

# MCTR: Multi Camera Tracking Transformer

Alexandru Niculescu-Mizil

alexnic@gmail.com

Deep Patel, Iain Melvin

NEC Laboratories America

{dpatel, iain}@nec-labs.com

## Abstract

*Multi-camera tracking plays a pivotal role in various real-world applications. While end-to-end methods have gained significant interest in single-camera tracking, multi-camera tracking remains predominantly reliant on heuristic techniques. In response to this gap, this paper introduces Multi-Camera Tracking tRansformer (MCTR), a novel end-to-end approach tailored for multi-object detection and tracking across multiple cameras with overlapping fields of view. MCTR leverages end-to-end detectors like DEtector TRansformer (DETR) to produce detections and detection embeddings independently for each camera view. The framework maintains set of track embeddings that encapsulate global information about the tracked objects, and updates them at every frame by integrating the local information from the view-specific detection embeddings. The track embeddings are probabilistically associated with detections in every camera view and frame to generate consistent object tracks. The soft probabilistic association facilitates the design of differentiable losses that enable end-to-end training of the entire system. To validate our approach, we conduct experiments on MMPTTrack and AI City Challenge, two recently introduced large-scale multi-camera multi-object tracking datasets.*

## 1. Introduction

Object tracking has long been a central challenge in computer vision attracting substantial attention in the research community due to its applicability in numerous real-world applications. While the majority of the research efforts have concentrated on multi-object tracking in single camera video feeds, there has been a rising demand for multi-object multi-camera tracking due to the increasing prevalence of multi-camera systems deployed in diverse applications such as security, monitoring, or sports analytics. In these applications multi-camera setups offer a multitude of advantages over single-camera counterparts, including increased coverage, reduction of blind spots and heightened tracking robustness, particularly in scenarios involving de-

tection failures or prolonged occlusions.

The predominant approach to multi-camera tracking thus far has revolved around heuristic techniques, which amalgamate various components such as person re-identification, single-camera tracking, homography estimation, and clustering. These techniques have demonstrated commendable performance; however, they are inherently heuristic, and their performance often hinges on the effectiveness of hand-crafted rules and the quality of heuristics used. Recently, the application of transformer-based algorithms to tracking has received significant attention [5, 31, 42, 52, 57] due to their ability to encapsulate the entirety of the tracking process within a unified, end-to-end framework. However, this transition toward end-to-end tracking primarily applies to single-camera scenarios. In the domain of multi-camera tracking, particularly when dealing with highly overlapping camera views, there has been little work in this direction.

In this paper, we propose the Multi Camera Tracking tRansformer (MCTR), a novel approach to multi-camera multi-object tracking that adopts an end-to-end architecture to track multiple objects across multiple camera feeds. MCTR builds on end-to-end object detection models such as DETR (DEtectors TRansformer) [7], adding two new components to facilitate multi-camera tracking: a tracking module and an association module. The flow of MCTR is conceptually simple: the object detectors are independently applied on every camera feed to generate view-specific detections and detection embeddings; the tracking module maintains a set of track embeddings and updates them using the information in the view-specific detection embeddings; and the association module generates an assignment of detections to tracks based on the respective detection and track embeddings. The model is trained end-to-end using a new loss formulation. This approach offers a more principled solution to multi-camera tracking, reducing the reliance on heuristic components and integrating the entire tracking process into a coherent, data-driven framework.

Multi-camera tracking introduces additional complexity over single-camera tracking as track consistency needs to be maintained both across time and across camera views. This makes it difficult to directly extend single-camera tech-

niques, that generally depend on the assumption that a track may only be associated with a single detection at every time step, like, for example, those employed by MOTR [52] which use the track embeddings as detection queries that are iteratively updated by DETR. Instead, MCTR maintains a separate set of track embeddings that encapsulate global information about the tracked objects across all views, and are distinct from the DETR object embeddings, which encapsulate local, view-specific information.

As is the case with all transformer-based models, the track embeddings as well as the detection embeddings do not have a particular order, and, thus, do not have a fixed relation with ground truth labels. This makes it challenging to devise a loss for training the model. DETR solves this issue for object detection by first finding the best assignment of detection embeddings to ground truth annotations using the Hungarian algorithm [25], and then computing classification and detection losses based on this assignment. MOTR and other approaches extend this idea to single-camera tracking by forcing each track query that has been previously associated with a ground truth object to keep detecting the same object in future frames. It is unclear, however, how to extend this procedure to keep track assignments consistent not only across time, but also across camera views. Furthermore, this approach is highly dependent on having a good association between detections and ground truth in the first frame, as it can not be changed in subsequent frames. In this paper we take a different approach to devise a training loss that avoids these issues. To calculate view-specific classification and detection losses we follow DETR’s approach and use the Hungarian algorithm in every frame and camera view independently to locally associate detections with ground truth. To maintain global tracking consistency, we use an approach akin to the attention mechanism in scaled dot-product attention to generate probabilistic assignments of local detections to global track embeddings. For pairs of detections in different cameras or different frames we calculate the probability, under the model, that the detections belong to the same track by integrating over the track assignment. We then use a negative log-likelihood loss to train the model to assign high probabilities to detection pairs from the same ground truth track and low probabilities to pairs from different tracks.

In summary, our contributions are: 1) An end-to-end framework for tracking multiple objects in multiple cameras with overlapping fields of view, 2) A method for probabilistic association of detections to tracks that is fully differentiable w.r.t input RGB frames, and 3) The formulation of a specialized loss function designed to guide the model in preserving consistent object identities, both temporally and across diverse camera perspectives.

The primary objective of this study is to test the feasibility of employing transformer-based end-to-end architec-

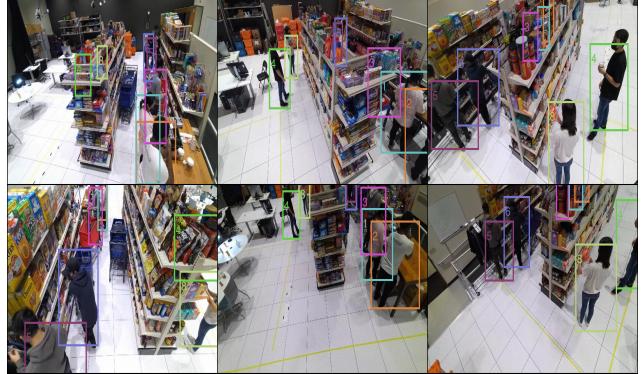


Figure 1. Example of multi-camera frame from the MMPTTrack dataset with 6 camera angles.

tures within the realm of multi-camera multi-object tracking problems. Our aim is to demonstrate the potential viability of such techniques and their applicability to these complex scenarios. We anticipate that our findings will inspire further exploration and research in this direction, potentially unlocking the full capabilities of these models in the future.

## 2. Related Works

### 2.1. Single Camera Tracking

Multi-object tracking within a single-camera setup has been extensively studied. Tracking by detection approaches [6, 27, 40, 46, 54] utilize object detectors [11, 15, 36], to identify objects in individual frames and perform track association using Kalman Filter [45] or Hungarian Matching [25]. Recently, end to end tracking methods such as [5, 31, 42, 51, 52, 55–57] have emerged, extending query based object detection for tracking [7]. For example, MOTR [52], MOTRv2 [55] and TrackFormer [31] propagate track queries across frames, iteratively updating them with image features for long-term detection and tracking. TransTrack [42] and P3Aformer [56] utilize a location-based cost matrix for bipartite matching [25]. In contrast to tracking by detection methods, end-to-end tracking methods are data-driven and avoid dependence on handcrafted heuristics. Despite their effectiveness, single-camera tracking encounters challenges, especially in cluttered and crowded environments with occlusions. The reliance on a single viewpoint limits the system’s robustness in complex scenarios. To alleviate this, we extend the tracking by query propagation paradigm for multi-camera multi-object tracking.

### 2.2. Multi-Camera Tracking

The multi-camera tracking (MCT) domain has witnessed diverse methodologies to tackle challenges in complex surveillance environments. A prevalent strategy involves a distributed approach, where single-camera tracking precedes hierarchical clustering [32], non-negative matrix fac-

torization (NMF) [44], or other merging and association algorithms for tracklets [4, 8, 20, 23, 24, 26, 35, 39, 41, 47, 48]. Alternatively, global or centralized methods bypass single-camera tracking, focusing on detecting individuals in each camera view and subsequently globally associating detections into tracklets [29, 49]. Occupancy maps derived from multi-camera views have been explored [12, 50], and trajectory prediction is employed in certain works, with variations in global or distributed approaches for multi-view association [22, 49]. Iguernaissi et al. [21] provides a comprehensive survey of the MCT papers.

Notable contributions in MCT include TRACTA [20], utilizing RNMF for cross-camera tracklet matching, and DMCT [50], incorporating a perspective-aware Ground-Point Network, occupancy heatmap estimation, and a glimpse network for person detection and tracking. Recently, ReST [9] proposed a two-stage association method using a reconfigurable graph model. In the vision-based autonomous driving systems domain, efforts focus on multi-view 3D object tracking, building upon end-to-end multi-view 3D object detection methodologies [10, 43]. Works such as MUTR3D [53], PF-Track [34], ViP3D [16], and DQTrack [28] utilizes a 3D track query-based approach for coherent object tracking across multiple cameras and frames. DQTrack [28] introduces a decoupled-query paradigm for camera-based 3D multi-object tracking, and ViP3D [16] pioneers a fully differentiable trajectory prediction approach. PF-Track [34] emphasizes spatio-temporal continuity.

In our approach, we conduct global track association for detections in each individual view, departing from single-camera tracklet-to-tracklet association. Diverging from existing global association methods, our fully end-to-end trainable framework based on transformers avoids heuristics or handcrafted pipelines for data-driven learning. Our track association loss is fully differentiable and trained end-to-end, similar to 3D object tracking. However, we focus on 2D person tracking across overlapping cameras.

### 3. Method

In this section, we provide a detailed description of the architecture of the proposed model. Figure 2 offers an overview of our model, highlighting its key components and their interactions.

The system consists of a detection module, a tracking module, and an association module. The detection module operates independently for each camera view and generates a set of detection embeddings. These embeddings encapsulate essential information regarding object detections within each individual view and are used to generate a bounding box prediction, a class prediction, and to inform the downstream modules. The tracking module maintains a set of track embeddings. The track embeddings contain global in-

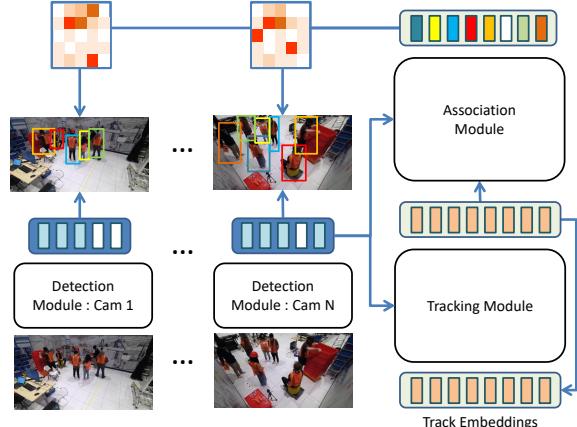


Figure 2. Model Overview.

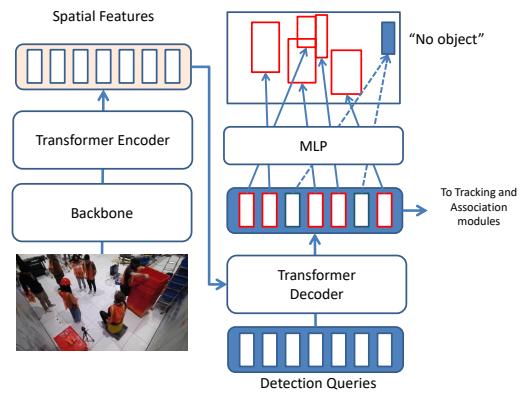


Figure 3. Detection Module: DETR.

formation about each object and are used to maintain consistent and coherent object identities across camera views and time. The track embeddings are updated at each frame using the information in the detection embeddings from all the camera views. The association module is tasked with producing a probabilistic assignment of detections to identities based on the information in the respective track and detection embeddings.

#### 3.1. Detection Module

In this paper we employ a vanilla DETR (DETection TRAnsformer) for the detection module, but any other end-to-end detector architecture can be used. A DETR model operates independently on each camera view. Briefly, the architecture of the DETR model includes a backbone to extract image features and a transformer encoder and decoder. The spatial features at the output of the backbone are added to a positional encoding and input to the encoder transformer which performs self attention and outputs a set

of spatial features. The DETR transformer decoder takes a set of detection query embeddings as input and outputs a set of detection embeddings. Each of the output embeddings are then passed through an MLP to predict the object class (or a "no object" class) and its bounding box. An overview of DETR is shown in figure 3. For more details see [7].

### 3.2. Tracking Module

The tracking module is tasked with updating the track embeddings with information from all the camera views in the current frame. The purpose of the track embeddings is to maintain global information about the tracked objects. The architecture of the tracking module is depicted in figure 4.

The track embeddings are initialised at the first frame to learned "track query embeddings". Like the "query inputs" to DETR, some of the embeddings will take on the role of representing one identity over time and others will remain unassigned. When a track embedding remains unassigned to a detection after a certain time, it is reset to its initial embedding. To update the track embeddings with detections from each camera, the cross-attention section has a set of multi-head cross-attention modules tailored to each camera view, with each module computing cross-attention between the current track embeddings and the detection embeddings from the corresponding view. The track embeddings act as "queries", and detection information as "keys" and "values". Because the relationship between objects and tracks depends on the camera position, each cross-attention module has its own distinct parameters (i.e. there is no parameter sharing between the cross-attention modules corresponding to the different views). The outputs of all the view specific cross-attention modules are averaged and passed through self-attention and feed-forward layers to obtain updated track embeddings.

The self-attention module in the tracking module is intended to introduce competition between the track embeddings. If one track query 'claims' a particular identity strongly, other track queries can "observe" this and are encouraged not to take on the same identity through the attention mechanism. This is similar to the way self attention works in the decoder for DETR to discourage 2 query embeddings from claiming the same output detection.

### 3.3. Association Module

The association module, depicted in figure 5, produces a probabilistic assignment of detections to tracks. The assignment is performed independently for each camera view through a mechanism that is the same as the attention mechanism in scaled dot-product attention, wherein the detections act as queries and tracks act as keys. The detection embeddings and track embeddings undergo a linear transformation, followed by the multiplication of the resulting matrices. Similar to scaled attention, the result is scaled

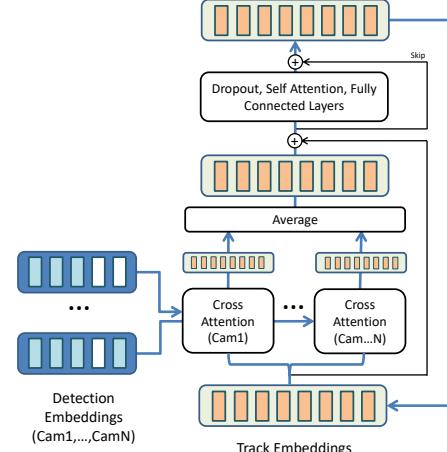


Figure 4. Tracking Module.

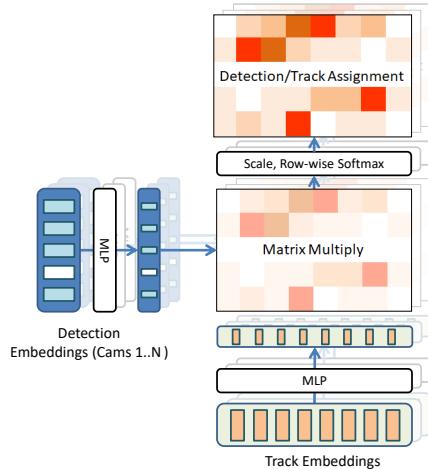


Figure 5. Association Module.

by the square root of the embedding dimension before applying a row-wise softmax operation. The result is a  $D \times T$  matrix  $A^v$  where each entry  $A_{d,t}^v$  represents the probability that detection  $d$  is associated with track  $t$  in view  $v$ .  $D$  is the number of detection queries while  $T$  is the number of track queries.

### 3.4. Training Loss

To train the model we use several loss functions, that we group into detection losses, track losses and auxiliary track losses.

**Detection loss.** The detection loss,  $\mathcal{L}_{det}$ , is the same as the one used by [7] to train DETR: a combination of negative log-likelihood classification loss, an IOU loss and a GIOU loss. Given ground truth annotations, Hungarian matching is used to find the bipartite assignment of detections to ground truth that incurs the lowest loss. This loss is used as a detection loss for each view. Following [7] we also use losses calculated at different layers of the DETR

transformer decoder as auxiliary losses,  $\mathcal{L}_{det\_aux}$ .

**Track losses.** The track losses enforce consistency of object identity between camera views and camera frames. Given two detections  $d_1$  and  $d_2$  from two different views  $v_1$  and  $v_2$ , the probability, according to the model, that  $d_1$  and  $d_2$  belong to the same track ID can be calculated by integrating over tracks:

$$P_{st}(d_1, d_2) = \sum_t A_{d_1, t}^{v_1} \cdot A_{d_2, t}^{v_2}$$

where  $A^{v_1}$  and  $A^{v_2}$  are the probabilistic assignment matrices for views  $v_1$  and  $v_2$  calculated by the association module. (See section 3.3.)

To obtain a label,  $y_{st}(d_1, d_2)$ , for a pair of detections, we leverage the detection to ground truth assignment provided by the Hungarian matching algorithm described above. We have three possible cases: 1) both  $d_1$  and  $d_2$  are associated with a ground truth annotation, and both ground truth annotations have the same track ID. In this case  $y_{st}(d_1, d_2) = 1$ ; 2) both  $d_1$  and  $d_2$  are associated with a ground truth annotation, but both ground truth annotations have the different track IDs. In this case  $y_{st}(d_1, d_2) = 0$ ; and 3) either  $d_1$  or  $d_2$  is not associated with any ground truth annotation. In this case  $y_{st}(d_1, d_2)$  is undefined.

The loss  $\mathcal{L}_{across\_cam}$  is the negative log-likelihood of the labels  $y_{st}$ , when they are defined:

$$\begin{aligned} \mathcal{L}_{across\_cams} = & -\frac{1}{N_p} \sum_{d_1, d_2} (y_{st}(d_1, d_2) \cdot \log(P_{st}(d_1, d_2)) \\ & + (1 - y_{st}(d_1, d_2)) \cdot \log(1 - P_{st}(d_1, d_2))) \end{aligned}$$

where the sum is taken over pairs of detections  $d_1$  and  $d_2$  from different camera views in the same frame, for which  $y_{st}(d_1, d_2)$  is defined.  $N_p$  is the number of such pairs.

Similarly, one can define  $\mathcal{L}_{across\_frames}$  by taking pairs of detections from the same camera view and different frames.

**Auxiliary track losses.** The probabilistic formulation of the assignment of detections to tracks enables the definition of auxiliary losses to encourage learning of more informative track embeddings. For example one can induce the track embeddings to encode information about the bounding boxes of corresponding object in all camera views. To achieve this the track embeddings are passed, for each view, through a three layer MLP with ReLU nonlinearity to predict the coordinates  $\hat{B}(t, v)$  of the bounding box for the respective track  $t$  in view  $v$ . The MLPs are view specific, since an object position would be different in different camera views. The track IOU loss is:

$$\mathcal{L}_{track\_IOU} = \frac{1}{V \cdot T} \sum_{v=1}^V \sum_{t=1}^T \sum_d A_{d, t}^v \cdot L_{IOU}(\hat{B}(t, v), B(d))$$

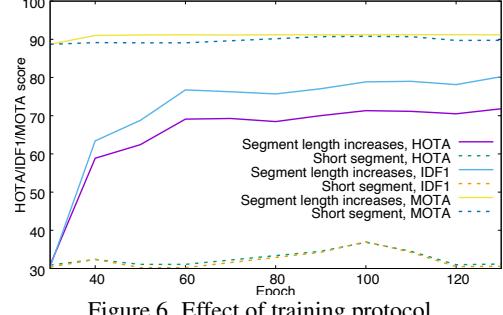


Figure 6. Effect of training protocol.

where the third sum is taken over all detections  $d$  that have been associated with a ground truth annotation by Hungarian matching for that view,  $B(d)$  is the bounding box of that ground truth annotation, and  $L_{IOU}$  is the IOU loss. Similarly one can define  $\mathcal{L}_{track\_GIOU}$  by replacing the IOU loss with the generalized IOU loss [37].

While the bounding boxes predicted from the track embeddings are not as accurate as the ones predicted from the detection embeddings, the auxiliary track losses serve an important role in ensuring consistency of object identities.

The final loss used to train the entire system end to end is:

$$\begin{aligned} \mathcal{L} = & \mathcal{L}_{det} + \mathcal{L}_{det\_aux} + \mathcal{L}_{across\_cams} + \mathcal{L}_{across\_frames} \\ & + \mathcal{L}_{track\_IOU} + \mathcal{L}_{track\_GIOU} \end{aligned}$$

### 3.5. Training Protocol

The training proceeds based on contiguous video segments which are sampled randomly from the training data. The video segments are split into non-overlapping four frame clips, with each clip serving as a training instance. At the beginning of the video segment the track embeddings are set to an initial embedding (which is learned). For each subsequent clip, the track embeddings are initialized with the final track embeddings from the previous clip.

Multi-camera tracking applications usually require systems to run continuously, and track targets over time spans of minutes or tens of minutes. This creates challenges for training an end-to-end system: if training is done only over short video segments, the mismatch between training and deployment conditions introduce a domain shift that may lower the performance. On the other hand, if training is performed over long video segments, the model would see very correlated data that lacks diversity and would tend to overfit to it. To address this problem we propose the following training protocol.

For the first thirty epochs the video segments are short, four frame clips. This stage of training is used to ensure that the model sees diverse data, which is especially important for the detector models. After this initial stage, the parameters of the detector model are frozen and the training of the

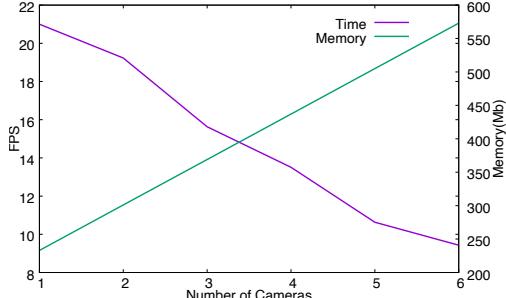


Figure 7. Inference time compute resources.

tracking and association modules continues on increasingly longer video segments. The length of the video segments is randomly chosen from a geometric distribution with the expected value increasing linearly as training progresses.

Figure 6 shows the benefit of this training protocol. As the length of the training video segments increases, performance on the test data improves, especially for the IDF1 and HOTA metrics which emphasize long term tracking accuracy. For MOTA score (which emphasizes detection accuracy) the performance remains relatively constant. For comparison, if the the training video segments are kept short, the IDF1 and HOTA metrics do not improve significantly, showing that the increase in performance is due to the devised training protocol, not simply due to extra training.

### 3.6. Inference

Inference on test data are made frame by frame in an online manner. In each view  $v$ , detections that have a predicted confidence above a threshold (0.9 for experiments in this paper) are kept and the rest are discarded. Then the Hungarian algorithm, with predicted association matrix  $A^v$  as the weight function, is used to find a bipartite matching between detections that have not been discarded and tracks. This aims to associate detections with tracks in a way that maximizes the total association probability. If a track has not been associated with any detection for more than four frames, its embedding is reset to the initial embedding. All other track embeddings remain the same.

Figure 7 shows the memory and time MCTR uses for inference as a function of the number of cameras. Both inference time and memory scale linearly with the number of cameras. With 1 camera, the model has about 44M parameters and takes 233Mb of memory and can run at 21 FPS on a single GeForce RTX 2080 Ti GPU. Each additional camera adds 1M parameters, 68Mb of memory and reduces the FPS by about 2. Thus MCTR would be well suited for real-time multi-camera tracking applications.

## 4. Experimental Results

We contrast MCTR with the single-camera end-to-end tracker MOTR [52], when the later is applied to each cam-

era view independently.<sup>1</sup> We also compare with ReST [9], a recent multi-stage multi-camera tracking approach. When using a pre-trained YOLO detector ReST is unable to maintain track identities over longer time frames leading to very low performance on our evaluation datasets. This highlights the shortcomings of multi-stage methods that require significant effort to ensure that all of the heuristic components are well tuned for the dataset. In contrast, end-to-end methods require much less manual tuning, learning to adapt to the characteristics of each dataset. To gauge tracking performance of ReST, with an ideal detector, we also report results for ReST when using the ground truth bounding boxes as detections (ReST-GT).

MOTR is trained on each camera view and scene using the default parameters suggested by the authors. During testing we discovered an interesting failure mode for MOTR. It tends to produce large number of highly overlapping duplicate detections and consequently tracks. This adversely affects the detection precision resulting in lower scores across all metrics especially MOTA [3]. To fix the issue, we apply non-maximum suppression (NMS) and filter boxes with high overlap using an IoU threshold of 0.95.

The models are evaluated using popular metrics in single-camera multi-object tracking literature: HOTA [30], IDF1 [38] and MOTA [3]. The metrics are computed independently for each camera view and test clip, and the results are averaged to get the final performance for each environment. This is similar to the protocol used in the MMP-Tracking challenge [2]. For MCTR and ReST we evaluate the cross-view association performance using the AIDF1 and MOAA metrics as proposed in [13, 17, 18]. AIDF1 is the association  $F_1$  score calculated as the geometric mean of precision and recall of predicted pairwise object associations averaged across all frames and all pairs of camera views. MOAA (Multi Object Association Accuracy) follows the calculation of MOTA, but ID switches are measured across all pairs of views rather than across frames.

### 4.1. MMPTtrack dataset

MMPTtrack [19] is a large scale multi-camera multi-person tracking dataset comprised of five environments: industry, retail, cafe, lobby and office. The data is densely labeled with the help of RGBD cameras and verified and corrected by human labelers.

One artifact of the way the MMPTtrack annotations have been collected is that ground-truth bounding boxes are given for all views, even if the person might be fully occluded in a particular view. In most environments occlu-

<sup>1</sup>While the performance of MOTR may be improved by using pre-trained object detectors to generate detection queries, as in MOTR v2 [55], the same techniques are equally applicable to MCTR. In fact MCTR can be used directly with detectors such as YOLO foregoing DETR altogether. However, this would distract from the goal of understanding the behaviour of end-to-end approaches.

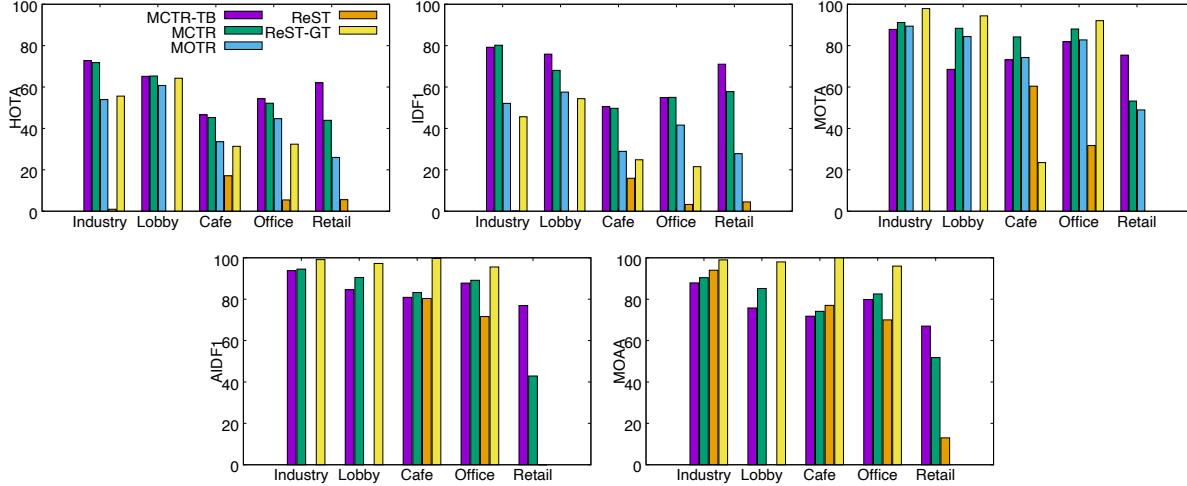


Figure 8. Results on MMPTtrack dataset.

sions are rare, making it less problematic. In the retail setting, however, a large fraction of people are occluded by shelves. Training a detector with these annotations would be ambiguous and lead to lower performance. To mitigate this issue, we filter out annotations of occluded persons: we employ a pre-trained YOLOX model [14] to detect people in each frame, and ground truth annotations with a lower IOU than 0.01 with a YOLOX detection are filtered out. While this approach may eliminate annotations for visible people that are undetected by YOLOX, it serves as a reasonable approximation. The models are evaluated on the original labels, so low recall is expected in this environment.

Figure 8 shows the performance of MCTR, MOTR, ReST and ReST-GT ( ReST with ground truth bounding boxes) on the validation clips from the five MMPTtrack environments <sup>2</sup>. These results represent averages across various cameras and clips. MCTR outperforms ReST and MOTR across all environments and metrics, showcasing the effectiveness of end-to-end methods for multi-camera tracking. Its higher IDF1 and HOTA scores highlight MCTR’s ability to use multi-view cues to handle occlusions and maintain consistent long-term tracks. Compared to ReST-GT, MCTR performs better on the HOTA and IDF1 metrics, showing that even with a perfect detector, ReST temporal tracking performance is lower than MOTR. On MOTA, MOAA and AIDF1, which weight detection accuracy more heavily, ReST-GT performs better than MCTR due to its perfect detector.

As mentioned above, in the retail environment there are a large number of ground-truth detections that are occluded in any particular camera view resulting in a significant number of missed detections. While a person might be occluded in one or more camera views, it is usually visible in at least

<sup>2</sup>ReST-GT inference did not finish on the Retail environment, and ReST did not finish on Lobby environment

one camera allowing MCTR the opportunity to detect that person. As discussed in section 3.4 MCTR is trained to predict bounding boxes for all views based on the global track embeddings. While these bounding boxes are primarily utilized for computing an auxiliary loss to enhance model training, they can be also used to predict bounding boxes for occluded people that can not be detected by the view-specific detector. When using the track-level bounding boxes for prediction (denoted as MCTR-TB in figure 8), the performance improves significantly on the retail environment, mainly driven by an increase in recall from 54.28 to 88.73. This is remarkable, because the model has not seen any ground truth annotations in the occluded regions during training. This suggests that the model has learned an approximate geometry and is able to extrapolate it to regions where it has not seen any training data. On the other environments, the HOTA and IDF1 performance remains similar, but there is a decrease in the MOTA score due to a somewhat worse detection bounding box accuracy.

Besides the ability to predict the location of occluded people, the auxiliary losses described in section 3.4 are instrumental in improving the long term tracking performance of MCTR. Table 1 shows the performance, on the industry environment, of MCTR with and without using the auxiliary loss. Training the model to predict detection bounding boxes from the global track embeddings leads to significant improvements in the tracking focussed HOTA and IDF1 scores, with a slight decrease in the more detection focused metrics.

To the best of our knowledge, state of the art performance on the MMPTtrack dataset has been achieved during the MMPT-Tracking challenge [2]. The highest MOTA score on the leaderboard is 87, averaged across all five environments, and the highest IDF1 score is 88.18 [1]. Besides using a plethora of engineered heuristics, both these

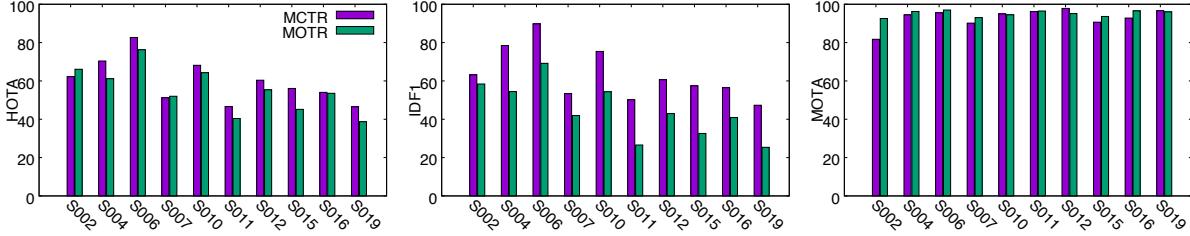


Figure 9. Results on AI City Challenge dataset.

Table 1. Performance of MCTR with and without using auxiliary losses for training

Model	HOTA	IDF1	MOTA	AIDF1	MOAA
MCTR	71.81	80.21	91.19	94.52	90.39
MCTR-noaux	68.00	75.26	92.11	95.00	91.30

Table 2. Comparison with top methods in the MMP-Tracking challenge, camera view evaluation. The ranks MCTR would have achieved for each metric are shown in parentheses.

Model	IDF1	MOTA	Online	Camera parameters	FPS
Hikvision	86.36	87	No	Yes	-
Alibaba	88.18	78	No	Yes	< 1
MCTR	62.30 (5)	81.01 (3)	Yes	No	9

methods are offline (i.e. they can use information from the future), and make heavy use of camera calibration parameters. MCTR is online, and does not use camera parameters. With a combined MOTA of 81.01 and a combined IDF1 of 62.30, MCTR would rank third in MOTA and fifth in IDF1.<sup>3</sup> While the MOTA performance is respectable, there is a large gap in the IDF1 performance. This indicates that maintaining consistent tracking over a long period of time is still challenging for an end-to-end method.

#### 4.2. AI City Challenge dataset

We also evaluate on data from 2023 AI City Challenge [33] Track 1, a synthetic dataset containing several indoor scenes with multiple overlapping cameras. Because our approach requires the same camera setup for training and test we deviate from the challenge setup, and instead use the first 70% of the clip for each scene for training and the rest of 30% for testing. We down-sample the resolution to 640x360 and use a 15 FPS frame rate.

Figure 9 shows the results for MCTR and MOTR on the AI City Challenge dataset.<sup>4</sup> MCTR again outperforms MOTR on IDF1 and HOTA metrics indicating improved long-term tracking and robustness to occlusions. MOTR

<sup>3</sup>Our performance is measured on the validation set as we do not have access to the test set. We assume that validation and test performances are comparable.

<sup>4</sup>We do not run ReST on this dataset because of its poor performance on MMPTrack, and the lack of ground truth camera calibration parameters which are required by ReST.

on the other hand tends to have better MOTA score, mainly due to a higher detection precision and recall. The improved detection performance of MOTR is likely due to their use of deformable DETR [58] which is known to perform better than standard DETR at detecting small objects, and to a prevalence of highly occluded persons (e.g. only head is visible behind shelves) in the dataset leading to a significant number of small objects.

#### 5. Conclusion & Limitations

In this paper we presented the Multi-Camera Tracking tRansformer (MCTR) a novel architecture that integrates detection and tracking across multiple cameras with overlapping fields of view in one coherent, end-to-end trainable system. MCTR is build on three key ideas: using separate track and detection embeddings, learning a probabilistic association between tracks and detections, and using a loss function based on the model’s prediction of pairs of detections belonging to the same track. MCTR outperforms end-to-end single camera approaches on the MMPTrack and AI City Challenge datasets, generating robust tracks that are consistent across time as well as camera views.

Compared to highly optimized heuristic systems that led the MMP-Tracking competition, MCTR does lag behind, particularly in IDF1 score. The lower long-term performance may stem from MCTR’s Markovian approach, where track embeddings are updated only with current and previous frame information. Early attempts to use longer temporal history did not improve results. Finding better ways to integrate more temporal and motion information could enhance long-term tracking performance.

Another limitation of MCTR, and of end-to-end systems in general, is the need for extensive labeled data for training. Moreover, MCTR’s reliance on camera-view specific cross-attention modules means it is closely tied to the particular camera setup used during training. An interesting question is whether a more general, camera setup independent, end-to-end model can be devised. This could involve using camera calibration parameters to adapt a single cross-attention module for all views, rather than separate modules. We leave this direction for future work.

## References

[1] "https://competitions.codalab.org/competitions/33729#results". 7

[2] https://iccv2021-mmp.github.io/. 6, 7

[3] Keni Bernardin and Rainer Stiefelhagen. Evaluating multiple object tracking performance: the clear mot metrics. *EURASIP Journal on Image and Video Processing*, 2008:1–10, 2008. 6

[4] James Black, Tim Ellis, and Paul Rosin. Multi view image surveillance and tracking. In *Workshop on Motion and Video Computing, 2002. Proceedings.*, pages 169–174. IEEE, 2002. 3

[5] Jiarui Cai, Mingze Xu, Wei Li, Yuanjun Xiong, Wei Xia, Zhiwen Tu, and Stefano Soatto. Memot: Multi-object tracking with memory. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8090–8100, 2022. 1, 2

[6] Jinkun Cao, Jiangmiao Pang, Xinshuo Weng, Rawal Khirodkar, and Kris Kitani. Observation-centric sort: Rethinking sort for robust multi-object tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9686–9696, 2023. 2

[7] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020. 1, 2, 4

[8] Weihua Chen, Lijun Cao, Xiaotang Chen, and Kaiqi Huang. An equalized global graph model-based approach for multi-camera object tracking. *IEEE Transactions on Circuits and Systems for Video Technology*, 27(11):2367–2381, 2016. 3

[9] Cheng-Che Cheng, Min-Xuan Qiu, Chen-Kuo Chiang, and Shang-Hong Lai. Rest: A reconfigurable spatial-temporal graph model for multi-camera multi-object tracking, 2023. 3, 6

[10] Simon Doll, Richard Schulz, Lukas Schneider, Viviane Benzin, Markus Enzweiler, and Hendrik PA Lensch. Spatialdetr: Robust scalable transformer-based 3d object detection from multi-view camera images with global cross-sensor attention. In *European Conference on Computer Vision*, pages 230–245. Springer, 2022. 3

[11] Kaiwen Duan, Song Bai, Lingxi Xie, Honggang Qi, Qingming Huang, and Qi Tian. Centernet: Keypoint triplets for object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6569–6578, 2019. 2

[12] Francois Fleuret, Jerome Berclaz, Richard Lengagne, and Pascal Fua. Multicamera people tracking with a probabilistic occupancy map. *IEEE transactions on pattern analysis and machine intelligence*, 30(2):267–282, 2007. 3

[13] Yiyang Gan, Ruize Han, Liqiang Yin, Wei Feng, and Song Wang. Self-supervised multi-view multi-human association and tracking. In *Proceedings of the 29th ACM international conference on multimedia*, pages 282–290, 2021. 6

[14] Zheng Ge, Songtao Liu, Feng Wang, Zeming Li, and Jian Sun. Yolox: Exceeding yolo series in 2021. *arXiv preprint arXiv:2107.08430*, 2021. 7

[15] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015. 2

[16] Junru Gu, Chenxu Hu, Tianyuan Zhang, Xuanyao Chen, Yilun Wang, Yue Wang, and Hang Zhao. Vip3d: End-to-end visual trajectory prediction via 3d agent queries. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5496–5506, 2023. 3

[17] Ruize Han, Wei Feng, Yujun Zhang, Jiewen Zhao, and Song Wang. Multiple human association and tracking from ego-centric and complementary top views. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(9):5225–5242, 2021. 6

[18] Ruize Han, Wei Feng, Jiewen Zhao, Zicheng Niu, Yujun Zhang, Liang Wan, and Song Wang. Complementary-view multiple human tracking. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 10917–10924, 2020. 6

[19] Xiaotian Han, Quanzeng You, Chunyu Wang, Zhizheng Zhang, Peng Chu, Houdong Hu, Jiang Wang, and Zicheng Liu. Mmptrack: Large-scale densely annotated multi-camera multiple people tracking benchmark. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 4860–4869, 2023. 6

[20] Yuhang He, Xing Wei, Xiaopeng Hong, Weiwei Shi, and Yihong Gong. Multi-target multi-camera tracking by tracklet-to-target assignment. *IEEE Transactions on Image Processing*, 29:5191–5205, 2020. 3

[21] Rabah Iguernaissi, Djamel Merad, Kheireddine Aziz, and Pierre Drap. People tracking in multi-camera systems: a review. *Multimedia Tools and Applications*, 78:10773–10793, 2019. 3

[22] Yuntae Jeon, Dai Quoc Tran, Minsoo Park, and Seunghee Park. Leveraging future trajectory prediction for multi-camera people tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 5398–5407, June 2023. 3

[23] Xiaoyan Jiang, Erik Rodner, and Joachim Denzler. Multi-person tracking-by-detection based on calibrated multi-camera systems. In *Computer Vision and Graphics: International Conference, ICCVG 2012, Warsaw, Poland, September 24–26, 2012. Proceedings*, pages 743–751. Springer, 2012. 3

[24] Philipp Kohl, Andreas Specker, Arne Schumann, and Jurgen Beyerer. The mta dataset for multi-target multi-camera pedestrian tracking by weighted distance aggregation. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2020. 3

[25] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955. 2

[26] Quoc Cuong Le, Donatello Conte, and Moncef Hidane. Online multiple view tracking: Targets association across cameras. In *6th Workshop on Activity Monitoring by Multiple Distributed Sensing (AMMDS 2018)*, 2018. 3

[27] Laura Leal-Taixé, Cristian Canton-Ferrer, and Konrad Schindler. Learning by tracking: Siamese cnn for robust target association. In *Proceedings of the IEEE conference on*

computer vision and pattern recognition workshops, pages 33–40, 2016. 2

[28] Yanwei Li, Zhiding Yu, Jonah Philion, Anima Anandkumar, Sanja Fidler, Jiaya Jia, and Jose Alvarez. End-to-end 3d tracking with decoupled queries. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 18302–18311, 2023. 3

[29] Wenqian Liu, Octavia Camps, and Mario Sznajer. Multi-camera multi-object tracking. *arXiv preprint arXiv:1709.07065*, 2017. 3

[30] Jonathon Luiten, Aljosa Osep, Patrick Dendorfer, Philip Torr, Andreas Geiger, Laura Leal-Taixé, and Bastian Leibe. Hota: A higher order metric for evaluating multi-object tracking. *International journal of computer vision*, 129:548–578, 2021. 6

[31] Tim Meinhardt, Alexander Kirillov, Laura Leal-Taixé, and Christoph Feichtenhofer. Trackformer: Multi-object tracking with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8844–8854, 2022. 1, 2

[32] Fionn Murtagh and Pedro Contreras. Algorithms for hierarchical clustering: an overview. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 2(1):86–97, 2012. 2

[33] Milind Naphade, Shuo Wang, David C. Anastasiu, Zheng Tang, Ming-Ching Chang, Yue Yao, Liang Zheng, Mohammed Shaiqur Rahman, Meenakshi S. Arya, Anuj Sharma, Qi Feng, Vitaly Ablavsky, Stan Sclaroff, Pranamesh Chakraborty, Sanjita Prajapati, Alice Li, Shangru Li, Krishna Kunadharaju, Shenxin Jiang, and Rama Chellappa. The 7th ai city challenge. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2023. 8

[34] Ziqi Pang, Jie Li, Pavel Tokmakov, Dian Chen, Sergey Zagoruyko, and Yu-Xiong Wang. 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*, pages 17928–17938, 2023. 3

[35] Kha Gia Quach, Pha Nguyen, Huu Le, Thanh-Dat Truong, Chi Nhan Duong, Minh-Triet Tran, and Khoa Luu. Dyglip: A dynamic graph model with link prediction for accurate multi-camera multiple object tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13784–13793, 2021. 3

[36] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018. 2

[37] Hamid Rezatofighi, Nathan Tsoi, JunYoung Gwak, Amir Sadeghian, Ian Reid, and Silvio Savarese. Generalized intersection over union: A metric and a loss for bounding box regression. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 658–666, 2019. 5

[38] Ergys Ristani, Francesco Solera, Roger Zou, Rita Cucchiara, and Carlo Tomasi. Performance measures and a data set for multi-target, multi-camera tracking. In *European conference on computer vision*, pages 17–35. Springer, 2016. 6

[39] Ergys Ristani and Carlo Tomasi. Features for multi-target multi-camera tracking and re-identification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6036–6046, 2018. 3

[40] Samuel Schulter, Paul Vernaza, Wongun Choi, and Manmohan Chandraker. Deep network flow for multi-object tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6951–6960, 2017. 2

[41] Andreas Specker, Daniel Stadler, Lucas Florin, and Jürgen Beyerer. An occlusion-aware multi-target multi-camera tracking system. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4173–4182, 2021. 3

[42] Peize Sun, Jinkun Cao, Yi Jiang, Rufeng Zhang, Enze Xie, Zehuan Yuan, Changhu Wang, and Ping Luo. Transtrack: Multiple object tracking with transformer. *arXiv preprint arXiv:2012.15460*, 2020. 1, 2

[43] Yue Wang, Vitor Campagnolo Guizilini, Tianyuan Zhang, Yilun Wang, Hang Zhao, and Justin Solomon. Detr3d: 3d object detection from multi-view images via 3d-to-2d queries. In *Conference on Robot Learning*, pages 180–191. PMLR, 2022. 3

[44] Yu-Xiong Wang and Yu-Jin Zhang. Nonnegative matrix factorization: A comprehensive review. *IEEE Transactions on knowledge and data engineering*, 25(6):1336–1353, 2012. 3

[45] Greg Welch, Gary Bishop, et al. An introduction to the kalman filter. 1995. 2

[46] Nicolai Wojke, Alex Bewley, and Dietrich Paulus. Simple online and realtime tracking with a deep association metric. In *2017 IEEE international conference on image processing (ICIP)*, pages 3645–3649. IEEE, 2017. 2

[47] Yuanlu Xu, Xiaobai Liu, Yang Liu, and Song-Chun Zhu. Multi-view people tracking via hierarchical trajectory composition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4256–4265, 2016. 3

[48] Senquan Yang, Fan Ding, Pu Li, and Songxi Hu. Distributed multi-camera multi-target association for real-time tracking. *Scientific Reports*, 12(1):11052, 2022. 3

[49] Xipeng Yang, Jin Ye, Jincheng Lu, Chenting Gong, Minyue Jiang, Xiangru Lin, Wei Zhang, Xiao Tan, Yingying Li, Xiaoqing Ye, et al. Box-grained reranking matching for multi-camera multi-target tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3096–3106, 2022. 3

[50] Quanzeng You and Hao Jiang. Real-time 3d deep multi-camera tracking. *arXiv preprint arXiv:2003.11753*, 2020. 3

[51] En Yu, Tiancai Wang, Zhuoling Li, Yuang Zhang, Xiangyu Zhang, and Wenbing Tao. Motrv3: Release-fetch supervision for end-to-end multi-object tracking. *arXiv preprint arXiv:2305.14298*, 2023. 2

[52] Fangao Zeng, Bin Dong, Yuang Zhang, Tiancai Wang, Xiangyu Zhang, and Yichen Wei. Motr: End-to-end multiple-object tracking with transformer. In *European Conference on Computer Vision*, pages 659–675. Springer, 2022. 1, 2, 6

[53] Tianyuan Zhang, Xuanyao Chen, Yue Wang, Yilun Wang, and Hang Zhao. Mutr3d: A multi-camera tracking framework via 3d-to-2d queries. In *Proceedings of the IEEE/CVF*

*Conference on Computer Vision and Pattern Recognition*,  
pages 4537–4546, 2022. 3

- [54] Yifu Zhang, Peize Sun, Yi Jiang, Dongdong Yu, Fucheng Weng, Zehuan Yuan, Ping Luo, Wenyu Liu, and Xinggang Wang. Bytetrack: Multi-object tracking by associating every detection box. In *European Conference on Computer Vision*, pages 1–21. Springer, 2022. 2
- [55] Yuang Zhang, Tiancai Wang, and Xiangyu Zhang. Motrv2: Bootstrapping end-to-end multi-object tracking by pre-trained object detectors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22056–22065, 2023. 2, 6
- [56] Zelin Zhao, Ze Wu, Yueqing Zhuang, Boxun Li, and Jiaya Jia. Tracking objects as pixel-wise distributions. In *European Conference on Computer Vision*, pages 76–94. Springer, 2022. 2
- [57] Tianyu Zhu, Markus Hiller, Mahsa Ehsanpour, Rongkai Ma, Tom Drummond, Ian Reid, and Hamid Rezatofighi. Looking beyond two frames: End-to-end multi-object tracking using spatial and temporal transformers. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. 1, 2
- [58] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection, 2021. 8