

ByteTrack: Multi-Object Tracking by Associating Every Detection Box

Yifu Zhang¹, Peize Sun², Yi Jiang³, Dongdong Yu³, Fucheng Weng¹,
Zehuan Yuan³, Ping Luo², Wenyu Liu¹, and Xinggang Wang^{1†}

¹ Huazhong University of Science and Technology

² The University of Hong Kong

³ ByteDance Inc.

Abstract. Multi-object tracking (MOT) aims at estimating bounding boxes and identities of objects in videos. Most methods obtain identities by associating detection boxes whose scores are higher than a threshold. The objects with low detection scores, *e.g.* occluded objects, are simply thrown away, which brings non-negligible true object missing and fragmented trajectories. To solve this problem, we present a simple, effective and generic association method, tracking by associating almost every detection box instead of only the high score ones. For the low score detection boxes, we utilize their similarities with tracklets to recover true objects and filter out the background detections. When applied to 9 different state-of-the-art trackers, our method achieves consistent improvement on IDF1 score ranging from 1 to 10 points. To put forwards the state-of-the-art performance of MOT, we design a simple and strong tracker, named ByteTrack. For the first time, we achieve 80.3 MOTA, 77.3 IDF1 and 63.1 HOTA on the test set of MOT17 with 30 FPS running speed on a single V100 GPU. ByteTrack also achieves state-of-the-art performance on MOT20, HiEve and BDD100K tracking benchmarks. The source code, pre-trained models with deploy versions and tutorials of applying to other trackers are released at <https://github.com/ifzhang/ByteTrack>.

Keywords: multi-object tracking, data association, detection boxes

1 Introduction

Was vernünftig ist, das ist wirklich; und was wirklich ist, das ist vernünftig.

— G. W. F. Hegel

Tracking-by-detection is the most effective paradigm for multi-object tracking (MOT) in current. Due to the complex scenarios in videos, detectors are prone to make imperfect predictions. State-of-the-art MOT methods [2,18,46,1,71,6,69,13,3,84,58] need to deal with true positive / false positive trade-off in detection boxes to eliminate low confidence detection boxes [4,41]. However, is it the right way to eliminate all low confidence detection boxes? Our answer is NO: as Hegel said “What is reasonable is real; that which is real is reasonable.” Low confidence detection boxes sometimes indicate the existence of objects, *e.g.* the occluded objects. Filtering out these objects

[†] Corresponding author.

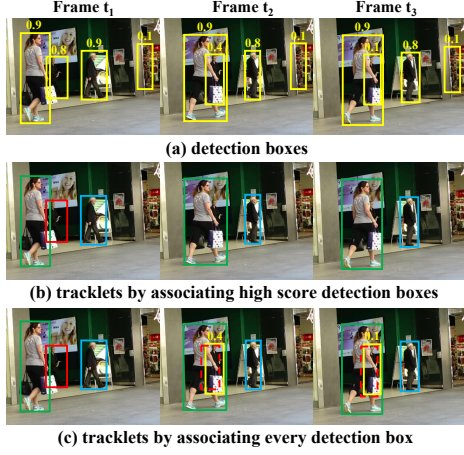


Fig. 1. Examples of our method which associates almost every detection box. (a) shows all the detection boxes with their scores. (b) shows the tracklets obtained by previous methods which associates detection boxes whose scores are higher than a threshold, *i.e.* 0.5. The same box color represents the same identity. (c) shows the tracklets obtained by our method. The dashed boxes represent the predicted box of the previous tracklets using Kalman filter. The two low score detection boxes are correctly matched to the previous tracklets based on the large IoU.

causes irreversible errors for MOT and brings non-negligible missing detection and fragmented trajectories.

Figure 1 (a) and (b) show this problem. In frame t_1 , we initialize three different tracklets as their scores are all higher than 0.5. However, in frame t_2 and frame t_3 when occlusion happens, red tracklet’s corresponding detection score becomes lower *i.e.* 0.8 to 0.4 and then 0.4 to 0.1. These detection boxes are eliminated by the thresholding mechanism and the red tracklet disappears accordingly. Nevertheless, if we take every detection box into consideration, more false positives will be introduced immediately, *e.g.*, the most right box in frame t_3 of Figure 1 (a). To the best of our knowledge, very few methods [31,61] in MOT are able to handle this detection dilemma.

In this paper, we identify that the similarity with tracklets provides a strong cue to distinguish the objects and background in low score detection boxes. As shown in Figure 1 (c), two low score detection boxes are matched to the tracklets by the motion model’s predicted boxes, and thus the objects are correctly recovered. At the same time, the background box is removed since it has no matched tracklet.

For making full use of detection boxes from high scores to low ones in the matching process, we present a simple and effective association method BYTE, named for each detection box is a basic unit of the tracklet, as byte in computer program, and our tracking method values every detailed detection box. We first match the high score detection boxes to the tracklets based on motion similarity or appearance similarity. Similar to [6], we adopt Kalman filter [30] to predict the location of the tracklets in the new frame. The similarity can be computed by the IoU or Re-ID feature distance of the predicted box and the detection box. Figure 1 (b) is exactly the results after the first matching. Then, we perform the second matching between the unmatched tracklets, *i.e.* the tracklet in red box, and the low score detection boxes using the same motion similarity. Figure 1 (c) shows the results after the second matching. The occluded person with low detection scores is matched correctly to the previous tracklet and the background (in the right part of the image) is removed.

As the integrating topic of object detection and association, a desirable solution to MOT is never a detector and the following association; besides, well-designed of their junction area is also important. The innovation of BYTE lies in the junction area of




Fig. 2. MOTA-IDF1-FPS comparisons of different trackers on the test set of MOT17. The horizontal axis is FPS (running speed), the vertical axis is MOTA, and the radius of circle is IDF1. Our ByteTrack achieves 80.3 MOTA, 77.3 IDF1 on MOT17 test set with 30 FPS running speed, outperforming all previous trackers. Details are given in Table 4.

detection and association, where low score detection boxes are bridges to boost both of them. Benefiting from this integration innovation, when BYTE is applied to 9 different state-of-the-art trackers, including the Re-ID-based ones [68,84,34,48], motion-based ones [87,70], chain-based one [49] and attention-based ones [58,79], notable improvements are achieved on almost all the metrics including MOTA, IDF1 score and ID switches. For example, we increase the MOTA of CenterTrack [87] from 66.1 to 67.4, IDF1 from 64.2 to 74.0 and decrease the IDs from 528 to 144 on the half validation set of MOT17.

Towards pushing forwards the state-of-the-art performance of MOT, we propose a simple and strong tracker, named ByteTrack. We adopt a recent high-performance detector YOLOX [25] to obtain the detection boxes and associate them with our proposed BYTE. On the MOT challenges, ByteTrack ranks 1st on both MOT17 [45] and MOT20 [17], achieving 80.3 MOTA, 77.3 IDF1 and 63.1 HOTA with 30 FPS running speed on V100 GPU on MOT17 and 77.8 MOTA, 75.2 IDF1 and 61.3 HOTA on much more crowded MOT20. ByteTrack also achieves state-of-the-art performance on HiEve [38] and BDD100K [78] tracking benchmarks. We hope the efficiency and simplicity of ByteTrack could make it attractive in real applications such as social computing.

2 Related Work

2.1 Object Detection in MOT

Object detection is one of the most active topics in computer vision and it is the basis of multi-object tracking. The MOT17 dataset [45] provides detection results obtained by popular detectors such as DPM [23], Faster R-CNN [51] and SDP [76]. A large number of methods [72,15,3,13,90,9,29] focus on improving the tracking performance based on these given detection results.

Tracking by detection. With the rapid development of object detection [51,27,50,36,11,24,59,57,21,22,63], more and more methods begin to utilize more powerful detectors to obtain higher tracking performance. The one-stage object detector RetinaNet [36] begin to be adopted by several methods such as [40,49]. CenterNet

[88] is the most popular detector adopted by most methods [87,84,70,86,66,61,64] for its simplicity and efficiency. The YOLO series detectors [50,8,63] are also adopted by a large number of methods [68,34,35,16,74] for its excellent balance of accuracy and speed. Most of these methods directly use the detection boxes on a single image for tracking.

However, the number of missing detections and very low scoring detections begin to increase when occlusion or motion blur happens in the video sequence, as is pointed out by video object detection methods [60,42]. Therefore, the information of the previous frames are usually leveraged to enhance the video detection performance.

Detection by tracking. Tracking can also adopted to help obtain more accurate detection boxes. Some methods [53,90,15,14,16,13] utilize single object tracking (SOT) [5] or Kalman filter [30] to predict the location of the tracklets in the following frame and fuse the predicted boxes with the detection boxes to enhance the detection results. Other methods [85,35] leverage tracked boxes in the previous frames to enhance feature representation of the following frame. Recently, Transformer-based [62,19,65,39] detectors [12,91] are adopted by several methods [58,43,79,10] for its strong ability to propagate boxes between frames. Our method also utilize the similarity with tracklets to strength the reliability of detection boxes.

After obtaining the detection boxes by various detectors, most MOT methods [68,84,48,40,34,70,58] only keep the high score detection boxes by a threshold, *i.e.* 0.5, and use those boxes as the input of data association. This is because the low score detection boxes contain many backgrounds which harm the tracking performance. However, we observe that many occluded objects can be correctly detected but have low scores. To reduce missing detections and keep the persistence of trajectories, we keep all the detection boxes and associate across every of them.

2.2 Data Association

Data association is the core of multi-object tracking, which first computes the similarity between tracklets and detection boxes and leverage different strategies to match them according to the similarity.

Similarity metrics. Location, motion and appearance are useful cues for association. SORT [6] combines location and motion cues in a very simple way. It first adopts Kalman filter [30] to predict the location of the tracklets in the new frame and then computes the IoU between the detection boxes and the predicted boxes as the similarity. Some recent methods [87,58,70,56] design networks to learn object motions and achieve more robust results in cases of large camera motion or low frame rate. Location and motion similarity are accurate in the short-range matching. Appearance similarity are helpful in the long-range matching. An object can be re-identified using appearance similarity after being occluded for a long period of time. Appearance similarity can be measured by the cosine similarity of the Re-ID features. DeepSORT [69] adopts a stand-alone Re-ID model to extract appearance features from the detection boxes. Recently, joint detection and Re-ID models [68,84,34,40,83,48,89] becomes more and more popular because of their simplicity and efficiency.

Matching strategy. After similarity computation, matching strategy assigns identities to the objects. This can be done by Hungarian Algorithm [32] or greedy assignment [87]. SORT [6] matches the detection boxes to the tracklets by once matching. DeepSORT [69] proposes a cascaded matching strategy which first matches the detection boxes to the most recent tracklets and then to the lost ones. MOTDT [13] first utilizes appearance similarity to match and then utilize the IoU similarity to match the unmatched tracklets. QDTrack [48] turns the appearance similarity into probability by a bi-directional softmax operation and adopts a nearest neighbor search to accomplish matching. Attention mechanism [62] can directly propagate boxes between frames and perform association implicitly. Recent methods such as [43, 79] propose track queries to find the location of the tracked objects in the following frames. The matching is implicitly performed in the attention interaction process without using Hungarian Algorithm.

All these methods focus on how to design better association methods. However, we argue that the way detection boxes are utilized determines the upper bound of data association and we focus on how to make full use of detection boxes from high scores to low ones in the matching process.

3 BYTE

We propose a simple, effective and generic data association method, BYTE. Different from previous methods [68, 84, 34, 48] which only keep the high score detection boxes, we keep almost every detection box and separate them into high score ones and low score ones. We first associate the high score detection boxes to the tracklets. Some tracklets get unmatched because they do not match to an appropriate high score detection box, which usually happens when occlusion, motion blur or size changing occurs. We then associate the low score detection boxes and these unmatched tracklets to recover the objects in low score detection boxes and filter out background, simultaneously. The pseudo-code of BYTE is shown in Algorithm 1.

The input of BYTE is a video sequence V , along with an object detector Det . We also set a detection score threshold τ . The output of BYTE is the tracks \mathcal{T} of the video and each track contains the bounding box and identity of the object in each frame.

For each frame in the video, we predict the detection boxes and scores using the detector Det . We separate all the detection boxes into two parts \mathcal{D}_{high} and \mathcal{D}_{low} according to the detection score threshold τ . For the detection boxes whose scores are higher than τ , we put them into the high score detection boxes \mathcal{D}_{high} . For the detection boxes whose scores are lower than τ , we put them into the low score detection boxes \mathcal{D}_{low} (line 3 to 13 in Algorithm 1).

After separating the low score detection boxes and the high score detection boxes, we adopt Kalman filter to predict the new locations in the current frame of each track in \mathcal{T} (line 14 to 16 in Algorithm 1).

The first association is performed between the high score detection boxes \mathcal{D}_{high} and all the tracks \mathcal{T} (including the lost tracks \mathcal{T}_{lost}). Similarity\#1 can be computed by either by the IoU or the Re-ID feature distances between the detection boxes \mathcal{D}_{high} and the predicted box of tracks \mathcal{T} . Then, we adopt Hungarian Algorithm [32] to finish the

Algorithm 1: Pseudo-code of BYTE.**Input:** A video sequence V ; object detector Det ; detection score threshold τ **Output:** Tracks \mathcal{T} of the video

```

1 Initialization:  $\mathcal{T} \leftarrow \emptyset$ 
2 for frame  $f_k$  in  $V$  do
    /* Figure 2(a) */
    /* predict detection boxes & scores */
3    $\mathcal{D}_k \leftarrow \text{Det}(f_k)$ 
4    $\mathcal{D}_{high} \leftarrow \emptyset$ 
5    $\mathcal{D}_{low} \leftarrow \emptyset$ 
6   for  $d$  in  $\mathcal{D}_k$  do
7       if  $d.\text{score} > \tau$  then
8            $\mathcal{D}_{high} \leftarrow \mathcal{D}_{high} \cup \{d\}$ 
9       end
10      else
11           $\mathcal{D}_{low} \leftarrow \mathcal{D}_{low} \cup \{d\}$ 
12      end
13  end

    /* predict new locations of tracks */
14  for  $t$  in  $\mathcal{T}$  do
15       $t \leftarrow \text{KalmanFilter}(t)$ 
16  end

    /* Figure 2(b) */
    /* first association */
17  Associate  $\mathcal{T}$  and  $\mathcal{D}_{high}$  using Similarity#1
18   $\mathcal{D}_{remain} \leftarrow$  remaining object boxes from  $\mathcal{D}_{high}$ 
19   $\mathcal{T}_{remain} \leftarrow$  remaining tracks from  $\mathcal{T}$ 

    /* Figure 2(c) */
    /* second association */
20  Associate  $\mathcal{T}_{remain}$  and  $\mathcal{D}_{low}$  using similarity#2
21   $\mathcal{T}_{re-remain} \leftarrow$  remaining tracks from  $\mathcal{T}_{remain}$ 

    /* delete unmatched tracks */
22   $\mathcal{T} \leftarrow \mathcal{T} \setminus \mathcal{T}_{re-remain}$ 

    /* initialize new tracks */
23  for  $d$  in  $\mathcal{D}_{remain}$  do
24       $\mathcal{T} \leftarrow \mathcal{T} \cup \{d\}$ 
25  end
26 end
27 Return:  $\mathcal{T}$ 

```

Track rebirth [69,87] is not shown in the algorithm for simplicity. In green is the key of our method.

matching based on the similarity. We keep the unmatched detections in \mathcal{D}_{remain} and the unmatched tracks in \mathcal{T}_{remain} (line 17 to 19 in Algorithm 1).

BYTE is highly flexible and can be compatible to other different association methods. For example, when BYTE is combined with FairMOT [84], Re-ID feature is added into `* first association *` in Algorithm 1, others are the same. In the experiments, we apply BYTE to 9 different state-of-the-art trackers and achieve notable improvements on almost all the metrics.

The second association is performed between the low score detection boxes \mathcal{D}_{low} and the remaining tracks \mathcal{T}_{remain} after the first association. We keep the unmatched tracks in $\mathcal{T}_{re-remain}$ and just delete all the unmatched low score detection boxes, since we view them as background. (line 20 to 21 in Algorithm 1). We find it important to use IoU alone as the `Similarity#2` in the second association because the low score detection boxes usually contains severe occlusion or motion blur and appearance features are not reliable. Thus, when apply BYTE to other Re-ID based trackers [68,84,48], we do not adopt appearance similarity in the second association.

After the association, the unmatched tracks will be deleted from the tracklets. We do not list the procedure of track rebirth [69,13,87] in Algorithm 1 for simplicity. Actually, it is necessary for the long-range association to preserve the identity of the tracks. For the unmatched tracks $\mathcal{T}_{re-remain}$ after the second association, we put them into \mathcal{T}_{lost} . For each track in \mathcal{T}_{lost} , only when it exists for more than a certain number of frames, *i.e.* 30, we delete it from the tracks \mathcal{T} . Otherwise, we remain the lost tracks \mathcal{T}_{lost} in \mathcal{T} (line 22 in Algorithm 1). Finally, we initialize new tracks from the unmatched high score detection boxes \mathcal{D}_{remain} after the first association. (line 23 to 27 in Algorithm 1). The output of each individual frame is the bounding boxes and identities of the tracks \mathcal{T} in the current frame. Note that we do not output the boxes and identities of \mathcal{T}_{lost} .

To put forwards the state-of-the-art performance of MOT, we design a simple and strong tracker, named ByteTrack, by equipping the high-performance detector YOLOX [25] with our association method BYTE.

4 Experiments

4.1 Setting

Datasets. We evaluate BYTE and ByteTrack on MOT17 [45] and MOT20 [17] datasets under the “private detection” protocol. Both datasets contain training sets and test sets, without validation sets. For ablation studies, we use the first half of each video in the training set of MOT17 for training and the last half for validation following [87]. We train on the combination of CrowdHuman dataset [55] and MOT17 half training set following [87,58,79,70]. We add Cityperson [81] and ETHZ [20] for training following [68,84,34] when testing on the test set of MOT17. We also test ByteTrack on HiEve [38] and BDD100K [78] datasets. HiEve is a large scale human-centric dataset focusing on crowded and complex events. BDD100K is the largest driving video dataset and the dataset splits of the MOT task are 1400 videos for training, 200 videos for validation and 400 videos for testing. It needs to track objects of 8 classes and contains cases of large camera motion.

Metrics. We use the CLEAR metrics [4], including MOTA, FP, FN, IDs, *etc.*, IDF1 [52] and HOTA [41] to evaluate different aspects of the tracking performance. MOTA

is computed based on FP, FN and IDs. Considering the amount of FP and FN are larger than IDs, MOTA focuses more on the detection performance. IDF1 evaluates the identity preservation ability and focus more on the association performance. HOTA is a very recently proposed metric which explicitly balances the effect of performing accurate detection, association and localization. For BDD100K dataset, there are some multi-class metrics such as mMOTA and mIDF1. mMOTA / mIDF1 is computed by averaging the MOTA / IDF1 of all the classes.

Implementation details. For BYTE, the default detection score threshold τ is 0.6, unless otherwise specified. For the benchmark evaluation of MOT17, MOT20 and HiEve, we only use IoU as the similarity metrics. In the linear assignment step, if the IoU between the detection box and the tracklet box is smaller than 0.2, the matching will be rejected. For the lost tracklets, we keep it for 30 frames in case it appears again. For BDD100K, we use UniTrack [67] as the Re-ID model. In ablation study, we use FastReID [28] to extract Re-ID features for MOT17.

For ByteTrack, the detector is YOLOX [25] with YOLOX-X as the backbone and COCO-pretrained model [37] as the initialized weights. For MOT17, the training schedule is 80 epochs on the combination of MOT17, CrowdHuman, Cityperson and ETHZ. For MOT20 and HiEve, we only add CrowdHuman as additional training data. For BDD100K, we do not use additional training data and only train 50 epochs. The input image size is 1440×800 and the shortest side ranges from 576 to 1024 during multi-scale training. The data augmentation includes Mosaic [8] and Mixup [80]. The model is trained on 8 NVIDIA Tesla V100 GPU with batch size of 48. The optimizer is SGD with weight decay of 5×10^{-4} and momentum of 0.9. The initial learning rate is 10^{-3} with 1 epoch warm-up and cosine annealing schedule. The total training time is about 12 hours. Following [25], FPS is measured with FP16-precision [44] and batch size of 1 on a single GPU.

4.2 Ablation Studies on BYTE

Similarity analysis. We choose different types of similarity for the first association and the second association of BYTE. The results are shown in Table 1. We can see that either IoU or Re-ID can be a good choice for *Similarity#1* on MOT17. IoU achieves better MOTA and IDs while Re-ID achieves higher IDF1. On BDD100K, Re-ID achieves much better results than IoU in the first association. This is because BDD100K contains large camera motion and the annotations are in low frame rate, which causes failure of motion cues. It is important to utilize IoU as *Similarity#2* in the second association on both datasets because the low score detection boxes usually contains severe occlusion or motion blur and thus Re-ID features are not reliable. From Table 1 we can find that using IoU as *Similarity#2* increases about 1.0 MOTA compared to Re-ID, which indicates that Re-ID features of the low score detection boxes are not reliable.

Comparisons with other association methods. We compare BYTE with other popular association methods including SORT [6], DeepSORT [69] and MOTDT [13] on the validation set of MOT17 and BDD100K. The results are shown in Table 2.

SORT can be seen as our baseline method because both methods only adopt Kalman filter to predict the object motion. We can find that BYTE improves the MOTA metric

Similarity#1	Similarity#2	MOT17			BDD100K		
		MOTA↑	IDF1↑	IDs↓	mMOTA↑	mIDF1↑	IDs↓
IoU	Re-ID	75.8	77.5	231	39.2	48.3	29172
IoU	IoU	76.6	79.3	159	39.4	48.9	27902
Re-ID	Re-ID	75.2	78.7	276	45.0	53.4	10425
Re-ID	IoU	76.3	80.5	216	45.5	54.8	9140

Table 1. Comparison of different type of similarity metrics used in the first association and the second association of BYTE on the MOT17 and BDD100K validation set.

Method	w/ Re-ID	MOT17			BDD100K		
		MOTA↑	IDF1↑	IDs↓	mMOTA↑	mIDF1↑	IDs↓
SORT		74.6	76.9	291	30.9	41.3	10067
DeepSORT	✓	75.4	77.2	239	24.5	38.2	10720
MOTDT	✓	75.8	77.6	273	26.7	39.8	14520
BYTE (ours)		76.6	79.3	159	39.4	48.9	27902
BYTE (ours)	✓	76.3	80.5	216	45.5	54.8	9140

Table 2. Comparison of different data association methods on the MOT17 and BDD100K validation set. The best results are shown in **bold**.

of SORT from 74.6 to 76.6, IDF1 from 76.9 to 79.3 and decreases IDs from 291 to 159. This highlights the importance of the low score detection boxes and proves the ability of BYTE to recover object boxes from low score one.

DeepSORT utilizes additional Re-ID models to enhance the long-range association. We surprisingly find BYTE also has additional gains compared with DeepSORT. This suggests a simple Kalman filter can perform long-range association and achieve better IDF1 and IDs when the detection boxes are accurate enough. We note that in severe occlusion cases, Re-ID features are vulnerable and may lead to identity switches, instead, motion model behaves more reliably.

MOTDT integrates motion-guided box propagation results along with detection results to associate unreliable detection results with tracklets. Although sharing the similar motivation, MOTDT is behind BYTE by a large margin. We explain that MOTDT uses propagated boxes as tracklet boxes, which may lead to locating drifts in tracking. Instead, BYTE uses low-score detection boxes to re-associate those unmatched tracklets, therefore, tracklet boxes are more accurate.

Table 2 also shows the results on BDD100K dataset. BYTE also outperforms other association methods by a large margin. Kalman filter fails in autonomous driving scenes and it is the main reason for the low performance of SORT, DeepSORT and MOTDT. Thus, we do not use Kalman filter on BDD100K. Additional off-the-shelf Re-ID models greatly improve the performance of BYTE on BDD100K.

Robustness to detection score threshold. The detection score threshold τ_{high} is a sensitive hyper-parameter and needs to be carefully tuned in the task of multi-object tracking. We change it from 0.2 to 0.8 and compare the MOTA and IDF1 score of BYTE and SORT. The results are shown in Figure 3. From the results we can see that




Fig. 3. Comparison of the performances of BYTE and SORT under different detection score thresholds. The results are from the validation set of MOT17.




Fig. 4. Comparison of the number of TPs and FPs in all low score detection boxes and the low score tracked boxes obtained by BYTE. The results are from the validation set of MOT17.

BYTE is more robust to the detection score threshold than SORT. This is because the second association in BYTE recovers the objects whose scores are lower than τ_{high} , and thus considers almost every detection box regardless of the change of τ_{high} .

Analysis on low score detection boxes. To prove the effectiveness of BYTE, we collect the number of TPs and FPs in the low score boxes obtained by BYTE. We use the half training set of MOT17 and CrowdHuman for training and evaluate on the half validation set of MOT17. First, we keep all the low score detection boxes whose scores range from τ_{low} to τ_{high} and classify the TPs and FPs using ground truth annotations. Then, we select the tracking results obtained by BYTE from low score detection boxes. The results of each sequence are shown in Figure 4. We can see that BYTE obtains notably more TPs than FPs from the low score detection boxes even though some sequences (*i.e.* MOT17-02) have much more FPs in all the detection boxes. The obtained TPs notably increases MOTA from 74.6 to 76.6 as is shown in Table 2.

Applications on other trackers. We apply BYTE on 9 different state-of-the-arts trackers, including JDE [68], CStrack [34], FairMOT [84], TraDes [70], QDTrack [48], CenterTrack [87], Chained-Tracker [49], TransTrack [58] and MOTR [79]. Among these trackers, JDE, CStrack, FairMOT, TraDes adopt a combination of motion and Re-ID similarity. QDTrack adopts Re-ID similarity alone. CenterTrack and TraDes predict the

Method	Similarity	w/ BYTE	MOTA↑	IDF1↑	IDs↓
JDE [68]	Motion(K) + Re-ID		60.0	63.6	473
	Motion(K) + Re-ID	✓	60.3 (+0.3)	64.1 (+0.5)	418
	Motion(K)	✓	60.6 (+0.6)	66.0 (+2.4)	360
CSTrack [34]	Motion(K) + Re-ID		68.0	72.3	325
	Motion(K) + Re-ID	✓	69.2 (+1.2)	73.9 (+1.6)	285
	Motion(K)	✓	69.3 (+1.3)	71.7 (-0.6)	279
FairMOT [84]	Motion(K) + Re-ID		69.1	72.8	299
	Motion(K) + Re-ID	✓	70.4 (+1.3)	74.2 (+1.4)	232
	Motion(K)	✓	70.3 (+1.2)	73.2 (+0.4)	236
TraDes [70]	Motion + Re-ID		68.2	71.7	285
	Motion + Re-ID	✓	68.6 (+0.4)	71.1 (-0.6)	259
	Motion(K)	✓	67.9 (-0.3)	72.0 (+0.3)	178
QDTrack [48]	Re-ID		67.3	67.8	377
	Motion(K) + Re-ID	✓	67.7 (+0.4)	72.0 (+4.2)	281
	Motion(K)	✓	67.9 (+0.6)	70.9 (+3.1)	258
CenterTrack [87]	Motion		66.1	64.2	528
	Motion	✓	66.3 (+0.2)	64.8 (+0.6)	334
	Motion(K)	✓	67.4 (+1.3)	74.0 (+9.8)	144
CTracker [49]	Chain		63.1	60.9	755
	Motion(K)	✓	65.0 (+1.9)	66.7 (+5.8)	346
TransTrack [58]	Attention		67.1	68.3	254
	Attention	✓	68.6 (+1.5)	69.0 (+0.7)	232
	Motion(K)	✓	68.3 (+1.2)	72.4 (+4.1)	181
MOTR [79]	Attention		64.7	67.2	346
	Attention	✓	64.3 (-0.4)	69.3 (+2.1)	263
	Motion(K)	✓	65.7 (+1.0)	68.4 (+1.2)	260

Table 3. Results of applying BYTE to 9 different state-of-the-art trackers on the MOT17 validation set. “K” is short for Kalman filter. In green are the improvements of at least +1.0 point.

motion similarity by the learned networks. Chained-Tracker adopts the chain structure and outputs the results of two consecutive frames simultaneously and associate in the same frame by IoU. TransTrack and MOTR adopt the attention mechanism to propagate boxes among frames. Their results are shown in the first line of each tracker in Table 3. To evaluate the effectiveness of BYTE, we design two different modes to apply BYTE to these trackers.

- The first mode is to insert BYTE into the original association methods of different trackers, as is shown in the second line of the results of each tracker in Table 3. Take FairMOT[84] for example, after the original association is done, we select all the unmatched tracklets and associate them with the low score detection boxes following the `* second association *` in Algorithm 1. Note that for the low score objects, the Re-ID features are not reliable so we only adopt the IoU between the detection boxes and the tracklet boxes after motion prediction as the similarity. We do

Tracker	MOTA↑	IDF1↑	HOTA↑	FP↓	FN↓	IDs↓	FPS↑
Tube_TK [47]	63.0	58.6	48.0	27060	177483	4137	3.0
MOTR [79]	65.1	66.4	-	45486	149307	2049	-
CTracker [49]	66.6	57.4	49.0	22284	160491	5529	6.8
CenterTrack [87]	67.8	64.7	52.2	18498	160332	3039	17.5
QuasiDense [48]	68.7	66.3	53.9	26589	146643	3378	20.3
TraDes [70]	69.1	63.9	52.7	20892	150060	3555	17.5
MAT [26]	69.5	63.1	53.8	30660	138741	2844	9.0
SOTMOT [86]	71.0	71.9	-	39537	118983	5184	16.0
TransCenter [73]	73.2	62.2	54.5	23112	123738	4614	1.0
GSDT [66]	73.2	66.5	55.2	26397	120666	3891	4.9
Semi-TCL [33]	73.3	73.2	59.8	22944	124980	2790	-
FairMOT [84]	73.7	72.3	59.3	27507	117477	3303	25.9
RelationTrack [77]	73.8	74.7	61.0	27999	118623	1374	8.5
PermaTrackPr [61]	73.8	68.9	55.5	28998	115104	3699	11.9
CSTrack [34]	74.9	72.6	59.3	23847	114303	3567	15.8
TransTrack [58]	75.2	63.5	54.1	50157	86442	3603	10.0
FUFET [54]	76.2	68.0	57.9	32796	98475	3237	6.8
SiamMOT [35]	76.3	72.3	-	-	-	-	12.8
CorrTracker [64]	76.5	73.6	60.7	29808	99510	3369	15.6
TransMOT [16]	76.7	75.1	61.7	36231	93150	2346	9.6
ReMOT [75]	77.0	72.0	59.7	33204	93612	2853	1.8
ByteTrack (ours)	80.3	77.3	63.1	25491	83721	2196	29.6

Table 4. Comparison of the state-of-the-art methods under the “private detector” protocol on MOT17 test set. The best results are shown in **bold**. MOT17 contains rich scenes and half of the sequences are captured with camera motion. **ByteTrack** ranks 1st among all the trackers on the leaderboard of MOT17 and outperforms the second one ReMOT by a large margin on almost all the metrics. It also has the highest running speed among all trackers.

not apply the first mode of BYTE to Chained-Tracker because we find it is difficult to implement in the chain structure.

- The second mode is to directly use the detection boxes of these trackers and associate using the whole procedure in Algorithm 1, as is shown in the third line of the results of each tracker in Table 3.

We can see that in both modes, BYTE can bring stable improvements over almost all the metrics including MOTA, IDF1 and IDs. For example, BYTE increases CenterTrack by 1.3 MOTA and 9.8 IDF1, Chained-Tracker by 1.9 MOTA and 5.8 IDF1, TransTrack by 1.2 MOTA and 4.1 IDF1. The results in Table 3 indicate that BYTE has strong generalization ability and can be easily applied to existing trackers.

4.3 Benchmark Evaluation

We compare ByteTrack with the state-of-the-art trackers on the test set of MOT17, MOT20 and HiEve under the private detection protocol in Table 4, Table 5 and Table 6, respectively. All the results are directly obtained from the official MOT Challenge evaluation server and the Human in Events server.

<https://motchallenge.net>

<http://humaninevents.org>

Tracker	MOTA↑	IDF1↑	HOTA↑	FP↓	FN↓	IDs↓	FPS↑
MLT [82]	48.9	54.6	43.2	45660	216803	2187	3.7
FairMOT [84]	61.8	67.3	54.6	103440	88901	5243	13.2
TransCenter [73]	61.9	50.4	-	45895	146347	4653	1.0
TransTrack [58]	65.0	59.4	48.5	27197	150197	3608	7.2
CorrTracker [64]	65.2	69.1	-	79429	95855	5183	8.5
Semi-TCL [33]	65.2	70.1	55.3	61209	114709	4139	-
CSTrack [34]	66.6	68.6	54.0	25404	144358	3196	4.5
GSDT [66]	67.1	67.5	53.6	31913	135409	3131	0.9
SiamMOT [35]	67.1	69.1	-	-	-	-	4.3
RelationTrack [77]	67.2	70.5	56.5	61134	104597	4243	2.7
SOTMOT [86]	68.6	71.4	-	57064	101154	4209	8.5
ByteTrack (ours)	77.8	75.2	61.3	26249	87594	1223	17.5

Table 5. Comparison of the state-of-the-art methods under the “private detector” protocol on MOT20 test set. The best results are shown in **bold**. The scenes in MOT20 are much more crowded than those in MOT17. **ByteTrack** ranks 1st among all the trackers on the leaderboard of MOT20 and outperforms the second one SOTMOT by a large margin on all the metrics. It also has the highest running speed among all trackers.

Tracker	MOTA↑	IDF1↑	MT↑	ML↓	FP↓	FN↓	IDs↓
DeepSORT [69]	27.1	28.6	8.5%	41.5%	5894	42668	2220
MOTDT [13]	26.1	32.9	8.7%	54.6%	6318	43577	1599
IOUtracker [7]	38.6	38.6	28.3%	27.6%	9640	28993	4153
JDE [68]	33.1	36.0	15.1%	24.1%	9526	33327	3747
FairMOT [84]	35.0	46.7	16.3%	44.2%	6523	37750	995
CenterTrack [87]	40.9	45.1	10.8%	32.2%	3208	36414	1568
ByteTrack (Ours)	61.7	63.1	38.3%	21.6%	2822	22852	1031

Table 6. Comparison of the state-of-the-art methods under the “private detector” protocol on HiEve test set. The best results are shown in **bold**. HiEve has more complex events than MOT17 and MOT20. **ByteTrack** ranks 1st among all the trackers on the leaderboard of HiEve and outperforms the second one CenterTrack by a large margin on all the metrics.

MOT17. ByteTrack ranks 1st among all the trackers on the leaderboard of MOT17. Not only does it achieve the best accuracy (*i.e.* 80.3 MOTA, 77.3 IDF1 and 63.1 HOTA), but also runs with highest running speed (30 FPS). It outperforms the second-performance tracker [75] by a large margin (*i.e.* +3.3 MOTA, +5.3 IDF1 and +3.4 HOTA). Also, we use less training data than many high performance methods such as [84,34,64,54,35] (29K images vs. 73K images). It is worth noting that we only leverage the simplest similarity computation method Kalman filter in the association step compared to other methods [84,34,48,66,79,58] which additionally adopt Re-ID similarity or attention mechanisms. All these indicate that ByteTrack is a simple and strong tracker.

MOT20. Compared with MOT17, MOT20 has much more crowded scenarios and occlusion cases. The average number of pedestrians in an image is 170 in the test set of MOT20. ByteTrack also ranks 1st among all the trackers on the leaderboard of MOT20 and outperforms other trackers by a large margin on almost all the metrics. For exam-

Tracker	split	mMOTA↑	mIDF1↑	MOTA↑	IDF1↑	FN↓	FP↓	IDs↓	MT↑	ML↓
Yu <i>et al.</i> [78]	val	25.9	44.5	56.9	66.8	122406	52372	8315	8396	3795
QDTrack [48]	val	36.6	50.8	63.5	71.5	108614	46621	6262	9481	3034
ByteTrack(Ours)	val	45.5	54.8	69.1	70.4	92805	34998	9140	9626	3005
Yu <i>et al.</i> [78]	test	26.3	44.7	58.3	68.2	213220	100230	14674	16299	6017
DeepBlueAI*	test	31.6	38.7	56.9	56.0	292063	35401	25186	10296	12266
Madamada*	test	33.6	43.0	59.8	55.7	209339	76612	42901	16774	5004
QDTrack[48]	test	35.5	52.3	64.3	72.3	201041	80054	10790	17353	5167
ByteTrack(Ours)	test	40.1	55.8	69.6	71.3	169073	63869	15466	18057	5107

Table 7. Comparison of the state-of-the-art methods on BDD100K test set. The best results are shown in **bold**. **ByteTrack** ranks 1st among all the trackers on the leaderboard of BDD100K and outperforms the second one QDTrack by a large margin on most metrics. The methods denoted by * are the ones reported on the leaderboard of BDD100K.

ple, it increases MOTA from 68.6 to 77.8, IDF1 from 71.4 to 75.2 and decreases IDs by 71% from 4209 to 1223. It is worth noting that ByteTrack achieves extremely low identity switches, which further indicates that associating every detection boxes is very effective under occlusion cases.

Human in Events. Compared with MOT17 and MOT20, HiEve contains more complex events and more diverse camera views. We train ByteTrack on CrowdHuman dataset and the training set of HiEve. ByteTrack also ranks 1st among all the trackers on the leaderboard of HiEve and outperforms other state-of-the-art trackers by a large margin. For example, it increases MOTA from 40.9 to 61.3 and IDF1 from 45.1 to 62.9. The superior results indicate that ByteTrack is robust to complex scenes.

BDD100K. BDD100K is multiple categories tracking dataset in autonomous driving scenes. The challenges include low frame rate and large camera motion. We utilize a simple ResNet-50 ImageNet classification model from UniTrack [67] to compute appearance similarity. ByteTrack ranks first on the leaderboard of BDD100K. It increases mMOTA from 36.6 to 45.5 on the validation set and 35.5 to 40.1 on the test set, which indicates that ByteTrack can also handle the challenges in autonomous driving scenes.

5 Conclusion

We present a simple yet effective data association method BYTE for multi-object tracking. BYTE can be easily applied to existing trackers and achieve consistent improvements. We also propose a strong tracker ByteTrack, which achieves 80.3 MOTA, 77.3 IDF1 and 63.1 HOTA on MOT17 test set with 30 FPS. ByteTrack is very robust to occlusion for its accurate detection performance and the help of associating low score detection boxes. It also sheds light on how to make the best use of detection results to enhance multi-object tracking. We hope the high accuracy, fast speed and simplicity of ByteTrack can make it attractive in real applications.

Acknowledgement This work was in part supported by NSFC (No. 61733007 and No. 61876212). Ping Luo is supported by the General Research Fund of HK No.27208720, No.17212120, and No.17200622.

References

1. Bae, S.H., Yoon, K.J.: Robust online multi-object tracking based on tracklet confidence and online discriminative appearance learning. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1218–1225 (2014) [1](#)
2. Berclaz, J., Fleuret, F., Turetken, E., Fua, P.: Multiple object tracking using k-shortest paths optimization. *IEEE transactions on pattern analysis and machine intelligence* **33**(9), 1806–1819 (2011) [1](#)
3. Bergmann, P., Meinhardt, T., Leal-Taixe, L.: Tracking without bells and whistles. In: ICCV. pp. 941–951 (2019) [1](#), [3](#)
4. Bernardin, K., Stiefelhagen, R.: Evaluating multiple object tracking performance: the clear mot metrics. *EURASIP Journal on Image and Video Processing* **2008**, 1–10 (2008) [1](#), [7](#)
5. Bertinetto, L., Valmadre, J., Henriques, J.F., Vedaldi, A., Torr, P.H.: Fully-convolutional siamese networks for object tracking. In: European conference on computer vision. pp. 850–865. Springer (2016) [4](#)
6. Bewley, A., Ge, Z., Ott, L., Ramos, F., Upcroft, B.: Simple online and realtime tracking. In: ICIP. pp. 3464–3468. IEEE (2016) [1](#), [2](#), [4](#), [5](#), [8](#)
7. Bochinski, E., Eiselein, V., Sikora, T.: High-speed tracking-by-detection without using image information. In: 2017 14th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS). pp. 1–6. IEEE (2017) [13](#)
8. Bochkovskiy, A., Wang, C.Y., Liao, H.Y.M.: YOLOv4: Optimal speed and accuracy of object detection. *arXiv preprint arXiv:2004.10934* (2020) [4](#), [8](#)
9. Brasó, G., Leal-Taixé, L.: Learning a neural solver for multiple object tracking. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6247–6257 (2020) [3](#)
10. Cai, J., Xu, M., Li, W., Xiong, Y., Xia, W., Tu, Z., Soatto, S.: Memot: Multi-object tracking with memory. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8090–8100 (2022) [4](#)
11. Cai, Z., Vasconcelos, N.: Cascade r-cnn: Delving into high quality object detection. In: CVPR. pp. 6154–6162 (2018) [3](#)
12. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: European Conference on Computer Vision. pp. 213–229. Springer (2020) [4](#)
13. Chen, L., Ai, H., Zhuang, Z., Shang, C.: Real-time multiple people tracking with deeply learned candidate selection and person re-identification. In: 2018 IEEE International Conference on Multimedia and Expo (ICME). pp. 1–6. IEEE (2018) [1](#), [3](#), [4](#), [5](#), [7](#), [8](#), [13](#)
14. Chu, P., Fan, H., Tan, C.C., Ling, H.: Online multi-object tracking with instance-aware tracker and dynamic model refreshment. In: 2019 IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 161–170. IEEE (2019) [4](#)
15. Chu, P., Ling, H.: Famnet: Joint learning of feature, affinity and multi-dimensional assignment for online multiple object tracking. In: ICCV. pp. 6172–6181 (2019) [3](#), [4](#)
16. Chu, P., Wang, J., You, Q., Ling, H., Liu, Z.: Transmot: Spatial-temporal graph transformer for multiple object tracking. *arXiv preprint arXiv:2104.00194* (2021) [4](#), [12](#)
17. Dendorfer, P., Rezatofighi, H., Milan, A., Shi, J., Cremers, D., Reid, I., Roth, S., Schindler, K., Leal-Taixé, L.: Mot20: A benchmark for multi object tracking in crowded scenes. *arXiv preprint arXiv:2003.09003* (2020) [3](#), [7](#)
18. Dicle, C., Camps, O.I., Szaier, M.: The way they move: Tracking multiple targets with similar appearance. In: Proceedings of the IEEE international conference on computer vision. pp. 2304–2311 (2013) [1](#)

19. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020) [4](#)
20. Ess, A., Leibe, B., Schindler, K., Van Gool, L.: A mobile vision system for robust multi-person tracking. In: CVPR. pp. 1–8. IEEE (2008) [7](#)
21. Fang, Y., Liao, B., Wang, X., Fang, J., Qi, J., Wu, R., Niu, J., Liu, W.: You only look at one sequence: Rethinking transformer in vision through object detection. Advances in Neural Information Processing Systems **34**, 26183–26197 (2021) [3](#)
22. Fang, Y., Yang, S., Wang, S., Ge, Y., Shan, Y., Wang, X.: Unleashing vanilla vision transformer with masked image modeling for object detection. arXiv preprint arXiv:2204.02964 (2022) [3](#)
23. Felzenszwalb, P., McAllester, D., Ramanan, D.: A discriminatively trained, multiscale, deformable part model. In: CVPR. pp. 1–8. IEEE (2008) [3](#)
24. Fu, J., Zong, L., Li, Y., Li, K., Yang, B., Liu, X.: Model adaption object detection system for robot. In: 2020 39th Chinese Control Conference (CCC). pp. 3659–3664. IEEE (2020) [3](#)
25. Ge, Z., Liu, S., Wang, F., Li, Z., Sun, J.: Yolox: Exceeding yolo series in 2021. arXiv preprint arXiv:2107.08430 (2021) [3](#), [7](#), [8](#)
26. Han, S., Huang, P., Wang, H., Yu, E., Liu, D., Pan, X., Zhao, J.: Mat: Motion-aware multi-object tracking. arXiv preprint arXiv:2009.04794 (2020) [12](#)
27. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: ICCV. pp. 2961–2969 (2017) [3](#)
28. He, L., Liao, X., Liu, W., Liu, X., Cheng, P., Mei, T.: Fastreid: A pytorch toolbox for general instance re-identification. arXiv preprint arXiv:2006.02631 (2020) [8](#)
29. Hornakova, A., Henschel, R., Rosenhahn, B., Swoboda, P.: Lifted disjoint paths with application in multiple object tracking. In: International Conference on Machine Learning. pp. 4364–4375. PMLR (2020) [3](#)
30. Kalman, R.E.: A new approach to linear filtering and prediction problems. J. Fluids Eng. **82**(1), 35–45 (1960) [2](#), [4](#)
31. Khurana, T., Dave, A., Ramanan, D.: Detecting invisible people. arXiv preprint arXiv:2012.08419 (2020) [2](#)
32. Kuhn, H.W.: The hungarian method for the assignment problem. Naval research logistics quarterly **2**(1-2), 83–97 (1955) [5](#)
33. Li, W., Xiong, Y., Yang, S., Xu, M., Wang, Y., Xia, W.: Semi-tcl: Semi-supervised track contrastive representation learning. arXiv preprint arXiv:2107.02396 (2021) [12](#), [13](#)
34. Liang, C., Zhang, Z., Lu, Y., Zhou, X., Li, B., Ye, X., Zou, J.: Rethinking the competition between detection and reid in multi-object tracking. arXiv preprint arXiv:2010.12138 (2020) [3](#), [4](#), [5](#), [7](#), [10](#), [11](#), [12](#), [13](#)
35. Liang, C., Zhang, Z., Zhou, X., Li, B., Lu, Y., Hu, W.: One more check: Making” fake background” be tracked again. arXiv preprint arXiv:2104.09441 (2021) [4](#), [12](#), [13](#)
36. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: ICCV. pp. 2980–2988 (2017) [3](#)
37. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: ECCV. pp. 740–755. Springer (2014) [8](#)
38. Lin, W., Liu, H., Liu, S., Li, Y., Qian, R., Wang, T., Xu, N., Xiong, H., Qi, G.J., Sebe, N.: Human in events: A large-scale benchmark for human-centric video analysis in complex events. arXiv preprint arXiv:2005.04490 (2020) [3](#), [7](#)
39. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. arXiv preprint arXiv:2103.14030 (2021) [4](#)

40. Lu, Z., Rathod, V., Votel, R., Huang, J.: Retinatrack: Online single stage joint detection and tracking. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 14668–14678 (2020) 3, 4
41. Luiten, J., Osep, A., Dendorfer, P., Torr, P., Geiger, A., Leal-Taixé, L., Leibe, B.: Hota: A higher order metric for evaluating multi-object tracking. International journal of computer vision **129**(2), 548–578 (2021) 1, 7
42. Luo, H., Xie, W., Wang, X., Zeng, W.: Detect or track: Towards cost-effective video object detection/tracking. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 33, pp. 8803–8810 (2019) 4
43. Meinhardt, T., Kirillov, A., Leal-Taixe, L., Feichtenhofer, C.: Trackformer: Multi-object tracking with transformers. arXiv preprint arXiv:2101.02702 (2021) 4, 5
44. Micikevicius, P., Narang, S., Alben, J., Diamos, G., Elsen, E., Garcia, D., Ginsburg, B., Houston, M., Kuchaiev, O., Venkatesh, G., et al.: Mixed precision training. arXiv preprint arXiv:1710.03740 (2017) 8
45. Milan, A., Leal-Taixé, L., Reid, I., Roth, S., Schindler, K.: Mot16: A benchmark for multi-object tracking. arXiv preprint arXiv:1603.00831 (2016) 3, 7
46. Milan, A., Roth, S., Schindler, K.: Continuous energy minimization for multitarget tracking. IEEE transactions on pattern analysis and machine intelligence **36**(1), 58–72 (2013) 1
47. Pang, B., Li, Y., Zhang, Y., Li, M., Lu, C.: 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. pp. 6308–6318 (2020) 12
48. Pang, J., Qiu, L., Li, X., Chen, H., Li, Q., Darrell, T., Yu, F.: Quasi-dense similarity learning for multiple object tracking. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 164–173 (2021) 3, 4, 5, 7, 10, 11, 12, 13, 14
49. Peng, J., Wang, C., Wan, F., Wu, Y., Wang, Y., Tai, Y., Wang, C., Li, J., Huang, F., Fu, Y.: Chained-tracker: Chaining paired attentive regression results for end-to-end joint multiple-object detection and tracking. In: European Conference on Computer Vision. pp. 145–161. Springer (2020) 3, 10, 11, 12
50. Redmon, J., Farhadi, A.: Yolov3: An incremental improvement. arXiv preprint arXiv:1804.02767 (2018) 3, 4
51. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: Advances in neural information processing systems. pp. 91–99 (2015) 3
52. Ristani, E., Solera, F., Zou, R., Cucchiara, R., Tomasi, C.: Performance measures and a data set for multi-target, multi-camera tracking. In: ECCV. pp. 17–35. Springer (2016) 7
53. Sanchez-Matilla, R., Poiesi, F., Cavallaro, A.: Online multi-target tracking with strong and weak detections. In: ECCV. pp. 84–99. Springer (2016) 4
54. Shan, C., Wei, C., Deng, B., Huang, J., Hua, X.S., Cheng, X., Liang, K.: Tracklets predicting based adaptive graph tracking. arXiv preprint arXiv:2010.09015 (2020) 12, 13
55. Shao, S., Zhao, Z., Li, B., Xiao, T., Yu, G., Zhang, X., Sun, J.: Crowddhuman: A benchmark for detecting human in a crowd. arXiv preprint arXiv:1805.00123 (2018) 7
56. Shuai, B., Berneshawi, A., Li, X., Modolo, D., Tighe, J.: Siammot: Siamese multi-object tracking. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12372–12382 (2021) 4
57. Sun, P., Jiang, Y., Xie, E., Shao, W., Yuan, Z., Wang, C., Luo, P.: What makes for end-to-end object detection? In: Proceedings of the 38th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 139, pp. 9934–9944. PMLR (2021) 3
58. Sun, P., Jiang, Y., Zhang, R., Xie, E., Cao, J., Hu, X., Kong, T., Yuan, Z., Wang, C., Luo, P.: Transtrack: Multiple-object tracking with transformer. arXiv preprint arXiv:2012.15460 (2020) 1, 3, 4, 7, 10, 11, 12, 13

59. Sun, P., Zhang, R., Jiang, Y., Kong, T., Xu, C., Zhan, W., Tomizuka, M., Li, L., Yuan, Z., Wang, C., et al.: Sparse r-cnn: End-to-end object detection with learnable proposals. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14454–14463 (2021) [3](#)
60. Tang, P., Wang, C., Wang, X., Liu, W., Zeng, W., Wang, J.: Object detection in videos by high quality object linking. *IEEE transactions on pattern analysis and machine intelligence* **42**(5), 1272–1278 (2019) [4](#)
61. Tokmakov, P., Li, J., Burgard, W., Gaidon, A.: Learning to track with object permanence. *arXiv preprint arXiv:2103.14258* (2021) [2](#), [4](#), [12](#)
62. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. In: Advances in neural information processing systems. pp. 5998–6008 (2017) [4](#), [5](#)
63. Wang, C.Y., Bochkovskiy, A., Liao, H.Y.M.: Yolov7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors. *arXiv preprint arXiv:2207.02696* (2022) [3](#), [4](#)
64. Wang, Q., Zheng, Y., Pan, P., Xu, Y.: Multiple object tracking with correlation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3876–3886 (2021) [4](#), [12](#), [13](#)
65. Wang, W., Xie, E., Li, X., Fan, D.P., Song, K., Liang, D., Lu, T., Luo, P., Shao, L.: Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. *arXiv preprint arXiv:2102.12122* (2021) [4](#)
66. Wang, Y., Kitani, K., Weng, X.: Joint object detection and multi-object tracking with graph neural networks. *arXiv preprint arXiv:2006.13164* (2020) [4](#), [12](#), [13](#)
67. Wang, Z., Zhao, H., Li, Y.L., Wang, S., Torr, P.H., Bertinetto, L.: Do different tracking tasks require different appearance models? *arXiv preprint arXiv:2107.02156* (2021) [8](#), [14](#)
68. Wang, Z., Zheng, L., Liu, Y., Li, Y., Wang, S.: Towards real-time multi-object tracking. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XI 16. pp. 107–122. Springer (2020) [3](#), [4](#), [5](#), [7](#), [10](#), [11](#), [13](#)
69. Wojke, N., Bewley, A., Paulus, D.: Simple online and realtime tracking with a deep association metric. In: 2017 IEEE international conference on image processing (ICIP). pp. 3645–3649. IEEE (2017) [1](#), [4](#), [5](#), [6](#), [7](#), [8](#), [13](#)
70. Wu, J., Cao, J., Song, L., Wang, Y., Yang, M., Yuan, J.: Track to detect and segment: An on-line multi-object tracker. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12352–12361 (2021) [3](#), [4](#), [7](#), [10](#), [11](#), [12](#)
71. Xiang, Y., Alahi, A., Savarese, S.: Learning to track: Online multi-object tracking by decision making. In: ICCV. pp. 4705–4713 (2015) [1](#)
72. Xu, J., Cao, Y., Zhang, Z., Hu, H.: Spatial-temporal relation networks for multi-object tracking. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3988–3998 (2019) [3](#)
73. Xu, Y., Ban, Y., Delorme, G., Gan, C., Rus, D., Alameda-Pineda, X.: Transcenter: Transformers with dense queries for multiple-object tracking. *arXiv preprint arXiv:2103.15145* (2021) [12](#), [13](#)
74. Yan, B., Jiang, Y., Sun, P., Wang, D., Yuan, Z., Luo, P., Lu, H.: Towards grand unification of object tracking. *arXiv preprint arXiv:2207.07078* (2022) [4](#)
75. Yang, F., Chang, X., Sakti, S., Wu, Y., Nakamura, S.: Remot: A model-agnostic refinement for multiple object tracking. *Image and Vision Computing* **106**, 104091 (2021) [12](#), [13](#)
76. Yang, F., Choi, W., Lin, Y.: Exploit all the layers: Fast and accurate cnn object detector with scale dependent pooling and cascaded rejection classifiers. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2129–2137 (2016) [3](#)
77. Yu, E., Li, Z., Han, S., Wang, H.: Relationtrack: Relation-aware multiple object tracking with decoupled representation. *arXiv preprint arXiv:2105.04322* (2021) [12](#), [13](#)

78. Yu, F., Chen, H., Wang, X., Xian, W., Chen, Y., Liu, F., Madhavan, V., Darrell, T.: Bdd100k: A diverse driving dataset for heterogeneous multitask learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2636–2645 (2020) [3](#), [7](#), [14](#)
79. Zeng, F., Dong, B., Wang, T., Chen, C., Zhang, X., Wei, Y.: Motr: End-to-end multiple-object tracking with transformer. arXiv preprint arXiv:2105.03247 (2021) [3](#), [4](#), [5](#), [7](#), [10](#), [11](#), [12](#), [13](#)
80. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. arXiv preprint arXiv:1710.09412 (2017) [8](#)
81. Zhang, S., Benenson, R., Schiele, B.: Citypersons: A diverse dataset for pedestrian detection. In: CVPR. pp. 3213–3221 (2017) [7](#)
82. Zhang, Y., Sheng, H., Wu, Y., Wang, S., Ke, W., Xiong, Z.: Multiplex labeling graph for near-online tracking in crowded scenes. IEEE Internet of Things Journal 7(9), 7892–7902 (2020) [13](#)
83. Zhang, Y., Wang, C., Wang, X., Liu, W., Zeng, W.: Voxeltrack: Multi-person 3d human pose estimation and tracking in the wild. IEEE Transactions on Pattern Analysis and Machine Intelligence (2022) [4](#)
84. Zhang, Y., Wang, C., Wang, X., Zeng, W., Liu, W.: Fairmot: On the fairness of detection and re-identification in multiple object tracking. International Journal of Computer Vision 129(11), 3069–3087 (2021) [1](#), [3](#), [4](#), [5](#), [7](#), [10](#), [11](#), [12](#), [13](#)
85. Zhang, Z., Cheng, D., Zhu, X., Lin, S., Dai, J.: Integrated object detection and tracking with tracklet-conditioned detection. arXiv preprint arXiv:1811.11167 (2018) [4](#)
86. Zheng, L., Tang, M., Chen, Y., Zhu, G., Wang, J., Lu, H.: Improving multiple object tracking with single object tracking. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2453–2462 (2021) [4](#), [12](#), [13](#)
87. Zhou, X., Koltun, V., Krähenbühl, P.: Tracking objects as points. In: European Conference on Computer Vision. pp. 474–490. Springer (2020) [3](#), [4](#), [5](#), [6](#), [7](#), [10](#), [11](#), [12](#), [13](#)
88. Zhou, X., Wang, D., Krähenbühl, P.: Objects as points. arXiv preprint arXiv:1904.07850 (2019) [4](#)
89. Zhou, X., Yin, T., Koltun, V., Krähenbühl, P.: Global tracking transformers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8771–8780 (2022) [4](#)
90. Zhu, J., Yang, H., Liu, N., Kim, M., Zhang, W., Yang, M.H.: Online multi-object tracking with dual matching attention networks. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 366–382 (2018) [3](#), [4](#)
91. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. arXiv preprint arXiv:2010.04159 (2020) [4](#)