

# TDT: Teaching Detectors to Track without Fully Annotated Videos

Shuzhi Yu<sup>1\*</sup> Guanhang Wu<sup>2</sup> Chunhui Gu<sup>2</sup> Mohammed E. Fathy<sup>2</sup>  
<sup>1</sup>Duke University <sup>2</sup>Google LLC

shuzhiyu@cs.duke.edu {guanhangwu, chunhui, msalem}@google.com

## Abstract

Recently, one-stage trackers that use a joint model to predict both detections and appearance embeddings in one forward pass received much attention and achieved state-of-the-art results on the Multi-Object Tracking (MOT) benchmarks. However, their success depends on the availability of videos that are fully annotated with tracking data, which is expensive and hard to obtain. This can limit the model generalization. In comparison, the two-stage approach, which performs detection and embedding separately, is slower but easier to train as their data are easier to annotate. We propose to combine the best of the two worlds through a data distillation approach. Specifically, we use a teacher embedder, trained on Re-ID datasets, to generate pseudo appearance embedding labels for the detection datasets. Then, we use the augmented dataset to train a detector that is also capable of regressing these pseudo-embeddings in a fully-convolutional fashion. Our proposed one-stage solution matches the two-stage counterpart in quality but is 3 times faster. Even though the teacher embedder has not seen any tracking data during training, our proposed tracker achieves competitive performance with some popular trackers (e.g. JDE) trained with fully labeled tracking data.

## 1. Introduction

Tracking people and objects in videos is an important task in computer vision and at the core of many applications. For tracking to be achieved, several sub-tasks need to be solved. The system has to identify (potentially many) objects of interest in the video, *i.e.* *object detection*, and it has to relate the locations of these objects as they move in the video, *i.e.* *data association*. Many approaches have been and continue to be proposed for solving the tracking problem [2, 6, 41, 51, 55, 61, 68]. One popular approach is tracking by detection. In this approach, detection is made on each video frame and the tracker associates the detec-

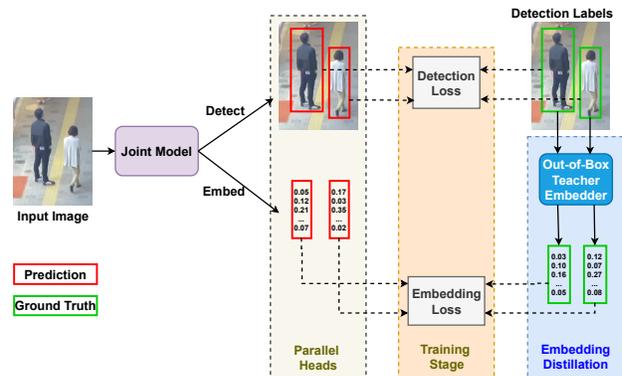


Figure 1. The training stage of our proposed one-stage tracking framework using only detection boxes as labels. We augment a detection example (top right image) with pseudo ground-truth embeddings (shown as the embedding vector in green) through a data distillation approach. These teacher embeddings are obtained by feeding the cropped image patches (green bounding boxes) into a teacher embedder (a pre-trained out-of-box Re-ID network). A joint model takes an input image, and jointly detects objects of interest (red bounding boxes) and generates an embedding (red embedding) for each object. The augmented dataset is used to supervise both predictions. This system is so flexible that different detectors and teacher embedders can be used.

tions from consecutive frames using various cues such as motion and appearance (a.k.a. appearance embeddings).

Many of the typical tracking algorithms use separately trained and operated components for detection and embedding [2, 6, 41, 55]. Recently, one-stage trackers that use a single convolutional neural network (CNN) for joint detection and appearance embedding were introduced [31, 51, 61]. As the detections and their embeddings are generated in one forward pass, these approaches are relatively faster than traditional two-stage ones. In addition, they achieve state-of-the-art performance on many Multi-Object Tracking (MOT) benchmarks [51, 61].

However, the success of these joint models depends heavily on the availability of fully annotated tracking data including both detection and identification labels. They need both labels to simultaneously supervise the detection

\*This work was done when Shuzhi Yu was an intern at Google

and embedding components. Such datasets are more restricted in nature and quantity than typical detection-only and identity-only datasets as they are much harder and more expensive to annotate. The lack of fully-annotated videos from diverse conditions can negatively impact the ability of one-stage trackers to generalize. In particular, they work very well on surveillance videos (*i.e.* the same data that is used to train them) but fail on videos from media and other domains where large people are likely to appear.

Reconsidering the two-stage models, although they are relatively slower, they do not require the expensive tracking labels for training. The two components, namely detector and embedder, are trained separately on their specialized datasets [58]. Such decoupling makes it possible to use less restrictive and more diverse detection and identification training sets, leading to better generalization of the individual components and the overall tracking system.

Accordingly, we propose in this paper a practical and effective approach (Fig. 1) that combines the speed of joint models with the lower training cost and better generalization of two-stage ones. In particular, we introduce a weak supervision framework that can train a joint model without fully annotated videos. The framework follows the teacher-student learning scheme [18,40] for training the embedders, *i.e.* embedding distillation. This is done by augmenting a given detection dataset with pseudo ground-truth embeddings generated by a Re-Identification (Re-ID) model independently trained on Re-ID datasets [65,66]. During training, the detector training loss function is augmented with an embedding distillation loss term to make sure the detection and embedding heads of the joint model are adequately supervised. This way, any detection dataset can be converted to a detection and embedding distillation one that can be used to train a joint model. Our experimental results suggest that training our model on the additional augmented detection dataset improves the overall detection and tracking performance (see Fig. 5). Furthermore, we show empirically that the overall tracking performance is enhanced with better teacher embedders (Tab. 3).

In addition, the training of the embedding is simplified in our weakly supervised framework. In general, there are two losses on embedding, namely the triplet loss [42,47,53] and the cross-entropy loss. The implementation of the triplet loss can be tricky since effectively selecting such triplets from a large sampling space is a non-trivial problem [42]. The cross-entropy loss models each person in the training videos as a class, which can make training harder and slower as we increase the number of labeled people in the training set. In comparison, our weakly supervised loss is straightforward and easy to implement. A simple least squares error between the predicted and pseudo ground truth embedding is shown to work well, without any sophisticated weighting algorithm to balance the detection and

embedding losses.

Another potential way of avoiding fully annotated tracking data is to alternatively train the detection and embedding heads of the joint model on detection and Re-ID datasets, respectively. Compared to our approach, this leads to a much more complicated training process involving more hyperparameters (*e.g.* learning rates, data sampling schemes, training schedules, *etc.* for different heads). In particular, it is not clear how to pick generally good settings that would work well for different combinations of datasets. Our method avoids these issues by using distillation to generate a single dataset, instead of two radically different ones.

To our best knowledge, this is the first work that trains joint tracking models without using any fully annotated tracking data. Furthermore, our tracker, named TDT-tracker (*i.e.* Teaching Detector to Track), runs 3 times faster than the counterpart two-stage model and stays competitive on the benchmark MOT challenges.

In the rest of the paper, Sec. 2 discusses the related work, Sec. 3 explains the proposed approach, Sec. 4 shows the empirical evaluation of our model, and Sec. 5 ends the paper with a conclusion.

## 2. Related work

**Two-stage tracking system:** Many methods took public detections [8, 14, 57] as input and focused on improving the data association step [41, 55]. They used historical states in the tracklet or hand-crafted features as cues to connect objects across frames. Some approaches [2, 6, 32, 54, 58] used CNNs to extract embeddings as a more robust representation of the objects. The detections and embeddings are produced sequentially, and thus, such systems are referred as the two-stage tracking system. Xiao *et al.* [56] proposed to use the low-level features more efficiently by sharing them between the detection and embedding branches. However, the speed is not close to the real time since the two components are essentially still sequential.

**One-stage tracking system:** A recent approach, named Joint Detection and Embedding (JDE), achieved almost real-time MOT by jointly predicting bounding boxes and their associated embeddings using the same CNN in a single forward pass [51]. A subsequent work, FairMOT [61] that is built on the CenterNet [68], further improved the performance on the benchmark tracking datasets by carefully balancing the two tasks during training. The joint learning scheme was further extended to include the data association step as well and achieved promising results [13, 34, 35, 48, 49]. Particularly, TubeTK [34] jointly modeled the spatial and temporal directions under a 3D CNN framework. Chained-Tracker chains paired objects from adjacent frames to directly integrate the data association step into the end-to-end training [35]. Wang *et al.* proposed to add

a learnable correlation module to the joint model [48] and achieved even better results. All of these trackers are trained on the tracking annotations [7, 9, 24, 33, 56, 63].

In comparison, our proposed training framework does not require expensive fully annotated tracking data. Instead, we combine a detector and an embedder through a data distillation approach [36]. We also borrow ideas from knowledge distillation, which aims to compress a complicated model to a simpler student model [18, 40]. We empirically found that an embedder head consisting of merely two convolutional layers is able to mimic the embedding of a much larger teacher embedder given the backbone features, and thus, it compresses the teacher embedder model.

Not much work focuses on mitigating the scarcity of the fully annotated tracking data on the MOT problem. Fabri *et al.* created a synthetic dataset named MOTSynth that contains ground-truth labels for detection and tracking [12], and empirically showed that the tracker can achieve promising results with training only on this dataset. However, whether these sythetic datasets generalize to real world is still an open question. A concurrent work, KDMOT [60], explores in a similar direction with different purposes. The method still pre-trains a joint model on some fully-annotated tracking data before using a teacher model to supervise its embedding head. In comparison, our goal is to understand how well such a joint model can perform without any limited tracking data. Thus, our model does not see any tracking label during training. In addition, our teacher embedder is only trained on one Re-ID dataset.

**Object detector and Re-ID networks:** We need to choose a suitable detector and appropriate teacher embedder for our proposed framework. In principle, most of the deep object detectors [15, 23, 28, 37, 38, 68] and any embedding-based Re-ID network [1, 25, 45, 46] can be used in our system. In our experiment, we choose the one-stage detector (*i.e.* without Region-of-Interest pooling) RetinaNet [27] as our backbone detector given its simple design. Since people appearing in videos can come in various sizes and different imaging conditions, we prefer to use Re-ID methods that cover multi-scale features and generalize well across domains [5, 29, 50, 65]. Among the many good choices, we choose OSNet [65] as our teacher embedder. Although it is not the state-of-the-art embedder on the Re-ID benchmarks anymore, its simplicity, lightweight design, and well-maintained code base [66] make it a good candidate to demonstrate the effectiveness of our method.

### 3. Method

#### 3.1. Problem formulation

Given a detection dataset and a pre-trained (teacher) Re-ID network, the goal is to train a one-stage tracker that

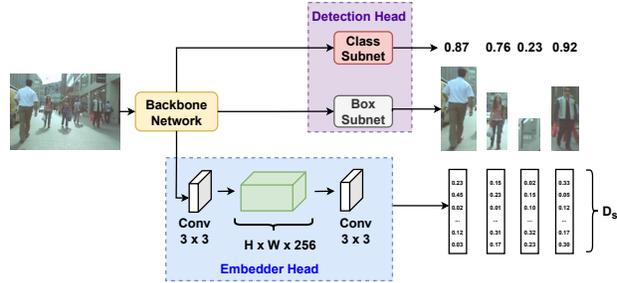


Figure 2. Our proposed joint model, adapted from RetinaNet [27]. The newly added embedder head, consisting of merely two convolutional layers, is parallel to the detection head. Given an input image, the backbone network extract multi-level features, which are fed into the three sub-networks simultaneously. The box subnet predicts the location of the bounding boxes containing the objects of interest. The class sub-net computes a confidence score of every class for each predicted bounding box. The embedder head generates an embedding for each bounding box, representing the discriminative features of that object.

achieves good performance on a target tracking dataset. The detection dataset only provides the ground-truth bounding boxes  $B_i \in \mathbb{R}^{K \times 4}$ , where  $K$  is the number of ground-truth bounding boxes in the  $i^{th}$  image. We do not use any identification labels of any tracking dataset in our network training.

#### 3.2. Model architecture

We choose RetinaNet [27] as our base detector, and add a simple embedder head to the RetinaNet, making it a joint model of detection and embedding. RetinaNet can be combined with different available architectures, *e.g.* MobileNet [19], ResNet [17], and SpineNet [10] to predict detections. Specifically, RetinaNet uses a Feature Pyramid Network (FPN) [26] to generate multi-level features as well as a classification subnet and a box regression one that utilize these multi-level features to make predictions on each of the FPN levels.

As shown in Fig. 2, an embedder head (highlighted in the blue rectangle) is added to RetinaNet, parallel to its detection head. The embedder branch is a cascade of 2 convolutional layers with a batch normalization layer [20] in between. This head predicts an embedding for each anchor box at each level. This means that the output embeddings for level  $l$  have size  $D_s \times S \times R \times H_l \times H_w$ , where  $S$  and  $R$  are the number of scales and aspect ratios in the anchors respectively,  $H_l$  and  $H_w$  are the height and width of level  $l$  in the feature pyramid respectively, and  $D_s$  is the dimension of the output embedding (*i.e.* student embedding).

At inference time, we apply non-maximum suppression (NMS) at threshold 0.5 to reduce the number of anchors. We only take the top 100 or less anchors, and remove those anchors with classification scores smaller than 0.05. The predicted embedding is normalized before being used in the

data association process.

### 3.3. Weak supervision framework

Figure 1 shows the proposed weak supervision framework. We generate a pseudo ground-truth embedding, *i.e.* the teacher embedding, for each ground-truth bounding box provided by the detection dataset. The areas corresponding to each bounding box are cropped independently and resized to the resolution expected by the teacher embedder, say  $256 \times 128$  as in our experiments [65]. These image patches are fed into the teacher embedder, and a teacher embedding  $f \in \mathbb{R}^{D_t}$  of dimension  $D_t$  (*i.e.* dimension of the teacher embedding) is generated and attached to each of these bounding boxes. The student embedding can have the same or less number of dimensions than the teacher embedding, that is  $D_s \leq D_t$ .

We create these pseudo embedding labels for each detection dataset beforehand and store them for the training process. Therefore, this is a one-time computational cost. The memory overhead depends on the density of the bounding boxes in the detection dataset and the number of dimensions of the teacher embeddings. For example, we use a teacher embedding of 512 dimensions [65, 66], and the storage overhead of the detection datasets used in our experiments ranges from 9.3% to 73.3% of the original storage.

Following RetinaNet, each anchor is at most assigned to one ground-truth bounding box based on the intersection-over-union (IoU). If an anchor is assigned to a ground-truth bounding box, it is also assigned the teacher embedding of that bounding box. There is no pseudo ground-truth embedding for those unassigned anchors.

**Teacher embedder** We choose a pre-trained out-of-box Re-ID network as our teacher embedder. The data association step in the tracking process requires the embeddings to be discriminative so that the same object occurrences can be associated across different frames. This is also required in the person re-identification task and thus, the Re-ID network is a good candidate as our teacher embedder. We choose OSNet [65, 66] as the teacher embedder in our proposed teacher-student framework due to its multi-scale design and well-maintained lightweight implementation [64].

### 3.4. Losses

We use different types of losses for these three types of predictions, namely classification, bounding box localization, and embeddings. These losses are measured on all the valid anchor boxes across all the levels defined by the FPN.

Specifically, we use focal loss [27]  $\mathcal{L}_c$  as the supervision signal on the predicted classification score by the classification subnet to mitigate the class imbalance problem. Although there may be many people in a video, the class

imbalance issue still exists due to the large number of anchors. The focal loss reduces the contribution of those easy samples to the final loss so that the network can learn on the hard samples.

Huber loss [16]  $\mathcal{L}_b$  is used on the box regression because of its robustness to outliers. We use the least square loss (L2 loss) as the loss measured between the teacher embeddings  $f$  and the predicted student embeddings  $\hat{f}$ , that is,

$$\mathcal{L}_e(f, \hat{f}) = \sum_{i=1}^{D_s} (f_i - \hat{f}_i)^2, \quad (1)$$

where  $f_i$  is the  $i^{\text{th}}$  element of the embedding. If  $D_s < D_t$ , we use the first  $D_s$  elements in the teacher embedding  $f$  as the supervision signal. Recall that  $D_t$  is the dimension of the teacher embedding. The re-normalization of the new teacher embedding is not required during training. The embedding loss  $\mathcal{L}_e$  of an image is averaged embedding loss over all the assigned anchors.

The total loss is weighted over the three losses,

$$\mathcal{L} = \alpha_c \mathcal{L}_c + \alpha_b \mathcal{L}_b + \alpha_e \mathcal{L}_e. \quad (2)$$

We find that it is robust to a range of different weighting schemes.

### 3.5. Data association

Our joint model can be easily integrated into different existing data association algorithms. Particularly, we use the association system proposed by FairMOT [61]. A tracklet represents the current state of the corresponding track. The first batch of tracklets are initialized based on the detections in the first frame. The Kalman filter [21] is used to predict a likely location of the tracklet in the next frame. Both the motion and the embedding of the objects are used to match the detections in the new frame and the existing tracklets in the tracklet pool, by the Hungarian algorithm [22]. We found that the threshold used in the matching distance can make a difference, since our embedding space is likely different from theirs. The unmatched tracklets and detections are further matched by thresholding the IoU at 0.5.

## 4. Experimental results

### 4.1. Datasets and metrics

Following JDE [51] and FairMOT [61], we integrate several datasets as a large training dataset to avoid training bias. Particularly, we use ETH dataset [11], CityPersons dataset [59], and CrowdHuman dataset [43], which only provide detection labels. In addition, we use fully annotated tracking datasets CalTech dataset [9], MOT17 dataset [33], CUHK-SYSU dataset [56], and PRW dataset [63]. However, we only use their detection labels for training our joint

models and ignore the ID labels. The embedding branch is supervised by the teacher embedders that are not trained on these tracking datasets either.

The teacher embedders used in our experiments are only trained on the MSMT17 dataset [52], which consists of 180 hours of videos and identity labels for person re-identification task. We use the teacher embedders to augment the mentioned detection datasets with pseudo embedding labels, following the procedure described in Sec. 3.3.

We use *BaseTrainSet* to denote the collection of all these augmented detection datasets excluding CrowdHuman. *BaseTrainSet* is used to train various models in our ablation studies (if not explicitly mentioned) and in Tab. 1. We use six video sequences from the MOT15 [24] training set as the validation dataset in the ablation studies and in Tab. 1, which do not include the overlapped sequences with the MOT17 dataset. We also evaluate our model, trained on both the *BaseTrainSet* and augmented CrowdHuman datasets, on the private MOT16 and MOT17 benchmark [24, 33] to compare with other work (Tab. 2). Our model is further fine-tuned on the MOT20 [7] training set before evaluated on the private MOT20 benchmark (Tab. 2).

We evaluate both the detection and tracking abilities of our models. We use Average Precision (AP) at IoU 0.5 to measure the performance of detection [51, 61]. Tracking performance is measured by the popular CLEAR metrics [3] and IDF1 score [39].

## 4.2. Implementation details

We use most of the default settings of RetinaNet [27] in our base detector. Particularly, the FPN consists of features from level 3 to 7. We use three scales and three aspect ratios (0.25, 0.5, and 1.0) of anchors. Both the class subnet and box subnet consist of four convolutional layers of feature size 128. The embedder head consists of two convolutional layers of filter size  $3 \times 3$ . The intermediate feature size is 256. We experiment on the output embedding size of 64, 128, 256, and 512 in an ablation study (Tab. 6), and use 128, combined with a ResNet-34 backbone architecture, in the benchmark evaluation (Tab. 2). The teacher embedder [65, 66] generates embeddings of size 512.

The input image size is  $608 \times 1088$  for ablation studies and Tab. 1, and  $1080 \times 1920$  for private benchmark evaluation (Tab. 2). Since our teacher embedder is not trained on the similar tracking datasets, we found that clearer human figures resulting from the larger resolution lead to better tracking performance. We use standard data augmentation techniques including random horizontal flipping and random scaling with scaling ratio from 0.5 to 2.0.

For the weighted loss, we set  $\alpha_c = 1$  for the focal loss on classification,  $\alpha_b = 50$  for the Huber loss on the bounding box regression, and  $\alpha_e = 10$  for the L2 loss on embedding. We found that both the detection and tracking performance

are robust to the weight of the embedding loss in a range of 2 to 10. We set  $\alpha$  as 0.25 and  $\gamma$  as 1.5 in the focal loss; the  $\delta$  in the Huber loss is 0.1.

We set initial learning rate as 0.15 and use the cosine decay schedule [30]. The warm up learning rate is 0.001 for the first 2000 iterations. We use batch size of 16. The models are trained with 300K to 500K iterations depending on the settings. All inferences are computed on a single Nvidia P100 GPU.

## 4.3. Main results

### 4.3.1 TDT-tracker versus two-stage tracker

Table 1 compares our TDT-tracker and its counterpart two-stage tracker. The detectors of both trackers have the same architecture and are trained under the same training settings. From Tab. 1, we see that the joint learning has a negative effect on the detection performance with 1.2% (0.924 to 0.913) decrease on the Average Precision (AP) of detection. However, the one-stage tracker achieves much better tracking performance with much faster speed. For example, the one-stage tracker achieves 3.3% better MOTA (0.760 to 0.785) and 37.4% (91 to 57) less ID switches with 3.2 times faster speed (3.03 to 9.61 frame per second).

Tracker	MOTA $\uparrow$	IDF1 $\uparrow$	IDs $\downarrow$	AP $\uparrow$	FPS $\uparrow$
Two-stage	0.760	0.769	91	<b>0.924</b>	3.03
TDT-tracker	<b>0.785</b>	<b>0.787</b>	<b>57</b>	0.913	<b>9.61</b>

Table 1. Comparison between the two-stage model and our joint model (TDT-tracker). Both models use the same detector architecture, and the joint model has an additional embedder head (Fig. 2). We use an OSNet [65, 66] as the embedder in the two-stage tracker as well as the teacher embedder for training our TDT-tracker.

### 4.3.2 MOT challenges

We evaluate our TDT-Tracker on the private benchmark tracking datasets [7, 33] to compare with some existing trackers (Tab. 2). There are two categories of trackers, namely the ones that use the fully annotated real tracking data in the training and the ones that do not. The state-of-the-art results are obtained by the former ones [49, 61].

From Tab. 2, our TDT-tracker outperforms other methods under the same category by a large margin on MOT17 [33] and MOT20 [7]. For example, TDT-tracker outperforms the second MOTA performance by 6.9% (59.7 to 63.8) and the second IDF1 by 17.1% (52.0 to 60.9) on MOT17. On MOT16 dataset, there are two two-stage trackers, namely CNNMTT [32] and POI [58], that achieve better MOTA and IDF1 than TDT-tracker (we still outperform on Mostly Tracked Targets (MT) and Mostly Lost Targets

Dataset	Tracker	MOTA $\uparrow$	IDF1 $\uparrow$	MT $\uparrow$	ML $\downarrow$	IDs $\downarrow$	Use Tracking Labels?	
MOT16	JDE* [51]	64.4	55.8	35.4%	20.0%	1544	Yes	
	ChainedTrackerV1* [35]	67.6	57.2	32.9%	23.1%	1897		
	FairMOT* [61]	<b>74.9</b>	<b>72.8</b>	<b>44.7%</b>	<b>15.9%</b>	<b>1074</b>		
	MOT16	SORTwHPD16 [4, 61]	59.8	53.8	25.4%	22.7%	1423	No
		DeepSORT_2 [54]	61.4	62.2	32.8%	18.2%	<b>781</b>	
		CNNMTT [32]	65.2	62.2	32.4%	21.3%	946	
		POI [58]	<b>66.1</b>	<b>65.1</b>	34.0%	20.8%	805	
TDT-tracker* (Ours)		64.1	61.5	<b>36.5%</b>	<b>16.5%</b>	1391		
MOT17	CenterTrack* [67]	67.8	64.7	34.6%	24.6%	<b>2583</b>	Yes	
	FairMOT* [61]	73.7	72.3	43.2%	17.3%	3303		
	CorrTracker* [48]	<b>76.5</b>	<b>73.6</b>	<b>47.6%</b>	<b>12.7%</b>	3369		
	MOT17	SST [44]	52.4	49.5	21.4%	30.7%	8431	No
		CenterTrack-MOTSynth* [12]	59.7	52.0	-	-	6035	
TDT-tracker* (Ours)		<b>63.8</b>	<b>60.9</b>	<b>35.4%</b>	<b>17.3%</b>	<b>4401</b>		
MOT20	GSDT [49]	<b>67.1</b>	67.5	53.1%	13.2%	<b>3133</b>	Yes	
	CorrTracker* [48]	65.2	<b>69.1</b>	<b>66.4%</b>	<b>8.9%</b>	5183		
	MOT20	Tracktor-MOTSynth [12]	43.7	39.7	-	-	<b>3467</b>	No
		TDT-tracker* (Ours)	<b>47.9</b>	<b>46.0</b>	18.5%	20.1%	5342	

Table 2. Comparison with existing trackers under the private MOT benchmark datasets [7, 33]. These trackers are further categorized into the ones that use (*i.e.* Yes in the last column) the expensive fully annotated real tracking data and the ones that do not (*i.e.* No). The one-stage trackers are marked by “\*\*”.

(ML) metrics). CNNMTT uses the publicly available detections by a state-of-the-art detector and use the ID labels of MOT16 training set to train its embedder. Those ID labels are very useful since the MOT16 training and testing sets contain similar types of people appearing. POI [58] uses different detection thresholds on different test sequences, while we only use one threshold. Adaptive thresholding helps achieve better performance on the benchmark datasets but it is not very practical.

In addition, while maintaining very similar MOTA as JDE [51] on MOT16 benchmark, our TDT-tracker achieves 10.2% better IDF1. JDE is a one-stage tracker and trained on some fully annotated tracking datasets.

## 4.4. Ablation studies

### 4.4.1 Importance of teacher embedders

Teacher Embedder	MOTA $\uparrow$	IDF1 $\uparrow$	IDs $\downarrow$	AP $\uparrow$
ResNet50 (22.8)	0.754	0.764	76	0.902
ON-Small (31.0)	0.763	0.766	76	0.911
ON-Large ( <b>43.3</b> )	<b>0.785</b>	<b>0.787</b>	<b>57</b>	<b>0.913</b>

Table 3. Detection and tracking performance of our TDT-tracker with different teacher embedders. All these models use the same architecture and training settings. The number in the parentheses adjacent to each teacher embedder is its mean Average Precision (mAP) evaluated on the Market1501 dataset [62].

The teacher embedder determines the embedding quality of our joint model as it provides the pseudo labels for training. In Tab. 3, we compare our TDT-tracker under three different teacher embedders, namely ResNet-50 [17, 64], OSNet-Small [65, 66], OSNet-Large [65, 66]. These three embedders are trained on the MSMT17 Re-ID dataset [52] (Fig. 3 left block) and achieve increasing mean Average Precision (mAP) on the Market1501 Re-ID dataset [62].

Interestingly, the one-stage tracker achieves increasing detection performance with better teacher embedders, and the largest AP increase is 1.2% (0.902 to 0.913). The tracking performance also increases with better teacher embedders. For example, comparing to using ResNet50 as the teacher embedder, TDT-tracker achieves 4.1% better MOTA and 25.0% lower ID switches.

Dataset	FairMOT		TDT-RNet34	
	AP $\uparrow$	MOTA $\uparrow$	AP $\uparrow$	MOTA $\uparrow$
Overall	<b>0.932</b>	<b>80.5</b>	0.926	78.4
KITTI13	<b>0.796</b>	<b>42.7</b>	0.792	32.3
ETH-SunnyDay	<b>0.997</b>	84.3	0.984	<b>88.7</b>

Table 4. Comparison between FairMOT [61] and TDT-tracker on the detection and tracking performance separately. Both models are trained on the same datasets including the BaseTrainSet and CrowdHuman. KITTI13 and ETH-Sunnyday are two sequences in the validation set.

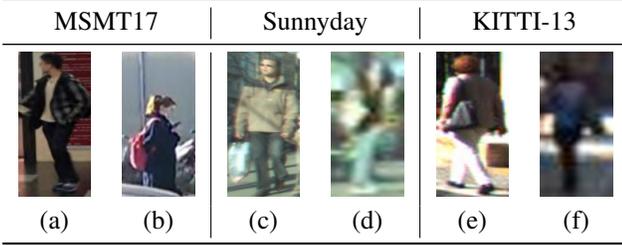


Figure 3. Left: Two examples from MSMT17 Re-ID dataset [52] that our teacher embedder is trained on. Middle and right: Examples from the Sunnyday sequence and KITTI-13 sequence of the MOT15 [24] training set that is used as the validation set in ablation studies. (a) is a typical sample in the MSMT17 dataset, and (b) is an occluded sample. (c) and (e) are two clear examples from Sunnyday and KITTI-13 respectively, and (d) and (f) are unclear ones. The clear examples are closer to the ones in the MSMT17, and they are more common in the Sunnyday sequence but not in the KITTI-13 sequence.

However, the current teacher embedder cannot support our joint model to achieve comparable tracking performance as the state-of-the-art tracker [48, 49, 61] (trained on the fully annotated tracking labels) on some cases. In Tab. 4, we compare the performance of a FairMOT [61] model and our TDT model on the validation dataset. Both models have been trained on the same datasets, namely the BaseTrain-Set and the CrowdHuman dataset, and use the ResNet-34 as the backbone network. In a sequence named KITTI13 (second row in Tab. 4) from the validation set, the most of the pedestrians appearing in the video are not clear (see Fig. 3 (f)), and our embedder does much worse (24.36% worse on MOTA) than FairMOT. In contrast, our embedder outperform FairMOT by 5.22% on MOTA on the sequence ETH-Sunnyday (third row) where most people are large and clear in the video (see Fig. 3 (c)). Both detectors perform comparably on both sequences, and thus, the embedding quality is the main cause. FairMOT does better on the KITTI13 likely attributes to their privilege of seeing similar samples during training. However, this privilege may not always exist.

Figure 4 shows qualitative results from our TDT-tracker on different types of video sequences in MOT17 [33] and MOT20 [7] test sets. Most of the objects are accurately detected and our TDT-tracker can track well even on occluded people (see the successful tracking of the fifth person from left with ID 461 in the starting frame of sequence (a) in Fig. 4).

#### 4.4.2 Backbone choices for the joint model

A suitable backbone architecture is essential to the success of our weakly supervised scheme. We need the backbone features to contain not only the information for detecting objects but also information on the distinguishing properties of the objects. In addition, these properties can be extracted

Model	MOTA↑	IDF1↑	IDs↓	AP↑	FPS↑
MNet	0.744	0.747	72	0.907	<b>11.18</b>
MNet-FPN	0.764	0.759	73	0.910	10.01
RNet34-FPN	<b>0.785</b>	<b>0.787</b>	<b>57</b>	<b>0.913</b>	9.61
RNet50-FPN	0.757	0.786	58	0.912	7.93

Table 5. Impacts of different backbone architecture on the performance of our TDT-tracker. MNet refers to MobileNet [19] and RNet is ResNet [17].

for people of different sizes in an image. In Tab. 5, we show the performance of detection and tracking under different choices of backbone architecture.

From Tab. 5, we can see that the feature fusion mechanism helps the overall tracking performance. With the FPN, MOTA improves 2.7% (0.744 to 0.764) on the MobileNet backbone, *i.e.* from MNet to MNet-FPN. A more powerful backbone architecture boosts the tracking performance by another 2.75% (0.764 to 0.785), *i.e.* from MNet-FPN to RNet34-FPN. The improved embedding quality is the main reason since the detection results are almost the same but the number of ID switches drops 21.9% (from 73 to 57). Interestingly, using ResNet-50 decreases the overall tracking performance measured by MOTA by a large margin. The reason is that the tracking precision (as in the CLEAR metrics [3]) decreases while other metrics stay similar. Similar phenomenon is also observed in FairMOT [61], where they also find that simply increasing the backbone power may not improve the overall tracking performance.

#### 4.4.3 Impact of embedding dimensionality

Dim	MOTA ↑	IDF1 ↑	IDs ↓	AP ↑	FPS ↑
64	0.739	0.772	67	0.902	<b>13.34</b>
128	0.773	0.782	<b>55</b>	0.908	11.82
256	0.762	0.777	63	0.910	10.76
512	<b>0.785</b>	<b>0.787</b>	57	<b>0.913</b>	9.61

Table 6. The effects of using different portion of the teacher embedding as the supervision signal. All the models share the same architecture except for the last layer of the embedding head since embeddings of different dimensions are generated.

A practical issue of using the proposed weakly supervised framework is that the dimension of the teacher embedder can be too large, which can create a large memory and computational overhead, especially for the anchor based detection models as they predict an embedding for each anchor box.

As mentioned in Sec. 3.4, a simple way to circumvent the situation is to use the first  $D_s$  elements of the teacher embedding. Recall that  $D_s$  represents the size of the desired

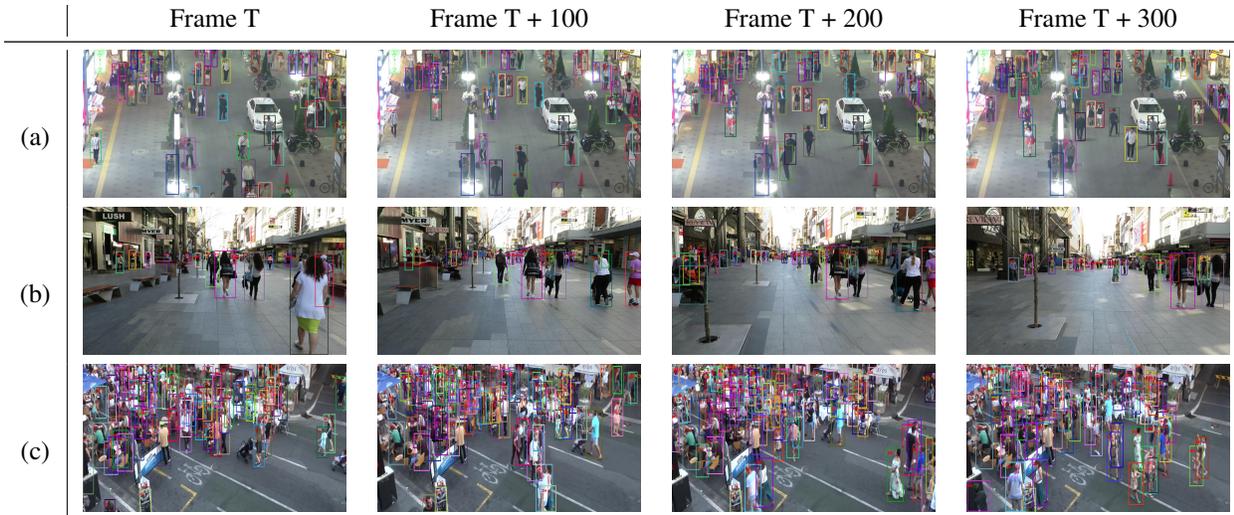


Figure 4. Three qualitative examples by our TDT-tracker on video sequences [7, 33] of different types. Sequence (a) and (b) contain high-resolution people with less occlusions. Sequence (c) shows the crowded scenario.

output embedding, *i.e.* student embedding. Tab. 6 shows the performance of setting  $D_s$  to be 64, 128, 256, and 512 respectively, given that  $D_t$  (dimension of the teacher embedding) is 512. First, there is a clear pattern of improving detection performance as larger portion of the teacher embedding is used. Interestingly, in comparison to the full teacher embeddings, our joint model can achieve 94.1% (0.739 out of 0.785) of the tracking performance measured in MOTA and 98.1% if measured in IDF1 by using only one eighth of them. However, increasing the embedding portion generally leads to better tracking performance and using the full teacher embedding still bests other settings. More dimension contains more information that can be helpful in discriminating people. In addition, the teacher embedding may not be evenly distributed, and the first part likely cannot represent the full information. Note that the  $D_s = 128$  and  $D_s = 256$  are against the trend, where the tracking performance decreases as  $D_s$  increases. We believe this is due to the randomness of data. On the other hand, there is a clear increasing pattern of the inference speed as  $D_s$  decreases. Thus, this is a typical speed and accuracy trade off situation.

#### 4.4.4 Training data

Under the proposed weakly supervised framework, any detection dataset can be easily augmented with some teacher embeddings and used as a tracking dataset to train a joint model for tracking. This is an advantage. We show in Fig. 5 that the overall tracking performance, measured in both MOTA and IDF1, improves as more augmented datasets are used in training. The detector also improves, although marginally, as well. This means our weakly super-

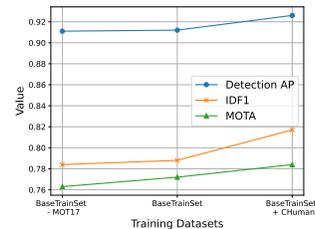


Figure 5. Improvement of the performance in both the detection (AP) and tracking (MOTA and IDF1) of our proposed system as more distilled detection datasets are used for training. These three models share the same architecture and training settings. The higher the better for these metrics.

vised framework benefits from the diversity of the detection datasets, which is promising because detection datasets are easier to annotate than tracking datasets.

## 5. Conclusion

In this paper, we propose a simple but effective embedding distillation framework that aims to mitigate the issue of expensive and scarce fully annotated real tracking data. Our TDT-tracker achieves competitive results on the benchmark datasets and we analyze various aspects of our framework. Particularly, we have shown that our tracker achieves better performance with better embedders. This is a promising direction since better Re-ID networks are actively being proposed, which directly improves our one-stage tracker. While our work has avoided using any tracking data to train the teacher embedder, it is still possible and practical to sample some pairs of frames from videos, have them annotated, and use their annotations as Re-ID examples to further improve the student tracker performance.

## References

- [1] Ejaz Ahmed, Michael Jones, and Tim K Marks. An improved deep learning architecture for person re-identification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3908–3916, 2015. 3
- [2] Seung-Hwan Bae and Kuk-Jin Yoon. Confidence-based data association and discriminative deep appearance learning for robust online multi-object tracking. *IEEE transactions on pattern analysis and machine intelligence*, 40(3):595–610, 2017. 1, 2
- [3] Keni Bernardin and Rainer Stiefelwagen. Evaluating multiple object tracking performance: the clear mot metrics. *EURASIP Journal on Image and Video Processing*, 2008:1–10, 2008. 5, 7
- [4] Alex Bewley, Zongyuan Ge, Lionel Ott, Fabio Ramos, and Ben Upcroft. Simple online and realtime tracking. In *2016 IEEE international conference on image processing (ICIP)*, pages 3464–3468. IEEE, 2016. 6
- [5] Xiaobin Chang, Timothy M Hospedales, and Tao Xiang. Multi-level factorisation net for person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2109–2118, 2018. 3
- [6] Long Chen, Haizhou Ai, Chong Shang, Zijie Zhuang, and Bo Bai. Online multi-object tracking with convolutional neural networks. In *2017 IEEE International Conference on Image Processing (ICIP)*, pages 645–649. IEEE, 2017. 1, 2
- [7] P. Dendorfer, H. Rezatofighi, A. Milan, J. Shi, D. Cremers, I. Reid, S. Roth, K. Schindler, and L. Leal-Taixé. Mot20: A benchmark for multi object tracking in crowded scenes. *arXiv:2003.09003[cs]*, Mar. 2020. arXiv: 2003.09003. 3, 5, 6, 7, 8
- [8] Piotr Dollár, Ron Appel, Serge Belongie, and Pietro Perona. Fast feature pyramids for object detection. *IEEE transactions on pattern analysis and machine intelligence*, 36(8):1532–1545, 2014. 2
- [9] Piotr Dollár, Christian Wojek, Bernt Schiele, and Pietro Perona. Pedestrian detection: A benchmark. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 304–311. IEEE, 2009. 3, 4
- [10] Xianzhi Du, Tsung-Yi Lin, Pengchong Jin, Golnaz Ghiasi, Mingxing Tan, Yin Cui, Quoc V. Le, and Xiaodan Song. Spinenet: Learning scale-permuted backbone for recognition and localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 3
- [11] Andreas Ess, Bastian Leibe, Konrad Schindler, and Luc Van Gool. A mobile vision system for robust multi-person tracking. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2008. 4
- [12] Matteo Fabbri, Guillem Brasó, Gianluca Maueri, Aljoša Ošep, Riccardo Gasparini, Orcun Cetintas, Simone Calderara, Laura Leal-Taixé, and Rita Cucchiara. Motsynth: How can synthetic data help pedestrian detection and tracking? In *International Conference on Computer Vision (ICCV)*, 2021. 3, 6
- [13] Kuan Fang, Yu Xiang, Xiaocheng Li, and Silvio Savarese. Recurrent autoregressive networks for online multi-object tracking. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 466–475. IEEE, 2018. 2
- [14] Pedro F Felzenszwalb, Ross B Girshick, David McAllester, and Deva Ramanan. Object detection with discriminatively trained part-based models. *IEEE transactions on pattern analysis and machine intelligence*, 32(9):1627–1645, 2009. 2
- [15] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015. 3
- [16] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning*. Springer Series in Statistics. Springer New York Inc., New York, NY, USA, 2001. 4
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 3, 6, 7
- [18] Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. *ArXiv*, abs/1503.02531, 2015. 2, 3
- [19] Andrew G. Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *CoRR*, abs/1704.04861, 2017. 3, 7
- [20] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. PMLR, 2015. 3
- [21] Rudolph Emil Kalman. A new approach to linear filtering and prediction problems. 1960. 4
- [22] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955. 4
- [23] Hei Law and Jia Deng. Cornernet: Detecting objects as paired keypoints. In *Proceedings of the European conference on computer vision (ECCV)*, pages 734–750, 2018. 3
- [24] L. Leal-Taixé, A. Milan, I. Reid, S. Roth, and K. Schindler. MOTChallenge 2015: Towards a benchmark for multi-target tracking. *arXiv:1504.01942 [cs]*, Apr. 2015. arXiv: 1504.01942. 3, 5, 7
- [25] Wei Li, Rui Zhao, Tong Xiao, and Xiaogang Wang. Deep-reid: Deep filter pairing neural network for person re-identification. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 152–159, 2014. 3
- [26] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017. 3
- [27] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision (CVPR)*, pages 2980–2988, 2017. 3, 4, 5
- [28] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C

- Berg. Ssd: Single shot multibox detector. In *European conference on computer vision*, pages 21–37. Springer, 2016. 3
- [29] Xihui Liu, Haiyu Zhao, Maoqing Tian, Lu Sheng, Jing Shao, Shuai Yi, Junjie Yan, and Xiaogang Wang. Hydraplus-net: Attentive deep features for pedestrian analysis. In *Proceedings of the IEEE international conference on computer vision*, pages 350–359, 2017. 3
- [30] Ilya Loshchilov and Frank Hutter. SGDR: stochastic gradient descent with warm restarts. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. 5
- [31] Zhichao Lu, Vivek Rathod, Ronny Votel, and Jonathan Huang. Retinatrack: Online single stage joint detection and tracking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14668–14678, 2020. 1
- [32] Nima Mahmoudi, Seyed Mohammad Ahadi, and Mohammad Rahmati. Multi-target tracking using cnn-based features: Cnnmtt. *Multimedia Tools and Applications*, 78:7077–7096, 2018. 2, 5, 6
- [33] A. Milan, L. Leal-Taixé, I. Reid, S. Roth, and K. Schindler. MOT16: A benchmark for multi-object tracking. *arXiv:1603.00831 [cs]*, Mar. 2016. arXiv: 1603.00831. 3, 4, 5, 6, 7, 8
- [34] Bo Pang, Yizhuo Li, Yifan Zhang, Muchen Li, and Cewu Lu. Tubetk: Adopting tubes to track multi-object in a one-step training model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6308–6318, 2020. 2
- [35] Jinlong Peng, Changan Wang, Fangbin Wan, Yang Wu, Yabiao Wang, Ying Tai, Chengjie Wang, Jilin Li, Feiyue Huang, and Yanwei Fu. Chained-tracker: Chaining paired attentive regression results for end-to-end joint multiple-object detection and tracking. In *European Conference on Computer Vision*, pages 145–161. Springer, 2020. 2, 6
- [36] Ilija Radosavovic, Piotr Dollár, Ross Girshick, Georgia Gkioxari, and Kaiming He. Data distillation: Towards omniscient supervised learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4119–4128, 2018. 3
- [37] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016. 3
- [38] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28:91–99, 2015. 3
- [39] 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. 5
- [40] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. In *In Proceedings of International Conference on Learning Representations (ICLR)*, 2015. 2, 3
- [41] Ricardo Sánchez-Matilla, Fabio Poiesi, and Andrea Cavallaro. Online multi-target tracking with strong and weak detections. In *ECCV Workshops*, 2016. 1, 2
- [42] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, pages 815–823, 2015. 2
- [43] Shuai Shao, Zijian Zhao, Boxun Li, Tete Xiao, Gang Yu, Xiangyu Zhang, and Jian Sun. Crowdhuman: A benchmark for detecting human in a crowd. *arXiv preprint arXiv:1805.00123*, 2018. 4
- [44] ShiJie Sun, Naveed Akhtar, HuanSheng Song, Ajmal Mian, and Mubarak Shah. Deep affinity network for multiple object tracking. *IEEE transactions on pattern analysis and machine intelligence*, 43(1):104–119, 2019. 6
- [45] Yifan Sun, Liang Zheng, Yi Yang, Qi Tian, and Shengjin Wang. Beyond part models: Person retrieval with refined part pooling (and a strong convolutional baseline). In *Proceedings of the European conference on computer vision (ECCV)*, pages 480–496, 2018. 3
- [46] Rahul Rama Varior, Mrinal Haloi, and Gang Wang. Gated siamese convolutional neural network architecture for human re-identification. In *European conference on computer vision*, pages 791–808. Springer, 2016. 3
- [47] Paul Voigtlaender, Michael Krause, Aljosa Osep, Jonathon Luiten, Berin Balachandar Gnana Sekar, Andreas Geiger, and Bastian Leibe. Mots: Multi-object tracking and segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7942–7951, 2019. 2
- [48] Qiang Wang, Yun Zheng, Pan Pan, and Yinghui Xu. Multiple object tracking with correlation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3876–3886, 2021. 2, 3, 6, 7
- [49] Yongxin Wang, Kris Kitani, and Xinshuo Weng. Joint object detection and multi-object tracking with graph neural networks. In *Proceedings of (ICRA) International Conference on Robotics and Automation*, May 2021. 2, 5, 6, 7
- [50] Yan Wang, Lequn Wang, Yurong You, Xu Zou, Vincent Chen, Serena Li, Gao Huang, Bharath Hariharan, and Kilian Q Weinberger. Resource aware person re-identification across multiple resolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8042–8051, 2018. 3
- [51] Zhongdao Wang, Liang Zheng, Yixuan Liu, Yali Li, and Shengjin Wang. Towards real-time multi-object tracking. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *European Conference on Computer Vision (ECCV)*, pages 107–122, Cham, 2020. Springer International Publishing. 1, 2, 4, 5, 6
- [52] Longhui Wei, Shiliang Zhang, Wen Gao, and Qi Tian. Person transfer gan to bridge domain gap for person re-identification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 79–88, 2018. 5, 6, 7

- [53] Kilian Q Weinberger and Lawrence K Saul. Distance metric learning for large margin nearest neighbor classification. *Journal of machine learning research*, 10(2), 2009. 2
- [54] 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, 6
- [55] Yu Xiang, Alexandre Alahi, and Silvio Savarese. Learning to track: Online multi-object tracking by decision making. In *Proceedings of the IEEE international conference on computer vision*, pages 4705–4713, 2015. 1, 2
- [56] Tong Xiao, Shuang Li, Bochao Wang, Liang Lin, and Xiaogang Wang. Joint detection and identification feature learning for person search. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3415–3424, 2017. 2, 3, 4
- [57] Fan Yang, Wongun Choi, and Yuanqing Lin. Exploit all the layers: Fast and accurate cnn object detector with scale dependent pooling and cascaded rejection classifiers. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2129–2137, 2016. 2
- [58] Fengwei Yu, Wenbo Li, Quanquan Li, Yu Liu, Xiaohua Shi, and Junjie Yan. Poi: Multiple object tracking with high performance detection and appearance feature. In *European Conference on Computer Vision*, pages 36–42. Springer, 2016. 2, 5, 6
- [59] Shanshan Zhang, Rodrigo Benenson, and Bernt Schiele. Citypersons: A diverse dataset for pedestrian detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3213–3221, 2017. 4
- [60] Wei Zhang, Lingxiao He, Peng Chen, Xingyu Liao, Wu Liu, Qi Li, and Zhenan Sun. Boosting end-to-end multi-object tracking and person search via knowledge distillation. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 1192–1201, 2021. 3
- [61] Yifu Zhang, Chunyu Wang, Xinggang Wang, Wenjun Zeng, and Wenyu Liu. Fairmot: On the fairness of detection and re-identification in multiple object tracking. *International Journal of Computer Vision*, pages 1–19, 2021. 1, 2, 4, 5, 6, 7
- [62] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. Scalable person re-identification: A benchmark. In *Proceedings of the IEEE international conference on computer vision*, pages 1116–1124, 2015. 6
- [63] Liang Zheng, Hengheng Zhang, Shaoyan Sun, Manmohan Chandraker, Yang Yang, and Qi Tian. Person re-identification in the wild. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3346–3355, 2017. 3, 4
- [64] Kaiyang Zhou and Tao Xiang. Torchreid: A library for deep learning person re-identification in pytorch. *arXiv preprint arXiv:1910.10093*, 2019. 4, 6
- [65] Kaiyang Zhou, Yongxin Yang, Andrea Cavallaro, and Tao Xiang. Omni-scale feature learning for person re-identification. In *ICCV*, 2019. 2, 3, 4, 5, 6
- [66] Kaiyang Zhou, Yongxin Yang, Andrea Cavallaro, and Tao Xiang. Learning generalisable omni-scale representations for person re-identification. *TPAMI*, 2021. 2, 3, 4, 5, 6
- [67] Xingyi Zhou, Vladlen Koltun, and Philipp Krähenbühl. Tracking objects as points. In *European Conference on Computer Vision*, pages 474–490. Springer, 2020. 6
- [68] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. In *arXiv preprint arXiv:1904.07850*, 2019. 1, 2, 3