



Delft University of Technology

**Document Version**

Final published version

**Licence**

Dutch Copyright Act (Article 25fa)

**Citation (APA)**

Liu, X., & Caesar, H. (2024). Offline Tracking with Object Permanence. In *Proceedings of the IEEE Intelligent Vehicles Symposium (IV 2024)* (pp. 1272-1279). (IEEE Intelligent Vehicles Symposium, Proceedings). IEEE. <https://doi.org/10.1109/IV55156.2024.10588572>

**Important note**

To cite this publication, please use the final published version (if applicable).

Please check the document version above.

**Copyright**

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.

Unless copyright is transferred by contract or statute, it remains with the copyright holder.

**Sharing and reuse**

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

**Takedown policy**

Please contact us and provide details if you believe this document breaches copyrights.

We will remove access to the work immediately and investigate your claim.

*This work is downloaded from Delft University of Technology.*

# Offline Tracking with Object Permanence

Xianzhong Liu<sup>1</sup> and Holger Caesar<sup>1</sup>

**Abstract**—To reduce the expensive labor costs of manually labeling autonomous driving datasets, an alternative is to automatically label the datasets using an offline perception system. However, objects might be temporarily occluded. Such occlusion scenarios in the datasets are common yet underexplored in offline auto labeling. In this work, we propose an offline tracking model that focuses on occluded object tracks. It leverages the concept of object permanence, which means objects continue to exist even if they are not observed anymore. The model contains three parts: a standard online tracker, a re-identification (Re-ID) module that associates tracklets before and after occlusion, and a track completion module that completes the fragmented tracks. The Re-ID module and the track completion module use the vectorized lane map as a prior to refine the tracking results with occlusion. The model can effectively recover the occluded object trajectories. It significantly improves the original online tracking result, demonstrating its potential to be applied in offline auto labeling as a useful plugin to improve tracking by recovering occlusions.

## I. INTRODUCTION

Supervised deep learning-based models have achieved good performance in autonomous driving. However, it usually requires a huge amount of labeled data with high quality to train and tune such data-hungry models. An effective way is to auto label datasets, where labels can be automatically provided by a trained perception system. Waymo first proposed to auto label data offline to improve the quality of the generated labels [17]. In online tracking, the location of an object is inferred only from past and present sensor data. Online trackers are thus likely to produce false associations under severe occlusions. Offline multi-object tracking (MOT) is acausal and the position of an object can be inferred from past, present, and future sensor data. A consistent estimate of the scene can thus be optimized globally using the data not limited to a short moment in the past, enabling accurate object tracking even under severe occlusions. Based on global information, [25, 17, 8, 15] have developed offline auto labeling pipelines that generate accurate object trajectories in 3D space from LiDAR point cloud sequence data.

Previous offline labeling methods mainly focused on using the observation sequence to refine the unoccluded trajectories (e.g. location, orientation, and size of the bounding boxes). However, object trajectories and point cloud sequences are sometimes partially missing because of the temporary occlusions on the objects. Severely occluded objects are likely to be missed by detectors and thus cannot be continuously tracked. Furthermore, the trackers in [25, 17] adopt a death memory strategy to terminate the unmatched tracks, which

means that a track not matched with any detections for a predefined number of frames will be terminated. Therefore, the identity of the same object may switch after reappearance from occlusion. In [21], the tracker learns object permanence from synthetic data by hallucinating the occluded motion. The authors use pseudo-labels with constant velocity to supervise the full occlusion scenarios, which fails to capture the nonlinear motion behind occlusions. The direct supervision using ground truth (GT) labels yields a worse performance than using the constant-velocity pseudo-labels. Without any extra prior knowledge, a neural network is not able to learn the complex motion and instead resorts to using a linear velocity baseline. Immortal Tracker [23] also captures object permanence by not terminating any unmatched tracks and prolonging their predictions. It has been applied in [8, 15] for its effectiveness in reducing identity switches (IDS). However, the data association is still simply done by using the constant-velocity prediction from a Kalman filter [11] which cannot capture the nonlinear motion under long occlusions either. Thus, tracking under long occlusion remains a challenge in point cloud-based offline auto labeling. Motion prediction models [13, 9], on the other hand, can produce accurate vehicle trajectories over longer horizons based on a semantic map. The lanes on the map serve as a strong prior knowledge to guide the motion of target vehicles and thus can be used to estimate motion under occlusion. Based on this insight, we not only associate tracklets together accurately but also complete the fragmented tracks. Unlike Immortal Tracker [23], we additionally interpolate and refine the occluded trajectory based on the lane map feature.

In this work, we propose an offline tracking model containing three modules, as shown in Fig. 1 to tackle the aforementioned problems. The first module is an off-the-shelf online tracker which produces the initial tracking result. Leveraging the idea of object permanence, the second Re-ID module tries to reassociate the terminated tracklets with the possible future candidate tracklets. Based on the association result, the last track completion module completes the missing trajectories. Taking the insight from the motion prediction task, we have extracted map information as a prior to enhance the association and track completion modules. Unlike common motion prediction methods with a fixed prediction horizon predefined [5, 10], our model can use a flexible prediction horizon at inference time to decode the predicted poses in the track completion module, to deal with variable occlusion durations. Compared to goal-based motion prediction methods [10], we take future tracklets as inputs instead of goal points, which enables the model to predict more accurately based on motion patterns. Previous offline

<sup>a</sup> Intelligent Vehicles section, Delft University of Technology, The Netherlands

auto-labeling methods primarily prioritize the precision of visible bounding boxes as their main metric, with MOT metrics receiving minor attention. In contrast, our work focuses on optimizing and evaluating MOT metrics.

The contributions of this paper are summarized as follows:

- We propose an offline tracking model for Re-ID (Sect. III-A) and occluded track completion (Sect. III-B) to track occluded vehicles.
- We are the first to utilize vectorized lane maps in the MOT task. By using these maps as a prior, we are able to recover the complex motion under occlusion.
- By encoding variable time queries (Sect. III-B), we present the first method that can decode future trajectories of variable duration, accommodating occlusions of differing lengths.
- We optimize our method for the MOT task and demonstrate the improvements relative to the original online tracking results on the nuScenes dataset (Sect. V-A), showing the potential of our model to be applied in offline auto labeling as a useful plugin to improve tracking and recover occlusions.
- The code and pre-trained models are publicly available at [github.com/tudelft-iv-students/offline-tracking-with-object-permanence](https://github.com/tudelft-iv-students/offline-tracking-with-object-permanence).

## II. RELATED WORK

### A. Tracking under occlusion

Estimating a target’s motion under occlusion is one of the main challenges in the MOT task. A common method adopts a constant velocity [24, 23] transition model and uses the Kalman filter [11] to update the tracks when detections are missed due to occlusion. CenterPoint [26] directly estimates displacement offset (or velocity) between frames. The unmatched track is updated with a constant velocity model. PermaTrack [21] supervises the occluded motion with pseudo-labels that keep constant velocity. Those methods heavily rely on motion heuristics. Such heuristics work well for short occlusion but cannot capture the nonlinear motion under long occlusion. Due to the large deviation error accumulated over time, the prediction of the unmatched track would not be associated with the correct detection. Most tracking models that have finite death memory [26, 24] would terminate such unmatched tracks, even if the object is detected again later. Such premature termination would lead to IDS. Immortal Tracker [23] effectively reduces such IDS by extending the life of unmatched tracks forever. Our model does not rely on any handcrafted motion heuristics to associate. Given any pair of future and history tracklets, the model outputs the learned affinity scores directly.

### B. Offline auto labeling from lidar point clouds

Manually annotating lidar point cloud data takes much effort due to the sparsity of the data and the temporal correlation of the sequence. Several works have attempted to tackle this problem by automatically labeling the dataset [3, 1, 7]. Auto4D and 3DAL [25, 17] fully automate annotation by taking initial tracklets generated from an online tracker

(i.e. AB3DMOT [24]). Point cloud and detection features are extracted to refine the initial tracks. Given full observation, such refinement utilizes future information to perform global optimization offline. Though these methods can produce accurate bounding boxes, their data association is still simply done online. The global information available in the offline setting has yet to be utilized in tracking. Such methods are likely to produce IDS and inconsistent tracks under occlusion. CTRL [8] uses the Immortal Tracker [23] to associate fragmented tracks due to long occlusion or missing detection. Then it backtraces the tracks by extrapolation to track the missing detections. DetZero [15] also uses the Immortal Tracker to perform bidirectional tracking on the time horizon and then ensemble the results with forward and reverse order tracking fusion. As a result, CTRL and DetZero have achieved better results which surpass the human annotation. However, their data association still relies on a Kalman filter with the constant-velocity model.

### C. Map-based prediction

Semantic maps are widely used as an input in motion prediction methods. As it contains accurate scene context from the surrounding environment, including road lanes, crosswalks, etc. With such information, models can predict how target vehicles navigate in the environment over a long horizon. HOME [10] rasterizes the map as multiple vector layers with distinct RGB values representing different map elements. Alternatively, VectorNet [9] proposes the vectorized approach which represents curves as vectorized polylines. It uses a subgraph network to encode each polyline as a node in a fully connected global interaction graph. It achieves better performance over the CNN baseline and reduces the model size. To capture higher resolution, LaneGCN [13] uses polyline segments as map nodes. It models a sparsely connected graph following the map topology. Similarly, PGP [5] constrains the connected edges such that any traversed path through the graph corresponds to a legal route for a vehicle.

## III. METHOD

As shown in Fig. 1, our model initially takes the detections from a detector as input. Then it uses an off-the-shelf online tracker to associate detections and generate initial tracklets. Next, the Re-ID module tries to associate the possible future tracklets with the terminated history tracklet. If a pair of tracklets are matched, the track completion module interpolates the gap between them by predicting the location and orientation of the missing box at each timestamp. Both modules extract motion information and lane map information to produce accurate results. The model finally outputs the track with refined new identities and completes the missing segments within the tracks.

### A. Re-ID

Based on the concept of object permanence, we assume any tracklets could potentially be terminated prematurely due to occlusion, and any other tracklets that appeared after the termination could be matched. The Re-ID model has two branches, the motion and the map branch. The motion branch takes tracklet motion as input and outputs affinity

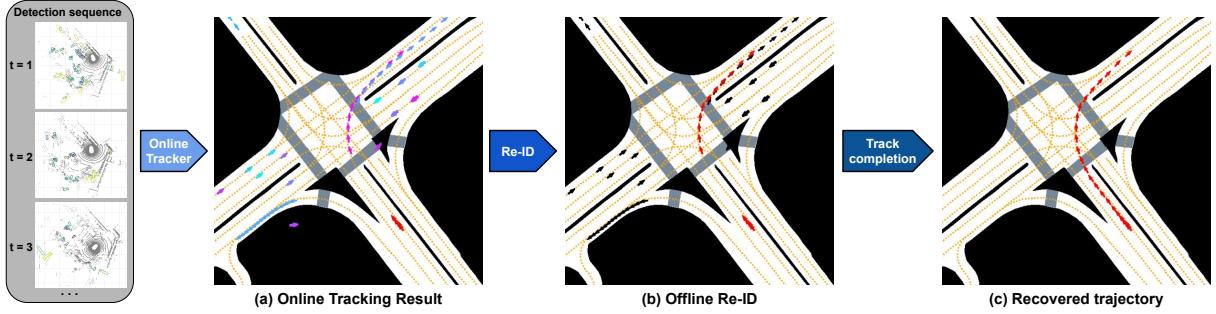


Fig. 1: A brief overview of the offline tracking model. **(a) Online tracking result:** Each tracklet is represented by a different color (the history tracklet is red). **(b) Offline Re-ID:** The matched pair of tracklets are red. The unmatched ones are black. **(c) Recovered trajectory.**

scores based on vehicle dynamics. The map branch takes motion and the lane map as inputs and outputs affinity scores based on map connectivity. Thus, for each target history tracklet, the model computes motion affinity and map-based affinity scores for it with all the possible future tracklets. The problem can thus be treated as a **binary classification task** for each single matching pair. The final score is a weighted sum of the motion and map affinity scores. We perform bipartite matching by greedily associating history tracklets with their future candidate tracklets based on the final scores. The matching pair which has a tracking score lower than a threshold is excluded from the association.

1) *Re-ID formulation:* The upstream online tracker provides a set of tracklets  $\mathbf{T}$  in the scene. In  $\mathbf{T}$ , history tracklets are the tracklets that end earlier before the last frame of the scene. They are represented as  $\mathbf{H} = [H_1, \dots, H_n]$ . For the  $i$ -th history tracklet  $H_i$ , it contains history observations of  $[h_{-T_i}^i, \dots, h_0^i]$ , where  $h_0^i$  is the last observation of  $H_i$  on the temporal horizon and  $h_{-T_i}^i$  the first observation. Each observation contains the location, orientation, size, uncertainty, and velocity. It is defined as:

$$h_t^i = (x_t^i, y_t^i, \theta_t^i, l_t^i, w_t^i, h_t^i, s_t^i, v_{xt}^i, v_{yt}^i) \quad (1)$$

where  $l, w, h$  represent the size of the  $j$ -th bounding box.  $s$  is the confidence score, and  $v_x, v_y$  the velocities on  $x, y$  directions. For each history tracklet  $H_i$ , it has a set of possible future candidate tracklets for matching:  $\mathbf{F}^i = [F_1^i, \dots, F_n^i]$ . Each future tracklet  $F_j^i$  starts after the termination of  $H_i$ :  $F_j^i = [f_{t_j}^i, \dots, f_{t_j+T_j}^i]$ , s.t.  $t_j > \tau$ . In practice, we set  $\tau$  as the horizon of the death memory of the online tracker since the GT future tracklets reappearing within such period are likely to be associated by the online tracker. Tracklet pairs violating this constraint will not be considered for the association. Each tracklet has a feature dimension of  $T \times 8$ , where  $T$  is the length of the tracklet. The 8 features correspond to  $[x, y, \theta, t, \cos(\theta), \sin(\theta), v_x, v_y]$ , where  $x, y$  are the local BEV coordinates,  $\theta$  the yaw angle,  $t$  the time relative to the last observation of the history tracklet and  $v_x, v_y$  the velocities on  $x$  and  $y$  direction.

The model also extracts information from the fully connected lane graph  $\mathcal{G} = (V, E)$  where each node in  $V$  is a lane section.  $E$  comprises edges connecting all node

pairs. Each lane node is encoded from a lanelet, which is a section of the lane with a maximal length of 20m. A lanelet is represented as several discrete lane poses with a resolution of 1m. Each lane pose contains the feature of  $[x_{lane}, y_{lane}, \theta_{lane}, \cos(\theta_{lane}), \sin(\theta_{lane}), \mathcal{D}, \mathcal{L}_{lane}]$ , where  $x_{lane}, y_{lane}$  are the BEV coordinates of the lane poses in the local frame,  $\theta_{lane}$  the yaw angle,  $\mathcal{D}$  the binary flag indicating whether the lane section ends. Following PGP [5], we have also included  $\mathcal{L}_{lane}$ , a 2-D binary vector indicating whether the pose lies on a stop line or crosswalk. The graph has a feature with the dimension of  $N_{lane} \times l \times 8$ , where  $N_{lane}$  is the number of lane nodes and  $l$  the length of a lanelet.

Given a history tracklet  $H_i$  and a future tracklet  $F_j^i$  and lane graph  $\mathcal{G}$ , the model outputs motion affinity scores and map-based affinity scores:

$$\begin{aligned} C_{motion}^{i,j} &= Net_{motion}(H_i, F_j^i) \in [0, 1] \\ C_{map}^{i,j} &= Net_{map}(H_i, F_j^i, \mathcal{G}) \in [0, 1] \end{aligned} \quad (2)$$

where  $Net_{motion}$  and  $Net_{map}$  represent the networks of the motion branch and map branch respectively. The final matching score for the association  $C^{i,j}$  is a weighted sum of  $C_{motion}^{i,j}$  and  $C_{map}^{i,j}$ . Therefore, we can construct a matching score matrix  $\mathbf{C} \in n \times N$  for tracklet association, where  $N$  is the number of tracklets  $\mathbf{T}$  in the scene and  $n$  the number of history tracklets  $\mathbf{H}$ . Each row in  $\mathbf{C}$  represents a history tracklet and each column represents a future tracklet. Based on the established affinity matrix, bipartite matching is formulated as a linear assignment problem such that the sum of the matching scores is maximized.

2) *Motion affinity:* We first transform the trajectories from the global frame to the local agent frame, where the origin aligns with the pose of the last observation of the history tracklet. The motion feature of each tracklet contains the location, orientation, time and velocity information at every time step. Given one history tracklet and all its possible future matching candidate tracklets, the model takes the motion features as input and outputs affinity scores in the range of  $[0, 1]$  for all future tracklet candidates. As in Fig. 3, the motion affinity branch comprises MLPs and GRUs. The GRU encoder encodes the history motion and outputs the last hidden state as history encoding. The history encoding is then used as the initial hidden state for the UGRU [18] to

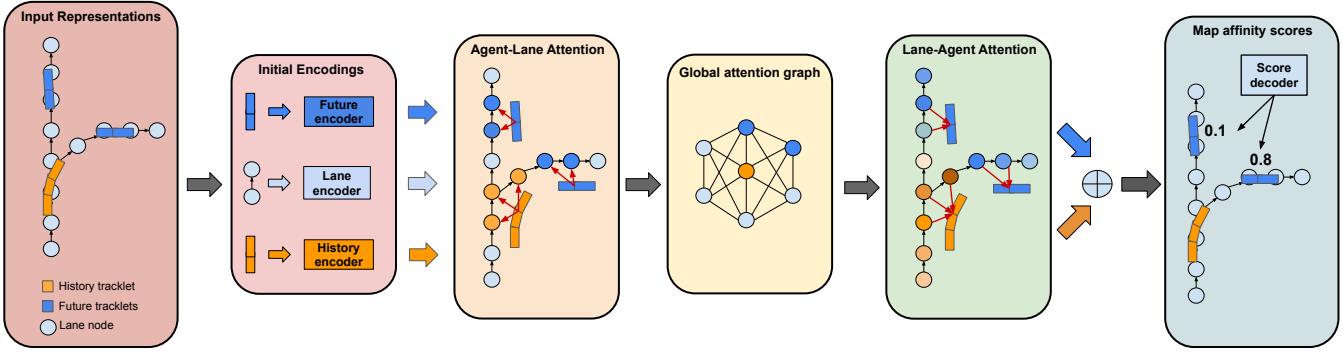


Fig. 2: A brief overview of the map branch. The branch starts with three parallel encoders which encode the future tracklet, lanes and history tracklet respectively. The model then propagates information between tracklets and the lane map by performing attention. Finally, map-based affinity scores are decoded from the tracklet features.

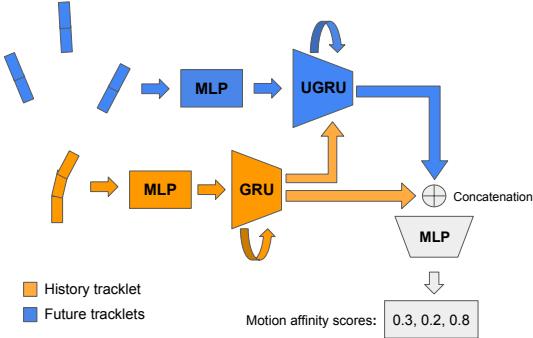


Fig. 3: The network structure of motion affinity branch. The history tracklet encoder is orange, whereas the future tracklets encoder is blue. Three possible future candidates correspond to the three outputted motion affinity scores.

encode future motions, enabling the awareness of historical information when encoding future motions. Unlike typical U-RNN used in motion prediction for encoding history motion features, our UGRU first does a forward pass and then does a backward pass in reverse order to encode the future motion features. The future motion encodings are concatenated with the history encoding to decode motion affinity scores.

3) *Map-based affinity*: To improve the association accuracy, we further extract information from the lane graph as another branch. The lane centerlines provide information about the direction of traffic flow and the possible path for the vehicles to follow, which can be utilized for occluded motion estimation over a long horizon. We represent the lane map as a vectorized graph as previously described in Sect. III-A.1. Compared with the rasterized representation in [10], this is a more efficient representation due to the sparsity of the map. We thus can easily extend the size of the map to incorporate all possible future tracklets far away in case of long occlusions. As in Fig. 2, the map branch starts by encoding the history and the future tracklets, using the same encoder previously introduced in the motion branch (Sect. III-A.2). The lane pose features are also encoded in parallel using a single-layer MLP. The model then performs agent-to-lane attention to fuse the information from the tracklets to the lanes. It selectively performs attention such that

the tracklet encodings are only aggregated to the surrounding lanes nearby. We used the spatial attention layer in [13] to aggregate the feature from tracklets to lane poses. The model aggregates the feature from each sub-node lane pose to get the lane node features. We apply a 2-layer-GRU to aggregate sub-node features. The first GRU layer is bidirectional so that the information can flow in both directions. Having the lane node features carrying the information of the tracklets, the model performs global attention [22] to propagate information globally in the graph. Next, the model aggregates the updated lane features back to tracklets by applying a lane-to-agent attention layer [13]. Finally, the future and history tracklet encodings updated with global map information are concatenated together. An MLP decodes map affinity scores from the concatenated features.

4) *Association*: To match  $n$  history tracklets with  $N$  future tracklets, we have an affinity score matrices  $C_{motion}$  and  $C_{map} \in n \times N$  from the two branches, we first filter the matching pairs whose affinity scores are lower than a threshold, then we get the final matching score matrix:

$$C = w \cdot C_{map} + (1 - w) \cdot C_{motion} \quad (3)$$

where  $w \in [0, 1]$  is a scalar weight. During experiments, we simply set  $w$  to 0.5. In practice, it can be optimized using validation data or decided based on domain knowledge about the relative reliability of motion versus map-based affinity scores.  $C$  is the weighted sum of the two affinity matrices. Finally, we solve the bipartite problem by maximizing the overall matching scores. Due to the sparsity of  $C$ , we simply perform a greedy matching.

#### B. Track completion

With the upstream Re-ID results, we have associated multiple tracklet pairs together. Therefore, we now have multiple fragmented tracks which have a missing segment in the middle. The track completion model interpolates the gaps within the tracks. Based on the history and future motion, the model first generates an initial trajectory in between. Then, a refinement head is applied to refine the initial trajectory based on motion and map features. The trajectory completion model learns a **regression task**.

We first transform the future and history tracklets from the global frame to the local agent frame. We set the origin as the midpoint of the line connecting the two endpoints of the missing segment. Most prediction models generally have a fixed prediction horizon [13, 5]. But the occlusion horizon in reality is never fixed. We thus use time features as input queries to decode poses at all the target timesteps. Specifically, each time query feature at time  $t$  is given as  $[t, t/T]$ , where  $T$  is the total prediction horizon. So that we can have a variable number of time queries corresponding to a variable prediction horizon.

1) *Track completion formulation*: Following the formulation in Sect. III-A.1, the model receives a history observation sequence  $H_i = [h_{-T_i}^i, \dots, h_0^i]$ , and a matched future observation sequence  $F_i = [f_{t_i}^i, \dots, f_{t_i+T'_i}^i]$ , where  $t_i$  is the occlusion horizon of the  $i$ -th trajectory,  $T_i$  and  $T'_i$  the lengths of history and future sequence respectively. The model predicts a motion sequence  $\mathcal{P}_i = [p_1^i, \dots, p_{t_i-1}^i]$  in between. Here, we choose to output the  $x, y$  coordinates in the BEV plane, and the yaw angle  $\theta$ .

$$p_t^i = Net_{completion}(F_i, H_i, \mathcal{G}, t) = [x_t^i, y_t^i, \theta_t^i] \quad (4)$$

where  $Net_{completion}$  represents the network of the track completion model. We use the same motion and map features as in the previous Sect. III-A.1, with the exception that we do not use the velocity explicitly. This results in a feature dimension of  $T_{input} \times 6$ . Given that the start and end pose of the predicted are already known, the model can implicitly predict the velocity in between. Hence, we discard the velocity features to reduce the possible noise from the imperfect online tracking result.

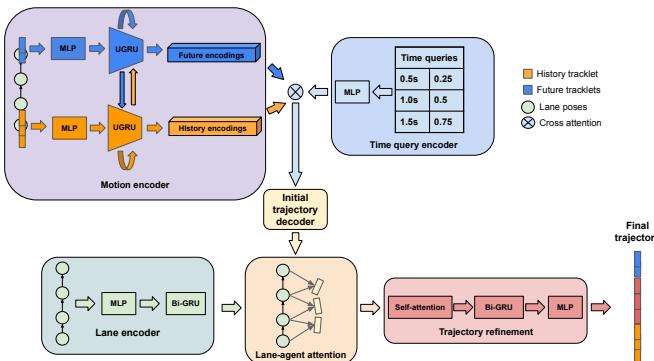


Fig. 4: A brief overview of the track completion model. The model stacks several sub-modules together to sequentially aggregate features and refine the results.

2) *Track completion network*: The overall structure of the track completion model is shown in Fig. 4. It encodes time features as queries and iteratively aggregates context information to the query features to refine the generated tracks. The motion encoder first encodes the history and future motion features using two parallel encoders. The UGRU in each encoder takes the final hidden state from the other, hence the information can propagate through the

temporal gap. The time features are encoded with a single-layer MLP as initial queries. A cross-attention layer then aggregates the motion information to the time queries. We also add a skip connection to concatenate the attention output with the motion encoding and time features. To improve the result, we added a refinement phase to the initial trajectory. The model encodes lane features and performs lane-to-agent attention [13] to fuse the extracted lane information with the generated trajectory information. The trajectory encodings carrying the lane information then go through the final trajectory refinement head which uses a self-attention layer with a skip connection to propagate information within the track. The following bidirectional GRU implicitly smooths the track by aggregating features from each pose to its adjacent poses. Finally, an MLP decodes the refined trajectory.

### C. Loss function

We adopt the focal loss [14] to train the Re-ID model. with  $\alpha_t$  and  $\gamma$  set as 0.5 and 2.0 respectively. For the track completion model, we use Huber loss (i.e. smooth L1 loss) as in [13] to train the  $x, y$  coordinate regression heads in both the initial trajectory decoder and the final refinement blocks. For yaw angle regression, we first adjust all the GT yaw angles such that their absolute differences with the predicted yaw angles are less than  $\pi$ . The final regression loss  $\mathcal{L}_{reg}$  is:

$$\mathcal{L}_{reg} = \alpha_{coord} \mathcal{L}_{coord} + \alpha_{yaw} \mathcal{L}_{yaw} \quad (5)$$

where  $\alpha_{coord}$  and  $\alpha_{yaw}$  are the weights for the coordinate regression loss  $\mathcal{L}_{coord}$  and yaw angle regression loss  $\mathcal{L}_{yaw}$  respectively. We set  $\alpha_{coord} = 1.0$  and  $\alpha_{yaw} = 0.5$ .

## IV. EXPERIMENT SETUP

### A. Training setup

**Training data:** While [17, 25] train their methods using online tracker outputs, we train only using the GT, but additionally inject pseudo-occlusions by masking a partial segment of each GT track. This gives two advantages: 1)

**More accurate inputs:** Unlike the imperfect trajectories generated from an online tracker, the trajectories from the available autonomous driving datasets are accurate and contain little noise. With proper data augmentation, our model can still deal with imperfect data when inferencing with the trajectories produced from an online tracker. 2) **More training samples:** Since we have longer intact tracks as input, we can generate data with longer occlusions.

**Pseudo-occlusion:** We randomize the length of the pseudo-occlusion for each track during the training of the Re-ID and track completion models. The pseudo-occlusion duration is selected such that the history tracklet and all the future tracklets have at least one observation.

**Data augmentation:** To address the potential imperfections from the online tracking results, we augment the training data by randomly rotating the local frame in each sample and adding random noise to the motion inputs.

### B. Evaluation setup

We train and evaluate our model on the nuScenes dataset [2]. The standard MOT evaluation on nuScenes could punish the fully occluded predictions since the evaluation

code filters out all the GT boxes with no points inside and then linearly interpolates the GT tracks. Therefore, some of the bounding boxes recovered from occlusion could be regarded as false positives (FP) as the GT boxes they are supposed to be matched with are linearly interpolated. We thus train and validate our model on the training split of the nuScenes dataset. The evaluation is separately done on the validation split and the test split so that we can not only benchmark our method in a standard way but also adjust the evaluation procedure for occlusion cases. In each of the experiments below, we use one of the following three evaluation setups as indicated.

**Official nuScenes setup:** Following the official nuScenes MOT evaluation protocol, we filter out all the empty GT boxes without any lidar or radar points inside. We thus focus on the visible GT boxes. Therefore, this evaluation can also be applied to the test split on the official server.

**All-boxes setup:** To get a comprehensive evaluation result, we do not filter any GT boxes. Therefore, the evaluation takes both the visible and the occluded boxes into account.

**Pseudo-occlusion setup:** To generate more occlusions with longer duration for evaluation, we take the GT vehicle tracks from the nuScenes prediction validation set and mask a partial segment within each track. Unlike the previous setups where the model takes the online tracking result as input, the model takes the unoccluded GT tracklets as inputs under this setup. In every evaluated sample, all future tracks are masked for random durations to create diverse pseudo-occlusions. The durations range from 1.5s to each track’s maximal length (12.5s at most) such that each future tracklet has at least one last pose visible. The history tracklet in each sample is also randomly deprecated such that it has at least one pose left as input and is at most 2.5 seconds long.

## V. EXPERIMENT

We first evaluate our offline tracking model on the imperfect data generated from the online tracking result. We show the relative improvements it brings to the original online tracking result on MOT metrics. In this setting, the experiments evaluate the potential of the offline tracking model to be applied in offline auto labeling. We use CenterPoint [26] as the initial off-the-shelf detector and tracker. We also perform an intra-class non-maxima suppression with an IoU threshold of 0.1 before the initial tracking. Next, we evaluate our offline tracking model using human-annotated data. We mask the vehicle tracks to create a large amount of pseudo-occlusion cases for evaluation. The human-annotated data have little input noise, which excludes the imperfections from the online tracker so the evaluation focuses solely on the offline tracking model we proposed. Furthermore, the real occlusion cases only take a relatively small portion of the dataset compared to the unoccluded cases and the occlusion duration is typically short. Unlike the real occlusions in the dataset, the generated pseudo-occlusions are abundant and typically have longer durations. We focus our method only on the vehicle classes. All the metric scores reported are derived from the epoch in a single training run that exhibited the lowest validation error during training.

## A. Quantitative evaluation with online tracking results

1) *Re-ID model evaluation:* In the sole evaluation of the Re-ID model, we adopt the **official nuScenes setup** which follows the standard evaluation protocol of nuScenes and filters out the empty GT boxes for evaluation. We only use the Re-ID model to reassociate the tracklets before and after occlusions, and we do not use the track completion model to interpolate the trajectories between the gaps. Instead, the nuScenes evaluation code automatically does a simple linear interpolation to fill all the gaps on both the GT tracks and the predicted tracks. For comparison, we select several top-performing lidar-based trackers that also use CenterPoint [26] detections from the nuScenes leaderboard.

From the results in Tab. I, our model effectively improves the original online tracking result from CenterPoint [26]. Please also note that the Re-ID model only refines the limited occlusion cases. Therefore, the improvement brought by the Re-ID model is upper bound by the number of occlusions. Yet the result after refinement already outperforms the other methods on AMOTA and AMOTP, indicating the Re-ID model can accurately reassociate the tracklets before and after occlusions. We thus have demonstrated the potential of our Re-ID model to improve the offline tracking result during auto labeling, especially for occlusion-rich scenes.

2) *Track completion model evaluation:* Based on the Re-ID results, the track completion model recovers the missing trajectories between all the matched tracklet pairs. Since we want to evaluate the ability of our model to recover the occluded trajectories, we adopt the **all-boxes setup** and thus do not filter out any GT bounding boxes. We take the Re-ID result refined by the motion and the map branch as input. Based on the Re-ID result, the track completion model interpolates the gaps and recovers the missing trajectories. If a gap within a track is spatially larger than 3m or temporally longer than 1.8 seconds, it is considered a possible occlusion, and the track completion model predicts the poses for the missing timestamps in between. Otherwise, the evaluation code automatically does a linear interpolation to fill the small gaps. We show the improvements our model brings to multiple online trackers in Tab. II. We have only tuned our model using the results from CenterPoint tracker [26] and then applied our model to the results from other trackers. The improvements have demonstrated that our model can be used as a general plugin to recover the imperfections caused by occlusions during offline autolabeling. Our model is compatible with any initial tracker, thus it can be easily integrated into the existing offline autolabeling frameworks [17, 25].

## B. Quantitative evaluation with pseudo-occlusions

1) *Re-ID model evaluation:* We evaluate the Re-ID model following the **pseudo-occlusion setup**. The number of matching candidate tracklets ranges from 2 to 65. We also include a constant velocity model (CVM) associator as our baseline. The CVM associator takes the last observable position and velocity as inputs and then predicts future trajectories with a constant velocity. The future tracklet which has the shortest distance to the constant velocity prediction will

TABLE I: Comparison of MOT metrics over the top-performing methods using CenterPoint [26] detections on the nuScenes test split (official nuScenes setup).

| Method                       | AMOTA       |             |             |           |             | AMOTP<br>↓ / m | Recall↑     | MOTA↑       | IDS↓      |
|------------------------------|-------------|-------------|-------------|-----------|-------------|----------------|-------------|-------------|-----------|
|                              | Overall     | Car         | Bus         | Truck     | Trailer     |                |             |             |           |
| CenterPoint [26]             | 69.8        | 82.9        | 71.1        | 59.9      | 65.1        | 0.596          | 73.5        | 59.4        | 340       |
| SimpleTrack [16]             | 70.0        | 82.3        | 71.5        | 58.7      | 67.3        | 0.582          | 73.7        | 58.6        | 259       |
| UVTR [12]                    | 70.1        | 83.3        | 67.2        | 58.4      | <b>71.6</b> | 0.636          | <b>74.6</b> | 59.3        | 381       |
| Immortal Tracker [23]        | 70.5        | 83.3        | 71.6        | 59.6      | 67.5        | 0.609          | 74.5        | 59.9        | 155       |
| NEBP [20]                    | 70.8        | 83.5        | 70.8        | 59.8      | 69          | 0.598          | 74.1        | <b>61.9</b> | <b>93</b> |
| 3DMOTFormer++ [6]            | 72.3        | 82.1        | 74.9        | 62.6      | 69.6        | 0.542          | 73.0        | 58.6        | 210       |
| ShaSTA [19]                  | 73.1        | 83.8        | 73.3        | <b>65</b> | 70.4        | 0.559          | 74.3        | 61.2        | 185       |
| Offline Re-ID (Motion + Map) | <b>73.4</b> | <b>84.2</b> | <b>75.1</b> | 64.1      | 70.3        | <b>0.532</b>   | 74.2        | 61.3        | 204       |

TABLE II: Track completion evaluation on the nuScenes validation split (All-boxes setup). w: with Re-ID and track completion. w/o: without Re-ID and track completion.

| Track Completion | AMOTA↑ |             |             | AMOTP↓ / m |              |               | IDS↓ |            |            | Recall↑ |             |             |
|------------------|--------|-------------|-------------|------------|--------------|---------------|------|------------|------------|---------|-------------|-------------|
|                  | w/o    | w           | Δ           | w/o        | w            | Δ             | w/o  | w          | Δ          | w/o     | w           | Δ           |
| CenterPoint [26] | 70.2   | <b>72.4</b> | <b>+2.2</b> | 0.634      | <b>0.615</b> | <b>-0.019</b> | 254  | <b>183</b> | <b>-71</b> | 73.7    | <b>74.5</b> | <b>+0.8</b> |
| SimpleTrack [16] | 70.0   | <b>71.0</b> | <b>+1.0</b> | 0.668      | <b>0.629</b> | <b>-0.039</b> | 210  | <b>170</b> | <b>-40</b> | 72.5    | <b>72.9</b> | <b>+0.4</b> |
| VoxelNet [4]     | 69.6   | <b>70.6</b> | <b>+1.0</b> | 0.710      | <b>0.665</b> | <b>-0.045</b> | 308  | <b>230</b> | <b>-78</b> | 72.8    | <b>72.9</b> | <b>+0.1</b> |
| ShaSTA [19]      | 72.0   | <b>72.6</b> | <b>+0.6</b> | 0.612      | <b>0.593</b> | <b>-0.019</b> | 203  | <b>174</b> | <b>-29</b> | 73.0    | <b>75.3</b> | <b>+2.3</b> |

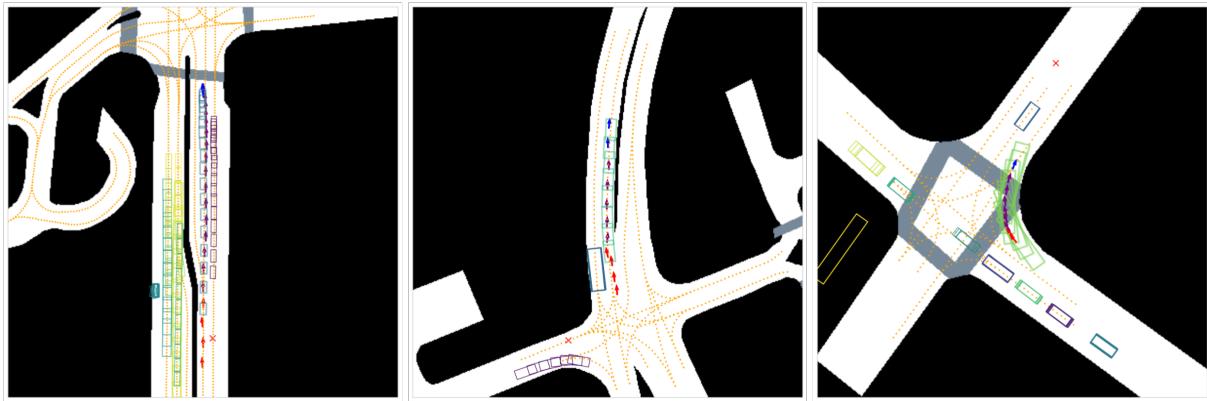


Fig. 5: Qualitative results of the offline tracking model. In each sample, GT boxes are plotted as rectangles. Each GT track is represented by a unique color. Model outputs are depicted with arrows, rather than rectangles, to differentiate from GT tracks. Red and blue arrows indicate history and future tracklets matched by the offline Re-ID model, respectively, while purple arrows show recovered trajectories. Orange dotted lines represent lanes. In the background, the white area is drivable. The gray area is the pedestrian crossing. The red cross is the average position of the ego vehicle during the occlusion.

TABLE III: Re-ID evaluation on the nuScenes validation split (Pseudo-occlusion setup).

| Method                 | CVM   | Map   | Motion | Motion+Map   |
|------------------------|-------|-------|--------|--------------|
| Association accuracy ↑ | 58.7% | 89.6% | 90.1%  | <b>90.3%</b> |

be matched. The association accuracies are listed in Tab. III. Given inputs without noise, our model in all settings achieves high association accuracies and outperforms the CVM associator by a large margin over 30%. With the perfect motion inputs, the motion branch has a higher association accuracy over the map branch by 0.5%. After the combination of the motion and the map branches, the accuracy increases by 0.2% compared to the result using only the motion branch. Hence, the two branches are complementary.

2) *Track completion model evaluation:* To standardize the evaluation, for each evaluated sample track, we mask 6 seconds of its trajectory and take 2 seconds of history tracklet and future tracklet as inputs. We chose HOME [10] as a baseline and re-implemented it on nuScenes. HOME originally decodes the predicted trajectories from sampled endpoints. To make the comparison fair, we directly give the

TABLE IV: Track completion evaluation on the nuScenes validation split (Pseudo-occlusion setup).

| Method       | ADE↓ / m     | Yaw error↓ / deg | MR↓          |
|--------------|--------------|------------------|--------------|
| HOME [10]    | 0.814        | -                | 24.3%        |
| Motion       | 0.705        | 2.38             | 14.9%        |
| Motion + map | <b>0.667</b> | <b>2.23</b>      | <b>13.3%</b> |

GT trajectory endpoints to the HOME, so that the model is also aware of the future. From the result in Tab.IV, the performance is improved on every metric after using the map information for trajectory refinement, indicating the extracted map prior improves the trajectory recovery. Our model also outperforms HOME on both ADE and MR, demonstrating its potential to recover trajectories under long occlusions.

### C. Qualitative results

As in Fig. 5, we applied our offline tracking model to improve the online tracking result by jointly performing Re-ID and track completion. The left plot shows a long occlusion with a duration of 8s and a gap distance of 71m. The middle plot shows a relatively short occlusion with a duration of 3s and a gap distance of 25m. The right plot shows an occlusion that happened on a turning vehicle at a crossroad.

The occlusion duration in the third sample is 3.5s and the gap distance is 13m. With the three samples, we show the ability of our offline tracking model to deal with long and short occlusions, as well as the non-linear motion under occlusion. Taking the online tracking results and the semantic map as inputs, our offline tracking model first correctly associated together the tracklets belonging to the same GT track. The associated future (blue) and history (red) tracklets in each plot are covered by a single GT track. Then the model recovered the occluded trajectories by interpolating the gaps between the associated tracklet pairs. Our model is not only capable of generating accurate trajectories using lane information, but it can also infer the trajectory based on motion features when the agent deviates from surrounding lanes, as shown in the right plot. The recovered trajectories (purple) also aligned well with the corresponding GT tracks.

## VI. CONCLUSION

While the previous point cloud-based offline auto labeling methods focused on generating accurate bounding boxes over visible objects, we have proposed a novel offline tracking method focusing on the occlusions on vehicles. The Re-ID module can effectively reduce IDS caused by premature termination under occlusion. Based on the Re-ID result, the track completion model recovers the occluded trajectories. Leveraging the idea from motion prediction, we innovatively extracted information from the lane map and used it as a prior to improve the performance of our models. We have demonstrated the ability of our model on the nuScenes dataset, using both the imperfect online tracking results and the handcrafted pseudo-occlusion data. The offline tracking model effectively improves the original online tracking result, showing its potential to be applied in auto labeling autonomous driving datasets. It also achieves good performances on the handcrafted pseudo-occlusions and outperforms the baselines by large margins. Future work will aim to unify the Re-ID and track completion modules into an end-to-end model and introduce a geometry affinity that utilizes cropped point cloud features to enhance association.

## REFERENCES

- [1] Holger Caesar et al. “nuPlan: A closed-loop ml-based planning benchmark for autonomous vehicles”. In: *CVPR ADP3 workshop* (2021).
- [2] Holger Caesar et al. “nuScenes: A multimodal dataset for autonomous driving”. In: *CVPR*. 2020.
- [3] Ming-Fang Chang et al. “Argoverse: 3d tracking and forecasting with rich maps”. In: *CVPR*. 2019.
- [4] Yukang Chen et al. “VoxelNeXt: Fully Sparse Voxel-Net for 3D Object Detection and Tracking”. In: *CVPR*. 2023.
- [5] Nachiket Deo, Eric Wolff, and Oscar Beijbom. “Multi-modal trajectory prediction conditioned on lane-graph traversals”. In: *CoRL*. 2022.
- [6] Shuxiao Ding et al. “3dmotformer: Graph transformer for online 3d multi-object tracking”. In: *ICCV*. 2023.
- [7] Scott Ettinger et al. “Large Scale Interactive Motion Forecasting for Autonomous Driving: The Waymo Open Motion Dataset”. In: *ICCV*. 2021.
- [8] Lue Fan et al. “Once Detected, Never Lost: Surpassing Human Performance in Offline LiDAR based 3D Object Detection”. In: *ICCV* (2023).
- [9] Jiyang Gao et al. “Vectornet: Encoding HD maps and agent dynamics from vectorized representation”. In: *CVPR*. 2020.
- [10] Thomas Gilles et al. “Home: Heatmap output for future motion estimation”. In: *ITSC*. 2021.
- [11] Rudolph Emil Kalman. “A New Approach to Linear Filtering and prediction Problems”. In: *Journal of Basic Engineering* 82.1 (1960).
- [12] Yanwei Li et al. “Unifying voxel-based representation with transformer for 3d object detection”. In: *NeurIPS* 35 (2022).
- [13] Ming Liang et al. “Learning lane graph representations for motion forecasting”. In: *ECCV*. 2020.
- [14] Tsung-Yi Lin et al. “Focal loss for dense object detection”. In: *ICCV*. 2017.
- [15] Tao Ma et al. “DetZero: Rethinking Offboard 3D Object Detection with Long-term Sequential Point Clouds”. In: *ICCV* (2023).
- [16] Ziqi Pang, Zhichao Li, and Naiyan Wang. “Simple-track: Understanding and rethinking 3d multi-object tracking”. In: *ECCV*. Springer. 2022.
- [17] Charles Ruizhongtai Qi et al. “Offboard 3d object detection from point cloud sequences”. In: *CVPR*. 2021.
- [18] Raphaël Rozenberg, Joseph Gesnouin, and Fabien Moutarde. “Asymmetrical bi-rnn for pedestrian trajectory encoding”. In: *arXiv preprint arXiv:2106.04419* (2021).
- [19] Tara Sadjadpour et al. “Shasta: Modeling shape and spatio-temporal affinities for 3d multi-object tracking”. In: *RA-L* (2023).
- [20] Victor Garcia Satorras and Max Welling. “Neural enhanced belief propagation on factor graphs”. In: *AISTATS*. 2021.
- [21] Pavel Tokmakov et al. “Learning to track with object permanence”. In: *ICCV*. 2021.
- [22] Ashish Vaswani et al. “Attention is all you need”. In: *NIPS* 30 (2017).
- [23] Qitai Wang et al. “Immortal Tracker: Tracklet Never Dies”. In: *arXiv preprint arXiv:2111.13672* (2021).
- [24] Xinshuo Weng et al. “3d multi-object tracking: A baseline and new evaluation metrics”. In: *IROS*. IEEE. 2020.
- [25] Bin Yang et al. “Auto4d: Learning to label 4d objects from sequential point clouds”. In: *arXiv preprint arXiv:2101.06586* (2021).
- [26] Tianwei Yin, Xingyi Zhou, and Philipp Krahenbuhl. “Center-based 3d object detection and tracking”. In: *CVPR*. 2021.