




JDT3D: Addressing the Gaps in LiDAR-Based Tracking-by-Attention

Brian Cheong^{*}, Jiachen Zhou^{*}, and Steven Waslander

University of Toronto

{brian.cheong, jason.zhou, steven.waslander}@robotics.utias.utoronto.ca

Abstract. Tracking-by-detection (TBD) methods achieve state-of-the-art performance on 3D tracking benchmarks for autonomous driving. On the other hand, tracking-by-attention (TBA) methods have the potential to outperform TBD methods, particularly for long occlusions and challenging detection settings. This work investigates why TBA methods continue to lag in performance behind TBD methods using a LiDAR-based joint detector and tracker called JDT3D. Based on this analysis, we propose two generalizable methods to bridge the gap between TBD and TBA methods: track sampling augmentation and confidence-based query propagation. JDT3D is trained and evaluated on the nuScenes dataset, achieving 0.574 on the AMOTA metric on the nuScenes test set, outperforming all existing LiDAR-based TBA approaches by over 6%. Based on our results, we further discuss some potential challenges with the existing TBA model formulation to explain the continued gap in performance with TBD methods. The implementation of JDT3D can be found at the following link: <https://github.com/TRAILab/JDT3D>.

Keywords: Multi-Object Tracking · Autonomous Vehicles · Computer Vision

1 Introduction

Multi-object tracking (MOT) is an essential component in the perception system of autonomous vehicles, as it allows autonomous agents to reason and plan appropriately in dynamic environments. The task involves identifying object trajectories, requiring that each object is accurately detected and consistently identified over time within the scene.

The majority of MOT methods adhere to the tracking-by-detection (TBD) paradigm. This paradigm is comprised of two separate processes: generating bounding box predictions without track IDs and associating these predictions to a set of maintained tracks based on some matching criteria. TBD methods currently achieve state-of-the-art results on 3D MOT benchmarks [12, 15, 23].

In contrast to the decoupled approach of TBD, joint detection and tracking (JDT) approaches have attempted to unify the detection and tracking tasks in an

^{*} Denotes equal contribution.

end-to-end manner [18, 19, 29, 31]. Recent works in autonomous driving perception [7, 17] and computer vision [8] demonstrate that networks trained end-to-end have shown superior performance compared to pipeline-based techniques, despite initially lower performance from earlier works [21, 27]. End-to-end methods can learn richer and more generalizable representations while minimizing the task-specific engineering of individual modules required in a pipeline-based approach. While TBD methods enforce a unidirectional flow of information from the detector to the tracker, the JDT formulation allows both the detector and tracker to exchange useful information to enhance the performance of both tasks.

Integrating detection and tracking into a joint learning task requires a unique approach to formulating the problem and training the model. One such approach is tracking-by-attention (TBA), illustrated in Fig. 1. In TBA, objects are represented as vector embeddings, or "queries", that are used to detect the same object over multiple frames. While recent TBA methods [18, 19, 29, 31] have been proposed as alternatives to TBD, they continue to underperform on the MOT task, especially in the LiDAR domain [29]. However, they have shown potential in certain cases, such as significantly fewer ID-switch errors [19]. The unexplored potential of TBA in the LiDAR domain motivates our work in addressing a key question: what is holding back LiDAR-based TBA? In this work, we propose a LiDAR-based TBA tracking method called JDT3D, based on which we explore different aspects that hold back LiDAR-based TBA.

One insight is that LiDAR-based tracking methods can suffer from sparse supervision signals due to point cloud and object sparsity. Inspired by object sampling augmentation [25] for LiDAR-based detection methods, we propose track sampling, a temporally consistent augmentation method that injects consistent objects over multiple LiDAR frames to enrich supervision signals while maintaining temporal consistency.

In addition, existing TBA methods maintain an inconsistency between training and inference regarding their query propagation method. During training, queries that are matched to ground truth samples are propagated to the next frame, but at inference, propagation is based on a confidence threshold. This inconsistency between training and inference could confuse the model to over-trust false positive queries. To address this, we propose a confidence threshold propagation strategy consistent for both training and inference and conduct comprehensive evaluations and discussions on the strategy design.

We summarize our contributions as follows:

JDT3D We present JDT3D, a LiDAR-based tracking-by-attention model that outperforms existing LiDAR-based tracking-by-attention methods on the nuScenes test set by over 6% on the AMOTA metric.

Confidence-based query propagation We perform ablations to validate the effectiveness and generalizability of the confidence-based propagation criterion.

Track Sampling Augmentation We propose and validate a novel data augmentation technique for multi-frame LiDAR methods that significantly improves performance and convergence rates during training.




Fig. 1: Illustration of tracking by attention. Each object is represented by a query. In the diagram, the yellow motorcycle leaves the frame, so it is removed from the set of maintained queries. To detect new objects in the scene, such as the tan bus, proposal queries are appended at each time step, represented by the blank squares.

Analysis Based on our experiments using JDT3D, we provide a detailed analysis to explain the continued gap in performance between LiDAR-based TBD and TBA. We attribute the gap to the reduced detection performance when training on multiple consecutive frames and the handling of queries that results in "temporal confusion".

2 Related Works

Multi-object tracking (MOT) methods either follow the tracking-by-detection paradigm, in which object detection is independently performed on single frames before the tracking module, or they perform detection and tracking jointly in an end-to-end fashion. MOT can also be further categorized into 2D and 3D MOT. The task of 2D MOT involves predicting object tracks in 2D space, while 3D MOT predicts the tracks in 3D space. 2D MOT methods generally take in camera images as sensor input, while 3D MOT methods may take in camera images, LiDAR point clouds, or both. In this work, we will be focusing our analysis on 3D MOT methods.

2.1 Tracking-by-detection (TBD)

In the tracking-by-detection paradigm, the detection and tracking problems are solved by two distinct modules. First, a single-frame object detector is used to detect objects in a given scene. Next, the detections are associated with a set of maintained tracks by a tracking module and used to update them with current detection information. Unassigned detections are used to initialize new tracks.

While different tracking methods exist among TBD methods, they are all consistently tied to the detector's performance. Several top-performing 3D tracking methods [1, 4, 15] are strong detectors that have been paired with the simple tracking method of CenterPoint [26]. The formulation of TBD inhibits the detector from taking advantage of temporal object information that would enhance object detection. As well, errors in the detection are guaranteed to be propagated

downstream to the tracking task. This motivates the exploration of methods that can jointly optimize detection and tracking to achieve MOT.

2.2 Joint Detection and Tracking (JDT)

Joint detection and tracking (JDT) is a recent paradigm in both 2D and 3D MOT that involves training a model to perform detection and tracking end-to-end. In contrast to the TBD methods that rely on a dedicated detector to generate detections, JDT methods simultaneously output object detections with associated tracking IDs. Without heuristic matching or hand-crafted rules, SimTrack [16] builds upon CenterPoint [26] by learning a hybrid-time centerness map between two consecutive LiDAR point clouds. AlphaTrack [28] adds a cross-modal fusion scheme to the CenterPoint tracking method [26] and applies an alternating training strategy to balance between the detection branch and appearance affinity. The above two methods perform detection and tracking jointly, but known tracks are not utilized to inform future detections. 3D DetecTrack [9] proposes tracklet-aware 3D detection that reconfigures the initial detection using the latest tracklets’ information and employs a spatio-temporal gated GNN for association.

Tracking-by-Attention (TBA) A growing trend within JDT is the tracking-by-attention (TBA) approach that emerged from TrackFormer [18] for 2D MOT. TBA uses the transformer architecture [22] and formulates tracking as a frame-to-frame set prediction problem that can be trained in an end-to-end manner.

TrackFormer [18] extends a 2D query-based detector, DETR [3] and its variant Deformable DETR [33], by designing autoregressive track queries that are passed across frames with associated object identities. MOTRv2 [31] uses a pre-trained YOLOX detector to initialize their object queries and uses the detections to anchor the decoded bounding box positions. In addition, track queries during training are propagated based on the prediction confidence rather than whether a prediction was matched to a ground truth object. This naturally creates false positive and false negative cases during training, removing the need for query injection and removal augmentations [18, 27]. JDT3D applies confidence-based track query propagation to LiDAR-based TBA, and we perform ablations showing the method’s improved performance over the ground truth matching criterion.

MUTR3D [30] and PF-Track [19] demonstrate the transfer of the tracking-by-attention framework to multi-view vision-only 3D MOT. PF-Track [19] also designs past and future reasoning modules to enhance the spatio-temporal representation of queries. MotionTrack [29] investigates the feasibility of extending the TBA paradigm to LiDAR-based and LiDAR-camera fusion-based 3D JDT by building upon TransFusion [1], a query-based 3D object detector. Rather than implicitly learning which queries should be suppressed through the transformer decoder, MotionTrack learns an explicit affinity matrix between tracks and objects in the scene and performs greedy matching to determine their association. This explicit matching echoes the decoupling found in TBD methods and potentially detracts from the benefits of an end-to-end method. While vision-only

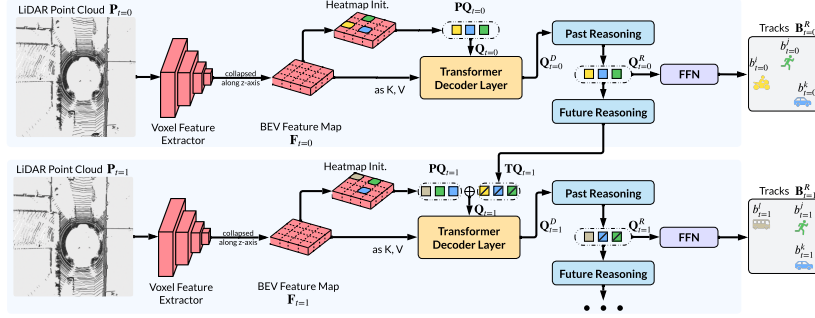


Fig. 2: JDT3D Architecture. At each time step, a BEV feature map is extracted and used to initialize a set of proposal queries. The proposal queries are concatenated to track queries passed from the previous frame and used to predict objects in the scene. Track queries detect the same unique objects in each time step, while proposal queries detect untracked or new objects.

TBA methods have shown surprising results on the MOT task, the baseline results set by MotionTrack’s LiDAR-only and fusion TBA methods underperform compared to LiDAR-based TBD baselines such as CenterPoint [26] and SimpleTrack [20], prompting further exploration of LiDAR TBA methods. Our work explores different potential hypotheses for the reported performance gap between LiDAR-based TBD and TBA methods and proposes improvements to the existing TBA methods.

2.3 LiDAR Data Augmentations

As with other machine learning tasks, data augmentation can greatly enhance the performance of LiDAR-based 3D object detection methods. Some common augmentations include random flipping, rotation, translation, and scaling of the point cloud [6]. In addition to these, object sampling is also a widely used augmentation method for LiDAR-based object detectors [1, 15, 26]. First introduced in SECOND [25], object sampling involves creating a database of ground truth object point cloud scans and bounding boxes and randomly sampling and injecting them into the scene during training to improve the convergence and performance overall. However, this augmentation does not apply directly to training TBA methods, since it would result in sporadic track instances in the scene in each training sequence. Our proposed track sampling augmentation samples consecutive track instances and injects them into a sequence of frames during training. This maintains the temporal nature of the object tracks while still offering the same benefits of faster convergence.

3 Method

To perform our analysis, we developed our own baseline LiDAR-based TBA method since no such implementations were made available as of the writing of this paper. We build on aspects from query-based 3D object detection [1] and vision-only TBA [19] to formulate our method, called JDT3D. JDT3D is summarized in Fig. 2.

3.1 Overview

At each time step, a feature extractor takes in the LiDAR scan of the scene, \mathbf{P} , and converts it into a BEV feature map representation, \mathbf{F} , using a standard voxel feature extractor [32]. This feature map is used to predict heatmaps of potential object locations and initialize a set of proposal queries, \mathbf{PQ} [1].

In the first time step, the proposal queries are passed through a transformer decoder layer with the BEV feature map to obtain a set of updated queries, \mathbf{Q}^D . The past reasoning module refines the updated queries with historical track information [19] before a feedforward network decodes the refined queries, \mathbf{Q}^R into bounding boxes and class predictions, \mathbf{B}^R . During training, \mathbf{Q}^D is decoded into a set of bounding boxes and class predictions, \mathbf{B}^D , for an auxiliary loss.

To create the track queries for the next frame, the predictions and their corresponding queries are filtered with a confidence threshold. The filtered queries are passed to the future reasoning module [19] which predicts a motion trajectory and projects each query into the next time step.

At every following time step, new proposal queries are initialized with the heatmap as described above and are concatenated with track queries from the previous time step, \mathbf{TQ} , before being passed to the transformer decoder.

To perform tracking, high-confidence predictions and associated queries are assigned an ID. Track queries are used to detect the same object consistently over multiple frames, while proposal queries are used to detect objects to be added to the track query set. Tracking is achieved implicitly through the direct relationship between each prediction and each query passed into the transformer decoder. Trajectories can be obtained by tracing a given track query back in time.

3.2 Query Propagation

During training, only the refined queries corresponding to predictions with a confidence score greater than τ_{pass} are passed to the next time step. This method of passing and pruning queries has the advantage of naturally creating instances of false negative and false positive track queries during training [31]. False negative track queries are missed tracks that should be re-detected in the next frame as if they were new objects entering the scene. False positive track queries simulate tracks that have left the scene and should be pruned in the next time step.

This method has the benefit of better matching the inference query propagation behaviour. As well, it creates more instances during training where the decoder must suppress track query predictions and output positive proposal query predictions, ensuring robust handling of track births and deaths.

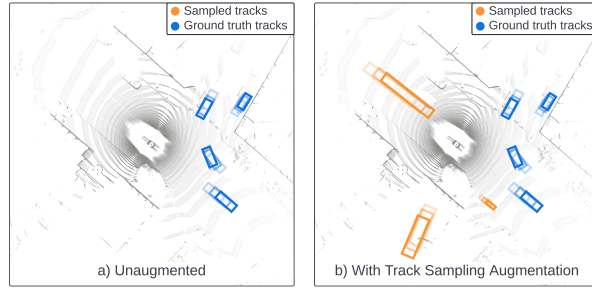


Fig. 3: An example of track sampling augmentation over three consecutive frames from the nuScenes dataset. The original trajectories and sampled trajectories are shown in blue and orange, respectively. Only a subset of the tracks and the third LiDAR scan are shown for visual clarity. The older boxes are shown with transparency.

3.3 Ground Truth Assignment

The ground truth assignment of the decoder predictions, \mathbf{B}_t^D , and the refined predictions, \mathbf{B}_t^R , are both done in the same manner using a two-stage assignment process. The first stage involves matching the track queries based on their previously matched ground truth objects [18]. If the previously matched ground truth track is not present in the current time step, such as in the case of a track leaving the scene, then the track query prediction is assigned to no object. False positive track queries were never assigned a ground truth track previously, so they are also assigned to no object.

The second stage of the assignment involves matching the proposal queries to the unmatched ground truths. The unmatched ground truths would include new objects in the scene and missed objects from previous time steps. A Hungarian matcher is used to compute the optimal assignment between the unmatched queries and ground truths based on an association cost [10].

3.4 Track Sampling Augmentation

To improve the convergence rate and performance, we introduce track sampling augmentation during training. Inspired by the object sampling augmentation [25], our track sampling involves adding additional tracks into a training clip from a database of tracks in a temporally consistent manner, as illustrated in Fig. 3. The original object sampling augmentation does not account for temporal consistency, and applying it naively to our multi-frame training scheme would result in disjointed track instances during training.

First, the training set is preprocessed to generate a database of ground truth tracks, including bounding boxes, object classes, track IDs, and LiDAR points contained in the bounding boxes. To perform track sample augmentation on a training clip, a set of tracks is randomly sampled, avoiding track IDs that overlap with the existing ground truth objects in the scene. This enforces unique IDs for

the tracks in the sequence. To inject a track, L consecutive ground-truth boxes are randomly selected from each sampled track. In each training frame, the LiDAR points corresponding to the sampled objects are used to replace regions in the original point cloud. The objects are injected in chronological order over the L frames. A collision check is performed to ensure that there are no collisions between injected objects and the original ground truth objects in the scene. In the case of a collision, the would-be sampled object is pruned from that frame, simulating an occlusion.

The benefits of track sampling mirror those of object sampling [25], with the extra ground truth objects in the scene providing more samples to learn from in a single frame. This results in a stronger training signal and more varied scenes.

3.5 Network Training and Losses

To train JDT3D, L LiDAR frames are passed sequentially to the model as a training sample. The loss for a given frame, \mathcal{L}_i , is computed as the weighted sum of the following losses:

$$\mathcal{L}_i = \lambda_h \mathcal{L}_h + \lambda_f \mathcal{L}_f + \sum_{m=\{D,R\}} (\lambda_{reg}^m \mathcal{L}_{reg}^m + \lambda_{cls}^m \mathcal{L}_{cls}^m) \quad (1)$$

The superscripts D and R correspond to predictions output by \mathbf{Q}^D and \mathbf{Q}^R , respectively. \mathcal{L}_{reg}^m is the L1 loss over the bounding boxes, and \mathcal{L}_{cls}^m is the focal loss [14] over the classification predictions. \mathcal{L}_h is the Gaussian Focal Loss on the predicted object heatmaps [11, 26], and \mathcal{L}_f is the L1 loss over the motion trajectory predicted by the future reasoning module [19]. Each loss has an associated scalar weight term, λ . The total loss is computed as the sum of each individual frame loss:

$$\mathcal{L}_{clip} = \sum_{i=1}^L \mathcal{L}_i \quad (2)$$

4 Experiments

4.1 Dataset and Metrics

JDT3D is trained and evaluated on the nuScenes dataset [2], a multi-modal autonomous driving dataset containing data from six multi-view cameras, five radars, and a 360-degree LiDAR. The dataset contains a total of 40,157 frames across 1000 scenes, of which 850 scenes are annotated for training (700) and validation (150) at a frame rate of 2 Hz. We train on the seven object classes annotated for the tracking task.

To evaluate JDT3D, the average multi-object tracking accuracy (AMOTA) and average multi-object tracking precision (AMOTP) tracking metrics [24] are computed across the seven tracking classes of nuScenes, evaluating the tracking accuracy and tracking position misalignment, respectively. Secondary metrics are

Table 1: A comparison of 3D multi-object tracking methods evaluated on the nuScenes [2] test set. **Bold** represents the best results among the Transformer-based TBA methods, while Underlined numbers are the best overall.

Method	TBA	Modality	AMOTA \uparrow	AMOTP \downarrow	FP \downarrow	FN \downarrow	IDS \downarrow
AB3DMOT [24]		L	0.151	1.501	15088	75730	9027
Chui et al. [5]		L	0.550	0.798	17533	33216	950
CenterPoint [26]		L	0.638	0.555	18612	<u>22928</u>	760
SimpleTrack [20]		L	<u>0.668</u>	<u>0.550</u>	17514	23451	575
PF-Track [19]	✓	C	0.434	1.252	19048	42758	<u>249</u>
DQTrack [13]	✓	C	0.523	1.096	-	-	1204
MotionTrack-L [29]	✓	L	0.51	0.99	-	-	9705
MotionTrack-LC [29]	✓	L+C	0.55	0.871	-	-	8716
JDT3D (ours)	✓	L	0.574	0.837	<u>11152</u>	<u>29919</u>	254

also computed by the nuScenes MOT benchmark [2], such as the number of false positives, false negatives, and identity switches. We also evaluate the detection performance across the tracking classes, with the mean average precision (mAP) and nuScenes Detection Score (NDS) metrics.

4.2 Implementation Details

VoxelNet [25, 32] is used as the LiDAR backbone to extract the BEV feature map. For each frame, the top 200 predictions from the heatmap are used to initialize the proposal queries. The training clip length L is set to 3 frames. The confidence threshold for track queries τ_{pass} is set to 0.4 for training and inference. The Hungarian matching cost for ground truth assignment consists of the sum of the L1 cost on the predicted box and the focal loss [14] on the predicted class confidence. All detections are generated in the car ego-frame, so the object centers of the track queries are updated to their relative positions in the current training frame based on the car’s ego-motion between time steps. A pretrained VoxelNet backbone from BEVFusion [15] is used to initialize the feature extractor. To train JDT3D, the model is first trained on the detection task using $L = 1$ for 20 epochs. It is then further trained on the tracking task using $L = 3$ for another 20 epochs. All experiments were trained and evaluated on eight NVIDIA V100 GPUs using random flipping, rotation, translation, and scaling on the training sequence, and track sampling augmentation, unless otherwise specified. JDT3D performs inference at 3.85 Hz on a Tesla V100.

4.3 Main Results

Our main results compare JDT3D with baseline TBD methods and relevant TBA methods evaluated on the nuScenes test set [2], summarized in Tab. 1. JDT3D outperforms the current state-of-the-art LiDAR-based TBA methods, with a relative improvement of 12.5% on the AMOTA metric and 15.5% on the

Table 2: Ablation on the track sampling augmentation, reported on the nuScenes validation set. The mAP metric is computed on the seven tracking categories.

Track Sampling Rate	AMOTA \uparrow	AMOTP \downarrow	IDS \downarrow	mAP \uparrow
0.0	0.314	1.377	388	0.225
0.25	0.600	0.864	210	0.3867
0.50	0.622	0.816	203	0.400
1.0	0.621	0.790	150	0.399

AMOTP metric on the test split. Moreover, with LiDAR input only, JDT3D also outperforms the fusion-based MotionTrack-LC [29] with a relative improvement of 4.36% on AMOTA and 3.79% on AMOTP on the test set. While JDT3D underperforms compared to baseline TBD methods on most metrics, it achieves the lowest number of ID switches among LiDAR-only methods, indicating a strong track coherence.

JDT3D has significantly more false negatives compared to TBD methods, but fewer false positives. The potential sources of false negatives are either due to missed detections or tracks diverging from the ground truths. However, the latter case would result in both high false negatives and false positives. Based on the low false positive results from JDT3D, we conclude that the false negatives are mainly due to missed detections.

4.4 Ablation Studies

Track Sampling. By comparing the performance of JDT3D on the nuScenes validation set with and without track sampling augmentation, we verify that sampling extra tracks into a training sequence significantly improves the convergence rate and the final performance. The comparison summarized in Tab. 2 shows a 97% relative improvement on the AMOTA metric from track sampling augmentation compared to no track sampling, with a similar trend across all other metrics, including detection metrics such as the mAP. We also analyzed the effect of varying the number of objects that are sampled by reducing the sampling frequency to 0.5 and 0.25 of the nominal rate. We observe that the AMOTA and mAP metrics plateau after a sampling rate of 0.5, although the AMOTP and IDS metrics see continued improvement with increased track sampling. The benefits of implementing track sampling are expected to generalize to other LiDAR-based TBA methods, but there are currently no available implementations to verify this hypothesis.

Training Clip Length. To validate our choice of clip length, the performances of JDT3D models trained with two to four frame sequences were compared and summarized in Tab. 3. The best performance is achieved when trained over $L = 3$ frames, with a significant jump when moving from $L = 2$ to $L = 3$. We hypothesize that the training process benefits from the longer-term, temporally consistent information, resulting in improved tracking performance. Moving from

Table 3: Ablation on the training sequence length, L , reported on the nuScenes validation set.

Clip Length (L)	AMOTA \uparrow	AMOTP \downarrow	IDS \downarrow	mAP \uparrow
2	0.397	1.294	446	0.214
3	0.621	0.790	150	0.399
4	0.599	0.783	173	0.392

Table 4: Ablation on the query propagation criteria, reported on the nuScenes [2] validation set. Passing queries with high confidence encourages the re-detection of poor detections and handling of false positives and false negatives.

Query Propagation	AMOTA \uparrow	AMOTP \downarrow	Recall \uparrow	IDS \downarrow
Matched to a GT	0.464	0.917	0.511	82
High Confidence	0.621	0.790	0.669	150

$L = 3$ to $L = 4$, the performance drops slightly while also significantly increasing the train time and required memory.

Query Propagation Method. In Tab. 4, we compare different query propagation methods; the first method is passing queries that have been matched to a ground truth [18, 19], and the second is passing based on the τ_{pass} threshold for high confidence predictions as described in Sec. 3.2. Results show that using confidence as a filtering criterion during training enhances the overall tracking metrics AMOTA and AMOTP by a large margin and significantly reduces the number of false positives and false negatives while increasing the number of true positives. While the ID switch errors are higher, this increase is less than 2% relative to the increase in true positives. Thus, we verify that using confidence as the query propagation criterion is more appropriate for the TBA paradigm.

In addition, we show the generalizability of this query propagation method by performing the same ablation on PF-Track. The results in Tab. 5 show improvements across most metrics, with a similar behaviour of more IDS errors.

Comparison with TBD Baseline. We compare the performance between our TBA framework and baseline TBD methods such as the CenterPoint greedy-based tracking method [26] and SimpleTrack [20]. To compare our tracking

Table 5: An ablation comparing the original ground truth-based propagation and our confidence-based propagation on PF-Track-S [19] evaluated on the nuScenes [2] validation set.

Query Propagation	AMOTA \uparrow	AMOTP \downarrow	Recall \uparrow	IDS \downarrow
Matched to a GT	0.408	1.343	0.507	166
High Confidence	0.426	1.316	0.558	306

Table 6: Comparison with "Tracking by Detection" on the nuScenes validation set, using the detections generated by JDT3D.

Tracking Method	AMOTA \uparrow	AMOTP \downarrow	IDS \downarrow
CenterPoint [26]	0.608	0.825	491
SimpleTrack [20]	0.573	0.877	5585
JDT3D (Ours)	0.621	0.790	150

Table 7: Comparison of detection performance between the single-frame pretraining and the three-frame fine-tuning. Evaluated on the seven tracking classes on the nuScenes validation set. The best results are typeset in boldface.

JDT3D variant	mAP \uparrow	NDS \uparrow
Single-frame	0.605	0.549
Three-frame	0.399	0.502

method to the TBD trackers independent of detection performance, we pass the detections generated by JDT3D to the tracker and evaluate its tracking performance compared to our joint approach. The results in Tab. 6 demonstrate clear advantages of our joint end-to-end approach that offers temporal consistency and data association and reaffirm that the gap in tracking performance lies in the missed detections.

4.5 Detection Performance

As mentioned in Sec. 4.2, JDT3D is pre-trained to perform single-frame detection and fine-tuned to perform tracking on sequences of multiple LiDAR frames. While the pre-training step achieves reasonable detection performance, the performance drops significantly after fine-tuning, as seen in Tab. 7. This indicates that the current TBA problem formulation may not be the best alignment of the detection and tracking objectives.

4.6 Proposal and Track Query Imbalances

Tab. 8 shows further analysis of the JDT3D outputs. During training, the positively matched proposal queries, or untracked newborn objects, are more confident in the first frame than in subsequent frames. Also, the positively matched track queries, or previously tracked objects, tend to be more confident than the newborn objects. This indicates that the model has a strong bias toward previously tracked objects. This bias could lead to more missed objects and fewer initialized tracks, explaining the high number of false negatives indicated in Tab. 1.

One hypothesis for this behaviour is that this is caused by the imbalance in training samples assigned to each type of query. Tab. 9 shows the average distribution of true positive predictions of both proposal and track queries in

Table 8: The average confidence of positively matched proposal and track predictions over a three-frame training sequence. There are no track queries in the first frame.

Avg. Confidence	Frame 1	Frame 2	Frame 3
Positively Matched Proposals	0.669	0.409	0.387
Track Predictions	n/a	0.759	0.760

Table 9: The distribution of positive and negative matches for proposal queries and track queries in the second and third frames of training.

	PQ Predictions	TQ Predictions
Avg Positive Matches	4 (2%)	47 (94%)
Avg Negative Matches	196 (98%)	3 (6%)

the second and third frames of training. Note that the proportion of positively matched proposal queries is significantly lower than for track queries. This is somewhat expected since the average track birthrate per frame is 0.7, so most objects in a scene were present in the previous frame. This imbalance in positive training samples between the proposal and track queries could explain the confidence imbalance observed.

5 Discussion

Based on our analysis from Tab. 6 we show that the reason LiDAR-based TBA methods continue to underperform compared to TBD methods is due to poorer detection performance. In addition, we note that there is a significant drop in the detection performance when fine-tuning the model on multiple frames in Tab. 7. In the following sections, we discuss hypotheses to explain the observed results and propose potential improvements to build on JDT3D.

5.1 Temporal Confusion

Our hypothesis for the observed drop in performance in Tab. 7 lies in the fact that the problem formulation of TBA is more complex than query-based detection. In query-based detection, the handling of the queries by the decoder is straightforward: if the query represents an object in the scene, it should output a high-confidence bounding box from the query. On the other hand, in TBA, the handling of queries is conditional: depending on whether the decoder receives a track query that represents the same object, the decoder needs to determine whether the incoming proposal query is a duplicate that should have a low confidence or a unique object that should have a high confidence, as illustrated in Fig. 4. This complex interaction results in "temporal confusion" such that the expected model behaviour is different depending on the query inputs. We hypothesize that the reason temporal confusion has not been an issue in vision-based TBA methods [18, 19, 30] is due to their larger decoder size enabling more



Fig. 4: An example of temporal confusion, where the decoder must handle the same proposal query differently based on the presence of track queries. In Case 1, the proposal query should be a positive prediction, while in Case 2, it should be a negative prediction.

complex reasoning about the queries. To test this hypothesis, further optimizations would be needed to enable a larger decoder with a LiDAR backbone due to hardware constraints.

5.2 Potential Improvements

With these results in mind, we propose two potential directions that can be further explored to bridge the gap between JDT and TBD methods.

Firstly, based on the ablation performed to evaluate the training clip size, we observed that three-frame training significantly improved performance over two-frame training, although extending to four frames reduced performance slightly. Further exploration of the effect of clip length on tracking performance may yield insights on how to better leverage longer term temporal information.

Secondly, optimizations to the model size and memory footprint would enable a larger and more complex decoder, that could potentially address the temporal confusion hypothesis.

6 Conclusion

This paper presents JDT3D, a novel LiDAR tracking method that makes significant progress in closing the gap between TBD and TBA methods. JDT3D is trained using track sampling augmentation, a generalizable data augmentation method that injects extra long-term tracks into training samples. We show that track sampling augmentation significantly improves training performance and convergence, and we hope that future LiDAR-based trackers can leverage these findings. As well, we show that the query propagation scheme during training also has a significant impact on the performance, showing that using a confidence threshold to filter low confidence predictions improves tracking performance greatly. Based on our experiments, we hypothesize that the complexity of tracking-by-attention combined with the memory constraints imposed by LiDAR feature extractors limiting the size of transformer decoders is the cause of the drop in detection performance between single and multi-frame variants of JDT3D. We propose further areas of exploration including exploring different training sequence lengths and enabling more decoder layers.

References

1. Bai, X., Hu, Z., Zhu, X., Huang, Q., Chen, Y., Fu, H., Tai, C.L.: Transfusion: Robust lidar-camera fusion for 3d object detection with transformers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1090–1099 (June 2022)
2. Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O.: nuscenes: A multimodal dataset for autonomous driving. 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 11618–11628 (2019), <https://api.semanticscholar.org/CorpusID:85517967>
3. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: European Conference on Computer Vision. pp. 213–229. Springer (2020)
4. Chen, Y., Yu, Z., Chen, Y., Lan, S., Anandkumar, A., Jia, J., Álvarez, J.M.: Focalformer3d : Focusing on hard instance for 3d object detection. ArXiv **abs/2308.04556** (2023), <https://api.semanticscholar.org/CorpusID:260736073>
5. Chiu, H.k., Prioletti, A., Li, J., Bohg, J.: Probabilistic 3d multi-object tracking for autonomous driving. arXiv preprint arXiv:2001.05673 (2020)
6. Hahner, M., Dai, D., Liniger, A., Gool, L.V.: Quantifying data augmentation for lidar based 3d object detection. CoRR **abs/2004.01643** (2020), <https://arxiv.org/abs/2004.01643>
7. Hu, Y., Yang, J., Chen, L., Li, K., Sima, C., Zhu, X., Chai, S., Du, S., Lin, T., Wang, W., et al.: Planning-oriented autonomous driving. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17853–17862 (2023)
8. Huang, T.E., Liu, Y., Van Gool, L., Yu, F.: Video task decathlon: Unifying image and video tasks in autonomous driving. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8647–8657 (2023)
9. Koh, J., Kim, J., Yoo, J.H., Kim, Y., Kum, D., Choi, J.W.: Joint 3d object detection and tracking using spatio-temporal representation of camera image and lidar point clouds. Proceedings of the AAAI Conference on Artificial Intelligence **36**(1), 1210–1218 (Jun 2022). <https://doi.org/10.1609/aaai.v36i1.20007>, <https://ojs.aaai.org/index.php/AAAI/article/view/20007>
10. Kuhn, H.W.: The hungarian method for the assignment problem. Naval Research Logistics (NRL) **52** (1955), <https://api.semanticscholar.org/CorpusID:9426884>
11. Law, H., Deng, J.: Cornernet: Detecting objects as paired keypoints. Int. J. Comput. Vis. **128**(3), 642–656 (2020). <https://doi.org/10.1007/S11263-019-01204-1>, <https://doi.org/10.1007/s11263-019-01204-1>
12. Li, X., Xie, T., Liu, D., Gao, J., Dai, K., Jiang, Z., Zhao, L., Wang, K.: Poly-mot: A polyhedral framework for 3d multi-object tracking. ArXiv **abs/2307.16675** (2023), <https://api.semanticscholar.org/CorpusID:260334560>
13. Li, Y., Yu, Z., Phillion, J., Anandkumar, A., Fidler, S., Jia, J., Alvarez, J.: End-to-end 3d tracking with decoupled queries. In: IEEE/CVF International Conference on Computer Vision (ICCV) (2023)
14. Lin, T.Y., Goyal, P., Girshick, R.B., He, K., Dollár, P.: Focal loss for dense object detection. 2017 IEEE International Conference on Computer Vision (ICCV) pp. 2999–3007 (2017), <https://api.semanticscholar.org/CorpusID:47252984>

15. Liu, Z., Tang, H., Amini, A., Yang, X., Mao, H., Rus, D., Han, S.: Bevfusion: Multi-task multi-sensor fusion with unified bird’s-eye view representation. 2023 IEEE International Conference on Robotics and Automation (ICRA) pp. 2774–2781 (2022), <https://api.semanticscholar.org/CorpusID:249097415>
16. Luo, C., Yang, X., Yuille, A.: Exploring simple 3d multi-object tracking for autonomous driving. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 10488–10497 (October 2021)
17. Mahmoud, A., Hu, J.S.K., Waslander, S.L.: Dense voxel fusion for 3d object detection. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 663–672 (January 2023)
18. Meinhardt, T., Kirillov, A., Leal-Taixé, L., Feichtenhofer, C.: Trackformer: Multi-object tracking with transformers. In: 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8834–8844 (2022). <https://doi.org/10.1109/CVPR52688.2022.00864>
19. Pang, Z., Li, J., Tokmakov, P., Chen, D., Zagoruyko, S., Wang, Y.X.: Standing between past and future: Spatio-temporal modeling for multi-camera 3d multi-object tracking. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 17928–17938 (June 2023)
20. Pang, Z., Li, Z., Wang, N.: SimpleTrack: Understanding and Rethinking 3D Multi-object Tracking. In: Karlinsky, L., Michaeli, T., Nishino, K. (eds.) Computer Vision – ECCV 2022 Workshops. pp. 680–696. Lecture Notes in Computer Science, Springer Nature Switzerland (2021). https://doi.org/10.1007/978-3-031-25056-9_43
21. Sun, P., Cao, J., Jiang, Y., Zhang, R., Xie, E., Yuan, Z., Wang, C., Luo, P.: Transtrack: Multiple-object tracking with transformer. arXiv preprint arXiv: 2012.15460 (2020)
22. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L.u., Polosukhin, I.: Attention is all you need. In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.) Advances in Neural Information Processing Systems. vol. 30. Curran Associates, Inc. (2017), <https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf>
23. Wang, L., Zhang, X.N., Qin, W., Li, X., Yang, L., Li, Z., Zhu, L., Wang, H., Li, J., Liu, H.: Camo-mot: Combined appearance-motion optimization for 3d multi-object tracking with camera-lidar fusion. IEEE Transactions on Intelligent Transportation Systems **24**, 11981–11996 (2022), <https://api.semanticscholar.org/CorpusID:252089635>
24. Weng, X., Wang, J., Held, D., Kitani, K.: 3d multi-object tracking: A baseline and new evaluation metrics. In: 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 10359–10366 (2020). <https://doi.org/10.1109/IROS45743.2020.9341164>
25. Yan, Y., Mao, Y., Li, B.: Second: Sparsely embedded convolutional detection. Sensors (Basel, Switzerland) **18** (2018), <https://api.semanticscholar.org/CorpusID:52957856>
26. Yin, T., Zhou, X., Krahenbuhl, P.: Center-based 3d object detection and tracking. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 11784–11793 (June 2021)
27. Zeng, F., Dong, B., Wang, T., Chen, C., Zhang, X., Wei, Y.: Motr: End-to-end multiple-object tracking with transformer. ArXiv **abs/2105.03247** (2021), <https://api.semanticscholar.org/CorpusID:234096063>

28. Zeng, Y., Ma, C., Zhu, M., Fan, Z., Yang, X.: Cross-modal 3d object detection and tracking for auto-driving. In: IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2021, Prague, Czech Republic, September 2021. IEEE (2021)
29. Zhang, C., Zhang, C., Guo, Y., Chen, L., Happold, M.: Motiontrack: End-to-end transformer-based multi-object tracking with lidar-camera fusion. 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW) pp. 151–160 (2023), <https://api.semanticscholar.org/CorpusID:259286967>
30. Zhang, T., Chen, X., Wang, Y., Wang, Y., Zhao, H.: Mutr3d: A multi-camera tracking framework via 3d-to-2d queries. In: 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). pp. 4536–4545 (2022). <https://doi.org/10.1109/CVPRW56347.2022.00500>
31. Zhang, Y., Wang, T., Zhang, X.: Motrv2: Bootstrapping end-to-end multi-object tracking by pretrained object detectors. 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 22056–22065 (2022), <https://api.semanticscholar.org/CorpusID:253581846>
32. Zhou, Y., Tuzel, O.: Voxelnet: End-to-end learning for point cloud based 3d object detection. 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition pp. 4490–4499 (2017), <https://api.semanticscholar.org/CorpusID:42427078>
33. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable DETR: deformable transformers for end-to-end object detection. In: 9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021 (2021)