTP\uparrow$	Miss \downarrow	Miss Match \downarrow	$FP\downarrow$
Sun et al. [46]	L	✓	✓	42.5 / 40.1	18.6 / 18.6	40.0 / 43.4	0.14 / 0.13	17.3 / 16.4
CenterPoint [36]	L	×	×	51.4 / 47.9	17.6 / 17.6	47.7 / 41.4	0.19 / 0.18	10.7 / 10.6
SimTrack [19]	L	×	×	53.1 / 49.6	17.4 / 17.4	35.5 / 39.8	0.20 / 0.19	11.2 / 10.5
3DMODT	L	×	×	55.9 / 51.2	18.9 / 18.9	33.1 / 36.7	0.13 / 0.12	10.3 / 10.2

TABLE IV: **Results on KITTI tracking benchmark.** Our tracker demonstrates the best 3D tracking performance on the KITTI cars validation set. Here, L implies LiDAR data, and I refers to camera feed.

Method	Input Data	External Detector	Complex Post-Processing	$sMOTA\uparrow$	$AMOTA\uparrow$	$AMOTP\uparrow$	$MOTA\uparrow$	$MOTP\uparrow$
mmMOT [6]	$L + I$	✓	✓	70.61	33.08	72.45	74.07	78.16
FANTrack [47]	$L + I$	✓	×	82.97	40.03	75.01	74.30	75.24
AB3DMOT [17]	L	✓	✓	93.28	45.43	77.41	86.24	78.43
PC-TCNN [48]	L	×	×	95.44	47.64	—	—	—
3D DetecTrack [49]	$L + I$	×	×	96.49	48.87	81.56	91.46	82.24
3DMODT	L	×	×	96.91	49.01	82.14	92.33	83.52

TABLE V: **Ablation on JRDB validation set.** Components are added sequentially to highlight their contribution.

Method	$MOTA\uparrow$	$IDS\downarrow$
Baseline	77.31	14513
+ Self-Attention	91.34	3562
+ Cross-Attention	96.23	1960

Training Details: Our model exhibits a transformer encoder $\phi(\cdot)$ [53] for feature extraction and point convolutions $\theta(\cdot)$ [54] at the output layers. We use a single Nvidia V100 GPU for training our model. For the JRDB dataset, we train the model for 40 epochs using Adam optimizer [55] with an initial learning rate of 0.001 and reduce it by a factor of 5 every 10 epochs. In case of the Waymo dataset, we train the model for 30 epochs with the same initial learning rate, reduced after every 10 epochs. For the KITTI dataset, we train the model for 60 epochs using the same initial learning rate, reduced every 15 epochs.

A. Results and Analysis

Table I presents results on the JRDB test set. We observe that the proposed method achieves a competitive MOTA

of 21.74 without the complexity of an external detector or extensive post-processing steps that other methods use. Similarly, the results on the Waymo dataset in Table III further affirm our network’s ability to generate effective tracking results. Table IV shows results on the KITTI car tracking benchmark. Unlike some of the existing approaches, our model is self contained in that it relies on the detection and motion cues (tracking offsets) generated within our network to track the objects. Moreover, we do not use additional modalities, thereby eliminating any overhead related to sensor calibrations. In Figure 3, we visualize the results on Waymo. Qualitative results indicate the ability of our network to perform well under varying tracking conditions, such as long sequences, and adverse environmental situations.

Next, we examine the generalization ability of our method across datasets. We train our model on the Waymo dataset and use JRDB and KITTI sequences for testing. In Table II, we see that our model achieves good tracking results in both cross dataset settings. Our method eliminates the need for an external detector by treating detection and tracking as joint tasks in point cloud-based MOT. Given this capability and its generalizability, our model can be effectively deployed in real-world applications.



Fig. 3: Qualitative results of our model on diverse video sequences from Waymo dataset. Sequence 1 has long tracks. Sequence 2 has poor lighting conditions. Sequence 3 has adverse weather. Tracking is performed using LiDAR point clouds. The RGB images are for illustration only.

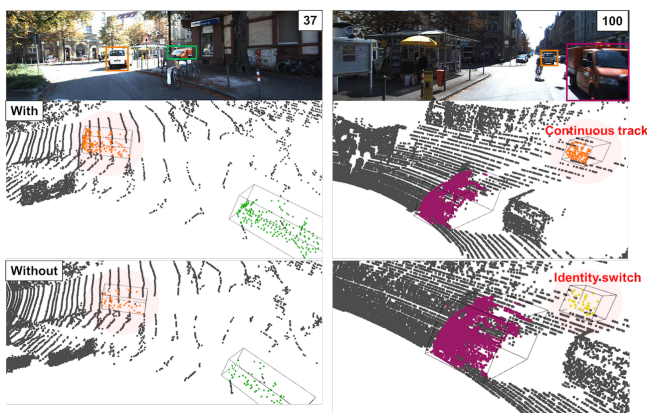


Fig. 4: Qualitative results illustrating the capability of attention-guided affinity module to boost the tracking performance. Row 1 depicts RGB sequence for illustration purposes where frame numbers are shown on the top right corners. Row 2 demonstrates ability of the network (with attention-guided affinities) to track the object in orange across time steps. Row 3 exhibits ID-switch encountered on the same object in the absence of attention-on-affinities.

B. Ablation Study

We sequentially add model components to validate the effectiveness of individual modules. Table V summarizes our results on the JRDB validation set. Here, we focus on the affinity matrices generated using transformer-based feature tokens. A single affinity matrix provides correspondences across two point cloud scans. First, we introduce only self-

attention on the affinities. Here, for a given affinity matrix, self-attention seeks consensus from the surrounding correlations of the same affinity matrix and suppresses erroneous correspondences. On the JRDB validation set, incorporating self-attention on affinity matrices boosts the MOTA by $\sim 14\%$. Secondly, we add cross-attention on affinities. In this case, the attention mechanism estimates the relevance of surrounding correlations across two different affinity matrices. Thus, using correlation cues from the subsequent affinity matrix, the quality of the given affinity matrix is improved. Intuitively, cross-attention on affinities leverages temporal context to boost tracking performance. We see an improvement of $\sim 5\%$, by introducing cross-attention on the affinities already refined by the self-attention mechanism (refer to the last row in Table V). Figure 4 illustrates the significance of attention-on-affinities on overall tracking performance.

V. CONCLUSION

We proposed a 3D point cloud-based joint detection and tracking model that eliminates the dependency on external detectors for object tracking. To overcome the challenges associated with sparsity and uneven density of 3D point cloud data, we introduced an attention mechanism and refined the affinity matrices that measure the correspondences across point cloud scans. Exhaustive experiments and ablations validated the effectiveness and generalizability of our approach in 3D MOT tasks. For future work, we will explore innovative techniques for reducing the computational requirements associated with attention-on-affinities, to allow us to incorporate a larger temporal context in the network.

REFERENCES

- [1] B. R. Kiran, I. Sobh, V. Talpaert, P. Mannion, A. A. Al Sallab, S. Yogamani, and P. Pérez, "Deep reinforcement learning for autonomous driving: A survey," *IEEE Transactions on Intelligent Transportation Systems*, 2021.
- [2] X. Weng, J. Wang, S. Levine, K. Kitani, and N. Rhinehart, "Sequential forecasting of 100,000 points," *arXiv preprint arXiv:2003.08376*, 2020.
- [3] A. Manglik, X. Weng, E. Ohn-Bar, and K. M. Kitani, "Forecasting time-to-collision from monocular video: Feasibility, dataset, and challenges," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2019, pp. 8081–8088.
- [4] X. Sun, X. Weng, and K. Kitani, "When we first met: Visual-inertial person localization for co-robot rendezvous," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2020, pp. 10 408–10 415.
- [5] A. Sheno, M. Patel, J. Gwak, P. Goebel, A. Sadeghian, H. Rezatofighi, R. Martin-Martin, and S. Savarese, "Jrmot: A real-time 3d multi-object tracker and a new large-scale dataset," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2020, pp. 10 335–10 342.
- [6] W. Zhang, H. Zhou, S. Sun, Z. Wang, J. Shi, and C. C. Loy, "Robust multi-modality multi-object tracking," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 2365–2374.
- [7] N. Wojke, A. Bewley, and D. Paulus, "Simple online and realtime tracking with a deep association metric," in *IEEE international conference on image processing*. IEEE, 2017, pp. 3645–3649.
- [8] K. Huang and Q. Hao, "Joint multi-object detection and tracking with camera-lidar fusion for autonomous driving," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2021, pp. 6983–6989.
- [9] J. Peng, C. Wang, F. Wan, Y. Wu, Y. Wang, Y. Tai, C. Wang, J. Li, F. Huang, and Y. Fu, "Chained-tracker: Chaining paired attentive regression results for end-to-end joint multiple-object detection and tracking," in *European conference on computer vision*. Springer, 2020, pp. 145–161.
- [10] Z. Lu, V. Rathod, R. Votel, and J. Huang, "Retinatrack: Online single stage joint detection and tracking," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 14 668–14 678.
- [11] J. Wu, J. Cao, L. Song, Y. Wang, M. Yang, and J. Yuan, "Track to detect and segment: An online multi-object tracker," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 12 352–12 361.
- [12] F. Zeng, B. Dong, T. Wang, X. Zhang, and Y. Wei, "Motr: End-to-end multiple-object tracking with transformer," *arXiv preprint arXiv:2105.03247*, 2021.
- [13] I. Mouawad and F. Odono, "Fastervideo: Efficient online joint object detection and tracking," in *International Conference on Image Analysis and Processing*. Springer, 2022, pp. 375–387.
- [14] Y. Wang, K. Kitani, and X. Weng, "Joint object detection and multi-object tracking with graph neural networks," in *2021 IEEE International Conference on Robotics and Automation*. IEEE, 2021, pp. 13 708–13 715.
- [15] R. Tapu, B. Mocanu, and T. Zaharia, "Deep-see: Joint object detection, tracking and recognition with application to visually impaired navigational assistance," *Sensors*, vol. 17, no. 11, p. 2473, 2017.
- [16] A. Kumar, J. Kini, M. Shah, and A. Mian, "Pc-dan: Point cloud based deep affinity network for 3d multi-object tracking," 2021.
- [17] X. Weng, J. Wang, D. Held, and K. Kitani, "3D Multi-Object Tracking: A Baseline and New Evaluation Metrics," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2020.
- [18] A. Kumar, J. Kini, A. Mian, and M. Shah, "Self supervised learning for multiple object tracking in 3d point clouds," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2022, pp. 3754–3761.
- [19] C. Luo, X. Yang, and A. Yuille, "Exploring simple 3d multi-object tracking for autonomous driving," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 10 488–10 497.
- [20] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," *Advances in neural information processing systems*, vol. 28, 2015.
- [21] Z. Ren, Z. Yu, X. Yang, M.-Y. Liu, A. G. Schwing, and J. Kautz, "Ufo: A unified framework towards omni-supervised object detection," in *European Conference on Computer Vision*. Springer, 2020, pp. 288–313.
- [22] Z. Ren, Z. Yu, X. Yang, M.-Y. Liu, Y. J. Lee, A. G. Schwing, and J. Kautz, "Instance-aware, context-focused, and memory-efficient weakly supervised object detection," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 10 598–10 607.
- [23] Y. Zou, X. Yang, Z. Yu, B. Kumar, and J. Kautz, "Joint disentangling and adaptation for cross-domain person re-identification," in *European Conference on Computer Vision*. Springer, 2020, pp. 87–104.
- [24] Z. Zheng, X. Yang, Z. Yu, L. Zheng, Y. Yang, and J. Kautz, "Joint discriminative and generative learning for person re-identification," in *proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 2138–2147.
- [25] R. E. Kalman, "A new approach to linear filtering and prediction problems," 1960.
- [26] X. Zhou, V. Koltun, and P. Krähenbühl, "Tracking objects as points," in *European Conference on Computer Vision*. Springer, 2020, pp. 474–490.
- [27] Y. Zhang, C. Wang, X. Wang, W. Zeng, and W. Liu, "Fairmot: On the fairness of detection and re-identification in multiple object tracking," *International Journal of Computer Vision*, vol. 129, no. 11, pp. 3069–3087, 2021.
- [28] Z. Wang, L. Zheng, Y. Liu, Y. Li, and S. Wang, "Towards real-time multi-object tracking," in *European Conference on Computer Vision*. Springer, 2020, pp. 107–122.
- [29] A. Arnab, M. Dehghani, G. Heigold, C. Sun, M. Lučić, and C. Schmid, "Vivit: A video vision transformer," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 6836–6846.
- [30] T. Meinhardt, A. Kirillov, L. Leal-Taixe, and C. Feichtenhofer, "Trackerformer: Multi-object tracking with transformers," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 8844–8854.
- [31] P. Sun, J. Cao, Y. Jiang, R. Zhang, E. Xie, Z. Yuan, C. Wang, and P. Luo, "Transtrack: Multiple object tracking with transformer," *arXiv preprint arXiv:2012.15460*, 2020.
- [32] 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.
- [33] Y. Yan, Y. Mao, and B. Li, "Second: Sparsely embedded convolutional detection," *Sensors*, vol. 18, no. 10, p. 3337, 2018.
- [34] 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.
- [35] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 11 553–11 562.
- [36] 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.
- [37] 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*. IEEE, 2021, pp. 14 227–14 233.
- [38] M. Liang, B. Yang, S. Wang, and R. Urtasun, "Deep continuous fusion for multi-sensor 3d object detection," in *Proceedings of the European conference on computer vision*, 2018, pp. 641–656.
- [39] M. Liang, B. Yang, Y. Chen, R. Hu, and R. Urtasun, "Multi-task multi-sensor fusion for 3d object detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 7345–7353.
- [40] S. Pang, D. Morris, and H. Radha, "Clocs: Camera-lidar object candidates fusion for 3d object detection," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2020, pp. 10 386–10 393.
- [41] J. H. Yoo, Y. Kim, J. Kim, and J. W. Choi, "3d-cvf: Generating joint camera and lidar features using cross-view spatial feature fusion for 3d object detection," in *European Conference on Computer Vision*. Springer, 2020, pp. 720–736.
- [42] X. Wang, R. Girshick, A. Gupta, and K. He, "Non-local neural

- networks,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 7794–7803.
- [43] S. Gao, C. Zhou, C. Ma, X. Wang, and J. Yuan, “Aiatrack: Attention in attention for transformer visual tracking,” *arXiv preprint arXiv:2207.09603*, 2022.
- [44] H. Zhao, L. Jiang, J. Jia, P. H. Torr, and V. Koltun, “Point transformer,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, October 2021, pp. 16 259–16 268.
- [45] C. Zhou, Z. Luo, Y. Luo, T. Liu, L. Pan, Z. Cai, H. Zhao, and S. Lu, “Ptr: Relational 3d point cloud object tracking with transformer,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 8531–8540.
- [46] P. Sun, H. Kretzschmar, X. Dotiwalla, A. Chouard, V. Patnaik, P. Tsui, J. Guo, Y. Zhou, Y. Chai, B. Caine, *et al.*, “Scalability in perception for autonomous driving: Waymo open dataset,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 2446–2454.
- [47] E. Baser, V. Balasubramanian, P. Bhattacharyya, and K. Czarnecki, “Fantrack: 3d multi-object tracking with feature association network,” in *2019 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2019, pp. 1426–1433.
- [48] H. Wu, Q. Li, C. Wen, X. Li, X. Fan, and C. Wang, “Tracklet proposal network for multi-object tracking on point clouds,” in *Proceedings of the International Joint Conference on Artificial Intelligence*, 2021, pp. 1165–1171.
- [49] J. Koh, J. Kim, J. H. Yoo, Y. Kim, D. Kum, and J. W. Choi, “Joint 3d object detection and tracking using spatio-temporal representation of camera image and lidar point clouds,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 1, 2022, pp. 1210–1218.
- [50] R. Martín-Martín, M. Patel, H. Rezatofighi, A. Sheno, J. Gwak, E. Frankel, A. Sadeghian, and S. Savarese, “Jrdb: A dataset and benchmark of egocentric visual perception for navigation in human environments,” *arXiv preprint arXiv:1910.11792*, 2019.
- [51] A. Geiger, P. Lenz, and R. Urtasun, “Are we ready for autonomous driving? the kitti vision benchmark suite,” in *Conference on Computer Vision and Pattern Recognition*, 2012.
- [52] X. Weng, J. Wang, D. Held, and K. Kitani, “3d multi-object tracking: A baseline and new evaluation metrics,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2020, pp. 10 359–10 366.
- [53] M.-H. Guo, J.-X. Cai, Z.-N. Liu, T.-J. Mu, R. R. Martin, and S.-M. Hu, “Pct: Point cloud transformer,” *Computational Visual Media*, vol. 7, no. 2, pp. 187–199, 2021.
- [54] H. Thomas, C. R. Qi, J.-E. Deschaud, B. Marcotegui, F. Goulette, and L. J. Guibas, “Kpconv: Flexible and deformable convolution for point clouds,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 6411–6420.
- [55] P. Kingma Diederik and J. B. Adam, “A method for stochastic optimization,” *Anon. International Conference on Learning Representations*, 2015.