

# RaTrack: Moving Object Detection and Tracking with 4D Radar Point Cloud

Zhijun Pan<sup>†</sup>, Fangqiang Ding<sup>†</sup>, Hantao Zhong<sup>†</sup>, and Chris Xiaoxuan Lu<sup>\*</sup>

**Abstract**—Mobile autonomy relies on the precise perception of dynamic environments. Robustly tracking moving objects in 3D world thus plays a pivotal role for applications like trajectory prediction, obstacle avoidance, and path planning. While most current methods utilize LiDARs or cameras for Multiple Object Tracking (MOT), the capabilities of 4D imaging radars remain largely unexplored. Recognizing the challenges posed by radar noise and point sparsity in 4D radar data, we introduce RaTrack, an innovative solution tailored for radar-based tracking. Bypassing the typical reliance on specific object types and 3D bounding boxes, our method focuses on motion segmentation and clustering, enriched by a motion estimation module. Evaluated on the View-of-Delft dataset, RaTrack showcases superior tracking precision of moving objects, largely surpassing the performance of the state of the art. We release our code and model at <https://github.com/LJacksonPan/RaTrack>.

## I. INTRODUCTION

Ensuring accurate perception of dynamic environments is pivotal for mobile autonomy. A crucial task in this domain is the consistent and robust tracking of moving objects in 3D space for autonomous vehicles. This ability acts as a cornerstone for subsequent autonomy tasks such as trajectory prediction [1]–[3], obstacle avoidance [4]–[6], and path planning [7]–[9].

State-of-the-art techniques in moving object tracking or Multiple Object Tracking (MOT) primarily use on-vehicle LiDARs [10]–[12], cameras [13]–[15], or their fusion [16]–[18]. Surprisingly, the potential of 4D mmWave radars remains under-explored. As an emerging automotive sensor, 4D mmWave radar is gaining traction due to its improved imaging ability, resilience against challenging weather and illumination conditions (e.g., fog, dust, darkness, glare), ability to measure object velocities and cost-effectiveness. These merits make 4D radar an appealing and robust supplement, or even alternative, to automotive LiDARs.

However, integrating 4D radars into moving object tracking presents non-trivial challenges. The prevalent approaches [10]–[12], [17], [19] often follow the *tracking-by-detection* paradigm. Such a paradigm involves first detecting objects in each frame independently and then linking

these *detected object types and 3D bounding boxes* across consecutive frames to form continuous object trajectories. Key to the *tracking-by-detection* success depends on the detection accuracy. This paradigm struggles when adapted to 4D radar data, due to the inherent radar noise and point sparsity, undermining accurate type classification and bounding box regression. Specifically, the non-negligible multi-path noises in radar data complicate the correct identification of objects while the sparsity of radar point clouds makes the object property (e.g. extension and orientation) regression even more difficult. As exhibited in [20], the mAP performance [21] of the 4D radar detection method is only 38.0,  $\sim 40\%$  inferior to its LiDAR counterpart in the same scene. Such poor 3D detection performance compromises the reliability of 4D radar-based tracking in real-world scenarios.

To address this, we present RaTrack, a first-of-its-kind tailored solution for moving object tracking using 4D automotive radars. Our approach stems from a critical insight: *for effective multi-object tracking, class-agnostic detection is often adequate, and the conventional reliance on 3D bounding boxes becomes redundant if distinct point clusters can be utilized*. Driven by this understanding, we restructure the moving object detection challenge into simpler motion segmentation and clustering tasks. This restructuring allows us to sidestep the complex tasks of object type classification and bounding box regression typically encountered with 4D radars. To further enhance our method’s performance, we integrate a point-wise motion estimation module, which enriches the inherently sparse radar data with point-level scene flow. Building on this, our data association module is precisely adapted to our clustering method and is calibrated to weigh different features for optimal matching. Our solution is architected as an end-to-end trainable network, with its training modelled as a multi-task learning endeavour. This encompasses motion segmentation, scene flow estimation, and affinity matrix computation.

Extensive experiments on the View-of-Delft dataset [20] validate the superiority of RaTrack over existing techniques in moving object detection and tracking precision. Moreover, our results underscore the merits of the cluster-based detection method and the instrumental role of scene flow estimation in both detection and data association phases.

## II. RELATED WORKS

Given the absence of prior work on 4D radar-based moving object tracking, we will touch on existing research in general 3D MOT. As an uplift of 2D MOT [22]–[26] in the 3D space, 3D MOT has attracted increasing interests

<sup>†</sup>Equal Contribution, listed randomly. <sup>\*</sup>Corresponding author.

Zhijun Pan is with the Computer Science Research Centre, Royal College of Art, United Kingdom. This work was partly done when he was with the School of Informatics, University of Edinburgh.

Fangqiang Ding is with the School of Informatics, University of Edinburgh, United Kingdom. This research is supported by the EPSRC, as part of the CDT in Robotics and Autonomous Systems hosted at the Edinburgh Centre of Robotics (EP/S023208/1).

Hantao Zhong is with the Department of Computer Science and Technology, University of Cambridge, United Kingdom.

Chris Xiaoxuan Lu is with the Department of Computer Science, University College London, United Kingdom.

due to its significant application to autonomous systems. Most online 3D MOT solutions adopt a tracking-by-detection approach, focusing on 3D bounding box detection and data association. The core of data association lies in extracting comprehensive tracking cues and matching new detections with previous tracklets.

**3D bounding box detection.** The premise of applying the tracking-by-detection pipeline is accurate 3D bounding box estimation. Thanks to recent advances in 3D object detection, many off-the-shelf detectors [11], [27]–[31] have already been employed as the front-end of modern tracking systems. According to their input modalities used for 3D detection, current 3D MOT systems can be classified into LiDAR point cloud-based [10], [12], [32]–[40], image-based [13], [14], [41], [42] and LiDAR-image fusion-based [16]–[19] methods. Different from these approaches, our method takes only 4D radar point clouds as the input and detects object instances as clusters of points instead of 3D bounding boxes for tracking.

**Tracking cues.** To exploit the 3D motion information, AB3DMOT [10] proposes a baseline method that models the motion of objects with a Kalman filter and predicts the displacement with a constant velocity model. The same strategy and its variants [12] are followed by later works [16], [17], [35]–[37] to induce motion cues for association. Another common strategy is to directly regress object velocities from the detectors [11], [18], [38]. In [19], [34], latent motion features are extracted by an LSTM network for tracked objects. Apart from motion cues, object appearance features are usually learned from neural networks, either from images [32], LiDAR point clouds [19], [38], [40] or both of them [12], [16], [18], [34], for data association. Unlike prior works, we estimate per-point scene flow vectors to obtain motion cues, which can not only help to match objects with similar motion but also benefit moving object detection. Besides scene flow, we aggregate complementary point-level features from each cluster for robust data association.

**Track-detection matching.** Given object motion or appearance cues, most methods generate an affinity map based on object motion or appearance cues, capturing matching scores for potential track-detection pairs. Some methods [10], [16], [32] use traditional distance metrics like cosine or  $L_2$  distances, while others [19], [34] employ networks for learnable distance metrics. Assignments are typically resolved using the Hungarian [43] or greedy algorithms. Recent techniques [37], [39] employ graph structures and Neural Message Passing [44] for more direct associations. In our approach, inspired by [19], [34], we use MLP networks to estimate cluster-pair similarities. Uniquely, we adopt the differentiable Sinkhorn algorithm [45] for bipartite matching, rendering the data association process fully differentiable, enhancing tasks like trajectory prediction and planning.

### III. PROBLEM FORMULATION

**Scope.** In a standard 3D MOT setup, all objects of interest, such as cars and motorcycles, are tracked regardless of their motion status [10], [36], [39]. Contrarily, this work solely

focuses on the moving objects. This focus stems from the premise that dynamic entities hold greater significance for tracking than static ones. Additionally, the inherent ability of radar sensors to measure velocity makes it a trivial task to distinguish between moving and stationary objects.

**Notation.** Given this context, we consider the problem of online *moving object detection and tracking* with 4D automotive radar. The input is an ordered 4D radar point cloud sequence  $\mathcal{P} = \{\mathbf{P}^t\}_{t=1}^T$  comprised of  $T$  frames captured by the same radar sensor mounted on a moving vehicle. A frame  $\mathbf{P}^t = [\mathbf{p}_1^t; \dots; \mathbf{p}_i^t; \dots; \mathbf{p}_{N^t}^t]$  contains  $N^t$  radar points. Each radar point  $\mathbf{p}_i^t$  is characterised by its 3D position  $\mathbf{x}_i^t \in \mathbb{R}^3$  in the metric space and auxiliary velocity features  $\mathbf{v}_i^t = [v_{r,i}^t, v_{c,i}^t]$ , where  $v_{r,i}^t$  and  $v_{c,i}^t$  are the measured relative radial velocity (RRV) and its ego-motion (assumed known) compensated variant. Given each radar point cloud  $\mathbf{P}^t$  from the stream, our objective is to detect multiple moving objects  $\mathbf{D}^t = \{\mathbf{d}_k^t\}_{k=1}^{K^t}$  in a class-agnostic manner without the need to regress their 3D bounding boxes. These detected objects are then associated with objects tracked in the previous frame  $\mathbf{O}^{t-1} = \{\mathbf{o}_m^{t-1}\}_{m=1}^{M^{t-1}}$ . The result of this process is a set of updated objects  $\mathbf{O}^t$  in track for the current frame.

## IV. PROPOSED METHOD

### A. Overview

We introduce RaTrack, a generic learning-based framework bespoke for 4D radar-based moving object detection and tracking. As seen in Fig. 1, in our network pipeline, we first apply a backbone (c.f. Sec. IV-B) to encode intra- and inter-frame radar point cloud features. With the extracted features, our point-wise motion estimation module (c.f. Sec. IV-C) infers point-level scene flow as an explicit complement to augment the latent features of radar point clouds. Our advocated idea of class-agnostic detection without bounding boxes is introduced in the object detection module (c.f. Sec. IV-D), in which moving points are first identified and then used to detect moving objects via clustering. Finally, our data association module (c.f. Sec. IV-E) computes the affinity matrix with a learnable distance metric and then optimises the bipartite matching problem. The entire network is end-to-end trainable with a multi-task loss that incorporates three supervised subtasks: motion segmentation, scene flow estimation, and affinity matrix computation.

### B. Backbone Network

On receiving a new 4D radar point cloud  $\mathbf{P}^t$ , our backbone neural network is used to extract representative latent features for each radar point to facilitate subsequent tasks, e.g., motion segmentation and scene flow estimation. To this end, we first extract point-level local-global features  $\mathbf{G}^t$  using a point feature encoder (PFE), which comprises a) three set abstraction layers [46] to extract local features at different scales in parallel, b) three MLP-based feature propagation layer to map local features into high-level representations, and c) the max-pooling operation to aggregate the global feature vector that is attached to per-point features. To further encode inter-frame point motion for the current frame, we

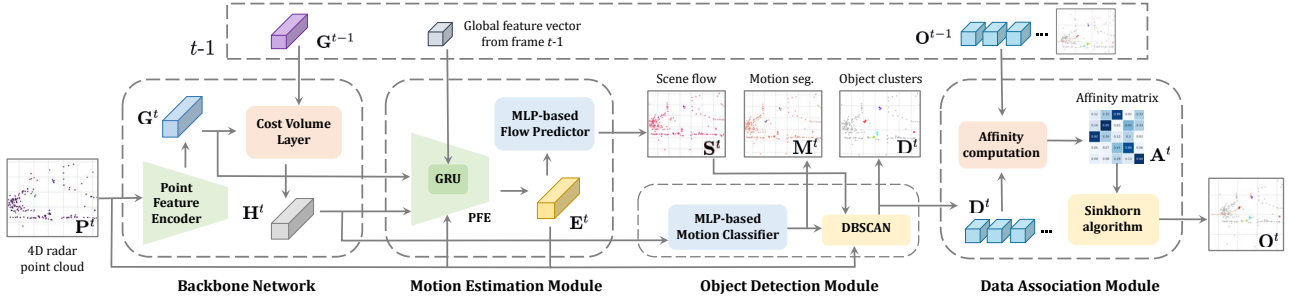


Fig. 1. Overall network pipeline of RaTrack. Note that radar 3D points are shown in the bird’s eye view for visualization.

recall the features  $\mathbf{G}^{t-1}$  from the last frame and correlate features across two frames by the cost volume layer [47], as seen in Fig. 1. The output is the cost volume  $\mathbf{H}^t$  that represents the motion information for each point in  $\mathbf{P}^t$ .

### C. Motion Estimation Module

Radar sensors’ inherent sparsity and noise result in latent features from radar point clouds that are deficient in informative geometric cues, complicating object detection and tracking. To address this, we introduce a point-wise motion estimation module that explicitly determines per-point motion vectors. Contrary to the conventional scene flow approach, which estimates per-point motion vectors in a forward direction (from frame  $t$  to  $t + 1$ ), we opt for a backward estimation (from frame  $t$  to  $t - 1$ ). This not only negates tracking latency but also ensures that the estimated flow vectors correspond to the points in the current frame  $t$ .

Before decoding scene flow, we first aggregate mixed features by integrating diverse backbone outputs: input point features  $\mathbf{P}^t$ , local-global features  $\mathbf{G}^t$  and cost volume  $\mathbf{H}^t$ . Subsequently, another PFE is employed, aiming to facilitate information exchange among these diverse features and enhance their spatial coherence, producing the flow embedding  $\mathbf{E}^t$  as shown in Fig. 1. Notably, within this PFE, a GRU network [48] is utilized to introduce temporal information into the global feature vector prior to its association with per-point features. The scene flow  $\mathbf{S}^t = [\mathbf{s}_1^t; \dots; \mathbf{s}_i^t; \dots; \mathbf{s}_{N^t}^t] \in \mathbb{R}^{N^t \times 3}$  is finally decoded with an MLP-based flow predictor. This module plays a crucial role in our framework and is leveraged to augment per-point features before clustering (c.f. Sec. IV-D) and to lead up data association as an extra motion clue (c.f. Sec. IV-E).

### D. Object Detection Module

In the classical tracking-by-detection paradigm, objects are first detected as class-specific 3D bounding boxes based on which tracklets are built up across frames. However, the performance of such approaches inevitably relies on accurate type classification and bounding box regression, which are hard to accomplish from the sparse and noisy 4D radar data.

Rather than relying on the error-prone 3D bounding box detection tailored for specific object types, we emphasize the fundamental necessity to simply group scattered points into clusters for effective tracking. Consequently, we champion a class-agnostic object detection approach that eschews

bounding boxes in our solution. By adopting this methodology, the fallible bounding box detector is replaced by a more dependable combination of motion segmentation and clustering, which proves to be particularly suited for radar point clouds. In other words, we detect objects in a bottom-up fashion, where points are first classified into moving and static (i.e., motion segmentation) and those close in the latent feature space are aggregated into object clusters. For motion segmentation, we leverage the cost volume  $\mathbf{H}^t$  provided by the backbone and compute the moving possibility score  $c_i^t$  for each point  $\mathbf{p}_i^t$  through an MLP-based motion classifier. The classification results are reliable enough as both the crucial RRV measurements and the inter-frame motion information are encoded by our backbone, yielding a robust motion representation. A fixed threshold  $\zeta_{mov}$  is further used to separate the moving targets from the static background, resulting in a motion segmentation mask  $\mathbf{M}^t = \{m_i^t \in \{0, 1\}\}_{i=1}^{N^t}$  as exhibited in Fig. 1. To delineate the boundaries of moving objects from the pinpointed moving points, we employ the classic clustering algorithm, DBSCAN [49]. This groups analogous points into object clusters, represented as  $\mathbf{D}^t = \{\mathbf{d}_k^t\}_{k=1}^{K^t}$  and we term  $\mathbf{D}^t$  as the detected (moving) objects hereafter. For robust clustering, we utilize the point cloud  $\mathbf{P}_t$ , their estimated scene flow  $\mathbf{S}^t$  and the flow embedding  $\mathbf{E}^t$  as the salient features and identify neighbour points. In this way, each detected object  $\mathbf{d}_k^t$  is represented as a cluster containing a subset of points  $\{\mathbf{p}_{k,j}^t\}_j$  with their corresponding scene flow and flow embedding vectors  $\{[\mathbf{s}_{k,j}^t, \mathbf{e}_{k,j}^t]\}_j$ . It is worth noting that we forego estimating explicit object categories (e.g., car, pedestrian). For the purposes of object tracking, such categorization is *not* imperative. Instead, we identify them simply as *class-agnostic* entities.

### E. Data Association Module

Given the objects  $\mathbf{D}^t$  detected in the current frame, our data association module endeavors to align them with the previously tracked objects  $\mathbf{O}^{t-1}$  from frame  $t - 1$ . For end-to-end learning and inspired by [12], [19], [34], we opt for the MLP over traditional hand-crafted metrics to derive the affinity matrix  $\mathbf{A}^t \in \mathbb{R}^{K^t \times M^t}$  for bipartite matching. Such a learnable distance metric can automatically adjust the weights of different features when calculating the similarity scores. For each pair of clusters  $\{\mathbf{d}_k^t, \mathbf{o}_m^{t-1}\}$ , its corresponding similarity score  $a_{k,m}^t$  can be computed as follows:

$$a_{k,m}^t = MLP(\mathbf{I}_k^t - \mathbf{I}_m^{t-1}) \quad (1)$$

where  $\mathbf{I}_k^t, \mathbf{I}_m^{t-1}$  are the aggregated features of two clusters respectively. Taking the object  $\mathbf{d}_k^t$  as the example, to generate its aggregated features  $\mathbf{I}_k^t$ , we concatenate a) the average and variance of its associated point subset  $\{\mathbf{p}_{k,j}^t\}_j$ , b) the max-pooling of the point-level scene flow and embedding vectors  $\{[\mathbf{s}_{k,j}^t, \mathbf{e}_{k,j}^t]\}_j$ . This process can effectively aggregate the information for each cluster and ensure the dimension consistency given clusters with various numbers of points.

Once the affinity matrix  $\mathbf{A}^t$  is computed, we identify the optimal matching pairs based on their similarity scores. To address this optimization challenge, we employ the Sinkhorn algorithm [45], as shown in Fig. 1. This method involves iterative normalization of  $\exp(\mathbf{A}^t)$  across both rows and columns, ensuring the entire data association process remains differentiable. Post optimization, we reassign object IDs to the successfully matched pairs, allocate new IDs for newly detected objects, and remove IDs associated with previously tracked objects absent in the current frame. Notably, we adopt the finalized scores of matched pairs as confidence scores for currently detected objects, given the inability of our detection module to provide such scores. Such confidence scores are used for certain metrics, such as AMOTA and AMOTP, which integrate results over all recall values.

#### F. End-to-End Training

We leverage labelled samples to end-to-end train our network with a multi-task loss of scene flow estimation, motion segmentation and affinity matrix computation:

$$\mathcal{L} = \alpha_1 \mathcal{L}_{flow} + \alpha_2 \mathcal{L}_{seg} + \alpha_3 \mathcal{L}_{aff} \quad (2)$$

where  $\alpha_1, \alpha_2, \alpha_3$  are hyperparameters to weigh different loss functions. For scene flow loss, we compute the  $L_2$  distance between the estimated scene flow  $\mathbf{S}^t$  and the ground truth one  $\tilde{\mathbf{S}}^t = \{\tilde{\mathbf{s}}_i^t\}_{i=1}^{N^t}$ :

$$\mathcal{L}_{flow} = \frac{1}{N^t} \sum_i \|\mathbf{s}_i^t - \tilde{\mathbf{s}}_i^t\|_2^2 \quad (3)$$

Given the ground truth motion segmentation mask  $\tilde{\mathbf{M}}^t = \{\tilde{m}_i^t\}_{i=1}^{N^t}$ , we separately supervise the classification scores of real moving and static points using the cross-entropy to address the low ratio ( $< 10\%$ ) of moving points in point clouds. The motion segmentation loss can be written as:

$$\mathcal{L}_{seg} = \beta \frac{\sum_i (1 - \tilde{m}_i^t) \log(1 - c_i^t)}{\sum_i 1 - \tilde{m}_i^t} + (1 - \beta) \frac{\sum_i \tilde{m}_i^t \log(c_i^t)}{\sum_i \tilde{m}_i^t} \quad (4)$$

where  $\beta$  is used to balance the influence of moving and static points. To supervise the computation of the affinity matrix  $\mathbf{A}^t$ , we formulate the prediction of the similarity score as a binary classification (matched or unmatched) problem and compute the binary cross-entropy loss as:

$$\mathcal{L}_{aff} = \frac{1}{K^t M^t} \sum_k \sum_m \tilde{a}_{k,m}^t \log(a_{k,m}^t) + (1 - \tilde{a}_{k,m}^t) \log(1 - a_{k,m}^t) \quad (5)$$

where  $\tilde{a}_{k,m}^t \in \{0, 1\}$  is the ground truth affinity score for object pair  $\{\mathbf{d}_k^t, \mathbf{o}_m^{t-1}\}$ . To see our ground truth label generation process, please refer to Sec. V-B.

Note that the above tasks are inherently intertwined and simultaneously optimized via end-to-end training. Supervising both scene flow estimation and motion segmentation directly aids the computation of the affinity matrix, which utilizes scene flow and clusters of moving points as input. Conversely, the gradients originating from the affinity matrix loss instruct the backbone to encode potent features from point clouds. This indirect guidance subsequently enhances both motion segmentation and scene flow estimation.

## V. EXPERIMENTS

### A. Evaluation Settings

**Dataset.** In our experiments, we demonstrate the effectiveness of RaTrack using the View-of-Delft (VoD) dataset [20], which includes essential components (i.e., 4D radar point clouds, odometry information, object bounding boxes and tracking IDs annotations) for our problem. As an official benchmark specific for 3D object detection, the annotations of its test split are not publicly available, thereby we evaluate our trained models with its validation split, which is unseen during our training process.

**Evaluation metrics.** To quantify our performance, we use the classical MOTA, MODA, MT, ML metrics [50], [51] and the popular sAMOTA, AMOTA, AMOTP metrics [10] for evaluation. To make these metrics adapt to our cluster-based object detections, we compute the IoU by counting the number of intersected and united radar points between the ground truth object and the predicted one. The threshold for our point-based IoU is set as 0.25 across all experiments.

**Baselines.** As there are no prior works for 4D radar-based moving object tracking, we select two state-of-the-art LiDAR-oriented 3D MOT methods, i.e., AB3DMOT [10] and CenterPoint [11] as our baselines. To ensure the comparison is fair, we keep their original settings and also train their models on the VoD training split. Note that baseline methods, though designed for LiDAR, also take 4D radar point clouds as input in this work. Specifically, we develop two augmented baselines AB3DMOT-PP and CenterPoint-PP by replacing the detector with PointPillars [52] for AB3DMOT and the backbone with that of PointPillars for CenterPoint.

### B. Implementation Details

**Label generation.** We follow [53]–[56] to generate *pseudo* scene flow labels using the ego-motion and object annotations. The ground truth motion segmentation mask can then be obtained by thresholding after compensating the ego-motion from the scene flow. To get the ground truth affinity matrix, we first match ground truth objects with our detected objects. A detected object that has a point-based IoU higher than 0.25 with any ground truth object will be assigned the same ID as the ground truth. Then, the ground truth affinity matrix can be constructed by assigning  $\tilde{a}_{k,m}^t = 1$  if the object pair  $\{\mathbf{d}_k^t, \mathbf{o}_m^{t-1}\}$  has the same ID, and vice versa.

**Hyperparameters.** Our hyperparameters are determined both empirically and according to the testing results. The threshold  $\zeta_{mov}$  (c.f. Sec. IV-D) is simply set as 0.5. For the DBSCAN algorithm, the neighbourhood radius is 1.5m while

TABLE I

PERFORMANCE OF RATrack AND BASELINES ON VOD. BASELINES WITH (A) REPRESENT METHODS TRAINED AND EVALUATED ON ALL OBJECTS, WHILE OTHERS ARE TRAINED AND EVALUATED ONLY ON MOVING OBJECTS, WHICH SERVE AS THE MAIN COMPARISON FOR RATrack.

| Method                        | sAMOTA [%] $\uparrow$ | AMOTA [%] $\uparrow$ | AMOTP [%] $\uparrow$ | MOTA [%] $\uparrow$ | MODA [%] $\uparrow$ | MT [%] $\uparrow$ | ML [%] $\downarrow$ |
|-------------------------------|-----------------------|----------------------|----------------------|---------------------|---------------------|-------------------|---------------------|
| CenterPoint [11] (A)          | 37.72                 | 8.77                 | 38.92                | 33.34               | 35.17               | 11.73             | 57.41               |
| CenterPoint-PP [11], [52] (A) | 42.64                 | 10.87                | 43.64                | 36.20               | 37.29               | 15.43             | 50.62               |
| AB3DMOT [10] (A)              | 33.56                 | 6.66                 | 33.34                | 31.00               | 31.20               | 14.81             | 64.20               |
| AB3DMOT-PP [10], [52] (A)     | 35.82                 | 7.67                 | 36.70                | 38.44               | 41.96               | 19.12             | 38.24               |
| CenterPoint [11]              | 43.21                 | 14.40                | 54.55                | 38.44               | 41.96               | 19.12             | 38.24               |
| CenterPoint-PP [11], [52]     | 44.54                 | 16.33                | 58.80                | 43.96               | 44.91               | 19.12             | 54.41               |
| AB3DMOT [10]                  | 51.23                 | 15.00                | 53.21                | 46.72               | 47.38               | 20.59             | 39.71               |
| AB3DMOT-PP [10], [52]         | 60.71                 | 21.51                | <b>62.75</b>         | 49.38               | 49.86               | 26.47             | 33.82               |
| <b>RaTrack</b>                | <b>74.16</b>          | <b>31.50</b>         | 60.17                | <b>67.27</b>        | <b>77.83</b>        | <b>42.65</b>      | <b>14.71</b>        |




Fig. 2. Qualitative results of RaTrack. We show on RGB images the ground truth 3D bounding boxes of moving objects (colors indicate different object classes) and projected radar point clouds (colors indicate the distances of points). On corresponding bird's eyes view figures, we show the detected moving objects and their predicted trajectories, where different colors are used to distinguish multiple object clusters.

the minimum number of points in a cluster is set as 2. For loss weights, we set  $\alpha_1, \alpha_2, \alpha_3$  in the overall loss (c.f. Eq. 2) to be 0.5, 0.5 and 1.0 respectively, and set  $\beta$  in the motion segmentation loss (c.f. Eq. 4) to be 0.4.

**Network training.** Training our network is non-trivial due to the dependencies of intermediate outputs across modules and frames. For example, the efficacy of our data association module depends on whether moving objects are correctly detected in the object detection module. To more effectively train our network, we separate the training into two stages: 1) We first train the backbone and the class predictor used for motion segmentation with Eq. 4 and keep other components frozen. The training keeps for 16 epochs with an initial learning rate of 0.001. This allows for fast learning of accurate moving object detection before data association. 2) We then train the whole network end-to-end for an additional 8 epochs with an initial learning rate of 0.0008. The Adam optimizer [57] is used for network parameter updates and the learning rate decays by 0.97 per epoch in both stages.

### C. Overall Performance

We evaluate RaTrack and our baseline methods and compare their results on different metrics, as shown in Table I. RaTrack exhibits a remarkable gain over baselines, increasing the sAMOTA, AMOTA and MOTA by 13.4%, 10.0%, 17.9% compared to the second-best scores. Such results demonstrate its superiority in 4D radar-based moving object detection and tracking. Notably, RaTrack achieves an improvement of 28.0% on the MODA metric that is used to measure the object detection accuracy. This supports that our class-agnostic object detection without bounding boxes is a better option than current 3D bounding box detectors [11], [29], [52] for recognizing and localizing moving objects in 4D radar point clouds. With more reliable object detection

and data association module, RaTrack also surpasses all baselines on the MT and ML metric, which demonstrates its ability to maintain long-term tracking of moving objects. It can be also observed that RaTrack has a slightly lower AMOTP than the AB3DMOT-PP baseline [10], [52]. As a precision metric, AMOTP only calculates the IoU between successfully matched object pairs, thus it becomes less important when the MOTA and MODA, which assess the tracking and detection accuracy, are considerably lower. We also show some qualitative results of RaTrack in Fig. 2.

As seen in Table I, we also train and evaluate baseline models with all annotated objects rather than only the moving ones to compare their performance in two settings. A consistent performance gap can be observed between the two settings of our baselines, which demonstrates that detecting and tracking moving objects is easier than that of general objects. Moreover, we find that our two augmented baselines (i.e., with -PP) outperform the original ones. We credit this to PointPillars [52] which is the state-of-the-art for 4D radar-based 3D object detection as exhibited in [20].

### D. Ablation Study

To validate the effectiveness of the object detection module (ODM) (c.f. Sec. IV-D) and motion estimation module (MEM) (c.f. Sec. IV-C), here we conduct an ablation study to see their impact. We also analyze the impact of the auxiliary velocity features. The results are shown in Table II.

**Object detection module.** By replacing our ODM with PointPillars detector [52], the performance on sAMOTA, MOTA and MODA degrades by 21.0%, 25.6% and 35.5% respectively. This supports again that our cluster-based object detection is more effective than previous detectors for moving object detection and can thus yield more accurate tracking results. However, the outcome of this ablated version

TABLE II

ABLATION STUDY RESULTS FOR RaTrack. FOR THE FIRST ROW, THE BOUNDING BOX DETECTION RESULTS FROM POINTPILLARS [52] ARE EMPLOYED TO GENERATE OBJECT CLUSTERS AS INPUT TO DATA ASSOCIATION.

| Method                   | sAMOTA [%] $\uparrow$ | AMOTA [%] $\uparrow$ | AMOTP [%] $\uparrow$ | MOTA [%] $\uparrow$ | MODA [%] $\uparrow$ | MT [%] $\uparrow$ | ML [%] $\downarrow$ |
|--------------------------|-----------------------|----------------------|----------------------|---------------------|---------------------|-------------------|---------------------|
| Replace ODM with PP [52] | 53.12                 | 17.47                | 53.65                | 41.66               | 42.32               | 24.21             | 38.95               |
| Remove MEM               | 68.45                 | 26.08                | 51.23                | 62.32               | 71.27               | 38.24             | 16.18               |
| Remove the velocity      | 31.50                 | 5.63                 | 16.83                | 24.17               | 32.45               | 11.76             | 30.88               |
| <b>RaTrack</b>           | <b>74.16</b>          | <b>31.50</b>         | <b>60.17</b>         | <b>67.27</b>        | <b>77.83</b>        | <b>42.65</b>      | <b>14.71</b>        |




Fig. 3. Result comparison between RaTrack and baselines on varying valid object thresholds (both predicted and ground truth objects are filtered) and point-based IoU thresholds.

is less optimal than anticipated. We credit this to the incompatibility of the bounding box detector with our framework. Specifically, the object clusters derived from bounding boxes might have more noisy points, which further interferes with our data association that utilizes point-level features.

**Motion estimation module.** The removal of our MEM yields the decrease of 5.7%, 5.4%, 6.6% on sAMOTA, AMOTA and MODA. Such a change in performance demonstrates that our estimated scene flow from MEM can facilitate robust object detection and benefit object temporal matching. Specifically, without MEM, our AMOTP drops by 8.9% which is a non-trivial change in the detection precision. This highlights the importance of our scene flow estimation as an additional motion cue in clustering. Points with similar scene flow vectors are prone to be grouped together.

**Auxiliary velocity feature.** By incorporating the velocity information into the input, a substantial gain is witnessed in all metrics, which demonstrates that the input velocity features are indispensable for RaTrack. Indeed, the velocity features are the key enabler to scene flow estimation and motion segmentation tasks in our network by providing point-level motion information to be encoded in the backbone features. On the other hand, RaTrack can sufficiently utilize such features with its bespoke network designed for 4D radar.

### E. Sensitivity Analysis

**Impact of valid object threshold.** During the evaluation, we ignore invalid objects ( $< 5$  points) from both predictions and ground truth as we are more interested in the objects that are *sufficiently* measured. Here we investigate the impact of this valid object threshold (i.e., the minimum number of points to identify an object as valid) on our evaluation results. As exhibited in Fig. 3 (left), RaTrack consistently outperforms all baseline methods, regardless of the valid object threshold enlarges, and continuously increase its scores as the threshold. We credited this to our cluster-based object detection that can group any number of neighbour points




Fig. 4. The effect of confidence values on the MOTA and the number of FN and FP. As each method has its own range of recall, we normalised their results from 0.25-1.0 for comparison.

into objects for tracking, where objects with more points are more easily being recognized and tracked. In contrast, it is hard for our baseline methods to produce comparable results with 3D bounding box-based detection strategy.

**Impact of point-based IoU threshold.** We also analyze the impact of the point-based IoU threshold (i.e., the minimum IoU to be identified as a true positive sample) for our evaluation results. As seen in Fig. 3 (right), RaTrack achieves the best results on all IoU thresholds with a consistent improvement of  $\sim 10\%$  over baselines, which further confirms the superiority of our pipeline.

**Impact of confidence threshold.** When calculating our metrics, AMOTA and MOTA, the results need to be evaluated repetitively on different confidence thresholds, which further corresponds to different recall values [10], [51]. Then the best MOTA score is reported for comparison and the average MOTA is computed as the AMOTA score. Here we show the impact of the confidence threshold on the MOTA and the number of FN and FP, which are two key components to compute the MOTA. As seen in Fig. 4, RaTrack shows higher MOTA scores over most recall values compared to two augmented baselines. Notably, after ignoring objects with less than five points in evaluation, RaTrack hardly exhibits any FPs, avoiding the compromises in balancing FPs and FNs faced by the baseline methods.

## VI. CONCLUSION

In this work, we unveiled the untapped potential of 4D mmWave radars for multiple moving object tracking. Addressing the challenges of noise and point sparsity in radar data, our approach, RaTrack, offers a fresh perspective on the tracking of moving objects, emphasizing the utility of motion segmentation and clustering over the conventional dependence on specific object types and bounding boxes. This restructured approach not only simplifies the process but also boosts the accuracy of multi-object tracking in complex dynamic scenarios. Extensive evaluations on a public dataset highlight the method's competitive performance.

## REFERENCES

- [1] H. Zhao, J. Gao, T. Lan, C. Sun, B. Sapp, B. Varadarajan, Y. Shen, Y. Shen, Y. Chai, C. Schmid, *et al.*, “Tnt: Target-driven trajectory prediction,” in *Proceedings of the Conference on Robot Learning*. PMLR, 2021, pp. 895–904.
- [2] N. Deo, E. Wolff, and O. Beijbom, “Multimodal trajectory prediction conditioned on lane-graph traversals,” in *Proceedings of the Conference on Robot Learning*. PMLR, 2022, pp. 203–212.
- [3] M. Wei and V. Isler, “Energy-efficient path planning for ground robots by and combining air and ground measurements,” in *Proceedings of the Conference on Robot Learning*. PMLR, 2020, pp. 766–775.
- [4] A. Li, L. Sun, W. Zhan, M. Tomizuka, and M. Chen, “Prediction-based reachability for collision avoidance in autonomous driving,” in *Proceedings of the IEEE International Conference on Robotics and Automation*. IEEE, 2021, pp. 7908–7914.
- [5] J. Lin, H. Zhu, and J. Alonso-Mora, “Robust vision-based obstacle avoidance for micro aerial vehicles in dynamic environments,” in *Proceedings of the IEEE International Conference on Robotics and Automation*. IEEE, 2020, pp. 2682–2688.
- [6] H. Zhang, H. Jin, Z. Liu, Y. Liu, Y. Zhu, and J. Zhao, “Real-time kinematic control for redundant manipulators in a time-varying environment: Multiple-dynamic obstacle avoidance and fast tracking of a moving object,” *IEEE Transactions on Industrial Informatics*, vol. 16, no. 1, pp. 28–41, 2019.
- [7] C. Petres, Y. Pailhas, P. Patron, Y. Petillot, J. Evans, and D. Lane, “Path planning for autonomous underwater vehicles,” *IEEE Transactions on Robotics*, vol. 23, no. 2, pp. 331–341, 2007.
- [8] R. Yonetani, T. Taniyai, M. Barekatin, M. Nishimura, and A. Kanazaki, “Path planning using neural a\* search,” in *Proceedings of the International Conference on Machine Learning*. PMLR, 2021, pp. 12 029–12 039.
- [9] H. Inotsume, T. Kubota, and D. Wettergreen, “Robust path planning for slope traversing under uncertainty in slip prediction,” *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3390–3397, 2020.
- [10] X. Weng, J. Wang, D. Held, and K. Kitani, “3d multi-object tracking: A baseline and new evaluation metrics,” in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2020, pp. 10 359–10 366.
- [11] 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.
- [12] H.-k. Chiu, J. Li, R. Amrbrug, and J. Bohg, “Probabilistic 3d multi-modal, multi-object tracking for autonomous driving,” in *Proceedings of the IEEE International Conference on Robotics and Automation*. IEEE, 2021, pp. 14 227–14 233.
- [13] T. Fischer, Y.-H. Yang, S. Kumar, M. Sun, and F. Yu, “Cc-3dt: Panoramic 3d object tracking via cross-camera fusion,” in *Proceedings of the Conference on Robot Learning*, 2022.
- [14] H.-N. Hu, Y.-H. Yang, T. Fischer, T. Darrell, F. Yu, and M. Sun, “Monocular quasi-dense 3d object tracking,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 2, pp. 1992–2008, 2022.
- [15] F. Ding, C. Fu, Y. Li, J. Jin, and C. Feng, “Automatic failure recovery and re-initialization for online uav tracking with joint scale and aspect ratio optimization,” in *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2020, pp. 5970–5977.
- [16] A. Shenoii, 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 *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2020, pp. 10 335–10 342.
- [17] A. Kim, A. Ošep, and L. Leal-Taixé, “Eagermot: 3d multi-object tracking via sensor fusion,” in *Proceedings of the IEEE International Conference on Robotics and Automation*. IEEE, 2021, pp. 11 315–11 321.
- [18] Y. Zeng, C. Ma, M. Zhu, Z. Fan, and X. Yang, “Cross-modal 3d object detection and tracking for auto-driving,” in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2021, pp. 3850–3857.
- [19] 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.
- [20] A. Palffy, E. Pool, S. Baratam, J. F. Kooij, and D. M. Gavrila, “Multi-class road user detection with 3+ 1d radar in the view-of-delft dataset,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4961–4968, 2022.
- [21] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, “Vision meets robotics: The kitti dataset,” *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1231–1237, 2013.
- [22] A. Milan, L. Leal-Taixé, I. Reid, S. Roth, and K. Schindler, “Mot16: A benchmark for multi-object tracking,” *arXiv preprint arXiv:1603.00831*, 2016.
- [23] G. Ciaparrone, F. L. Sánchez, S. Tabik, L. Troiano, R. Tagliaferri, and F. Herrera, “Deep learning in video multi-object tracking: A survey,” *Neurocomputing*, vol. 381, pp. 61–88, 2020.
- [24] Z. Wang, L. Zheng, Y. Liu, Y. Li, and S. Wang, “Towards real-time multi-object tracking,” in *Proceedings of the European Conference on Computer Vision*. Springer, 2020, pp. 107–122.
- [25] P. Voigtlaender, M. Krause, A. Ošep, J. Luiten, B. B. G. Sekar, A. Geiger, and B. Leibe, “Mots: Multi-object tracking and segmentation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 7942–7951.
- [26] T. Meinhardt, A. Kirillov, L. Leal-Taixe, and C. Feichtenhofer, “Trackformer: Multi-object tracking with transformers,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 8844–8854.
- [27] J. Ku, M. Mozifian, J. Lee, A. Harakeh, and S. L. Waslander, “Joint 3d proposal generation and object detection from view aggregation,” in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2018, pp. 1–8.
- [28] B. Zhu, Z. Jiang, X. Zhou, Z. Li, and G. Yu, “Class-balanced grouping and sampling for point cloud 3d object detection,” *arXiv preprint arXiv:1908.09492*, 2019.
- [29] S. Shi, X. Wang, and H. Li, “Pointcnnc: 3d object proposal generation and detection from point cloud,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 770–779.
- [30] 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*, June 2018.
- [31] W. Shi and R. Rajkumar, “Point-gnn: Graph neural network for 3d object detection in a point cloud,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 1711–1719.
- [32] E. Baser, V. Balasubramanian, P. Bhattacharyya, and K. Czarnecki, “Fantrack: 3d multi-object tracking with feature association network,” in *IEEE Intelligent Vehicles Symposium*. IEEE, 2019, pp. 1426–1433.
- [33] J. Pöschmann, T. Pfeifer, and P. Protzel, “Factor graph based 3d multi-object tracking in point clouds,” in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2020, pp. 10 343–10 350.
- [34] X. Weng, Y. Wang, Y. Man, and K. M. Kitani, “Gnn3dmot: Graph neural network for 3d multi-object tracking with 2d-3d multi-feature learning,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 6499–6508.
- [35] N. Benbarka, J. Schröder, and A. Zell, “Score refinement for confidence-based 3d multi-object tracking,” in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2021, pp. 8083–8090.
- [36] Z. Pang, Z. Li, and N. Wang, “Simpletrack: Understanding and rethinking 3d multi-object tracking,” *arXiv preprint arXiv:2111.09621*, 2021.
- [37] 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.
- [38] T. Wen, Y. Zhang, and N. M. Freris, “Pf-mot: Probability fusion based 3d multi-object tracking for autonomous vehicles,” in *Proceedings of the International Conference on Robotics and Automation*. IEEE, 2022, pp. 700–706.
- [39] A. Kim, G. Brasó, A. Ošep, and L. Leal-Taixé, “Polarmot: How far can geometric relations take us in 3d multi-object tracking?” in *Proceedings of the European Conference on Computer Vision*. Springer, 2022, pp. 41–58.
- [40] T. Sadjadpour, J. Li, R. Amrbrug, and J. Bohg, “Shasta: Modeling shape and spatio-temporal affinities for 3d multi-object tracking,” *arXiv preprint arXiv:2211.03919*, 2022.

- [41] M. Chaabane, P. Zhang, J. R. Beveridge, and S. O'Hara, "Deft: Detection embeddings for tracking," *arXiv preprint arXiv:2102.02267*, 2021.
- [42] X. Zhou, V. Koltun, and P. Krähenbühl, "Tracking objects as points," in *Proceedings of the European Conference on Computer Vision*. Springer, 2020, pp. 474–490.
- [43] H. W. Kuhn, "The hungarian method for the assignment problem," *Naval Research Logistics Quarterly*, vol. 2, no. 1-2, pp. 83–97, 1955.
- [44] G. Brasó and L. Leal-Taixé, "Learning a neural solver for multiple object tracking," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 6247–6257.
- [45] R. Sinkhorn and P. Knopp, "Concerning nonnegative matrices and doubly stochastic matrices," *Pacific Journal of Mathematics*, vol. 21, no. 2, pp. 343–348, 1967.
- [46] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, "Pointnet++: Deep hierarchical feature learning on point sets in a metric space," *Proceedings of the Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [47] W. Wu, Z. Y. Wang, Z. Li, W. Liu, and L. Fuxin, "Pointpwc-net: Cost volume on point clouds for (self-) supervised scene flow estimation," in *Proceedings of the European Conference on Computer Vision*. Springer, 2020, pp. 88–107.
- [48] K. Cho, B. van Merriënboer, D. Bahdanau, and Y. Bengio, "On the Properties of Neural Machine Translation: Encoder–Decoder Approaches," in *Proceedings of the Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*, 2014, pp. 103–111.
- [49] M. Ester, H.-P. Kriegel, J. Sander, X. Xu, *et al.*, "A density-based algorithm for discovering clusters in large spatial databases with noise," in *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, vol. 96, no. 34, 1996, pp. 226–231.
- [50] L. Leal-Taixé, A. Milan, I. Reid, S. Roth, and K. Schindler, "MOTchallenge 2015: Towards a benchmark for multi-target tracking," *arXiv preprint arXiv:1504.01942*, 2015.
- [51] K. Bernardin and R. Stiefelwagen, "Evaluating multiple object tracking performance: the clear mot metrics," *EURASIP Journal on Image and Video Processing*, vol. 2008, pp. 1–10, 2008.
- [52] 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.
- [53] S. A. Baur, D. J. Emmerichs, F. Moosmann, P. Pinggera, B. Ommer, and A. Geiger, "SLIM: Self-Supervised LiDAR Scene Flow and Motion Segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 13 126–13 136.
- [54] F. Ding, Z. Pan, Y. Deng, J. Deng, and C. X. Lu, "Self-supervised scene flow estimation with 4-d automotive radar," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 8233–8240, 2022.
- [55] F. Ding, A. Palffy, D. M. Gavrila, and C. X. Lu, "Hidden gems: 4d radar scene flow learning using cross-modal supervision," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 1–10.
- [56] P. Jund, C. Sweeney, N. Abdo, Z. Chen, and J. Shlens, "Scalable scene flow from point clouds in the real world," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 1589–1596, 2021.
- [57] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.