

Simple Cues Lead to a Strong Multi-Object Tracker

Jenny Seidenschwarz^{1*} Guillem Brasó^{1,2} Victor Castro Serrano¹ Ismail Elezi¹ Laura Leal-Taixé^{1†}

¹Technical University of Munich

²Munich Center for Machine Learning

Abstract

For a long time, the most common paradigm in Multi-Object Tracking was tracking-by-detection (TbD), where objects are first detected and then associated over video frames. For association, most models resorted to motion and appearance cues, e.g., re-identification networks. Recent approaches based on attention propose to learn the cues in a data-driven manner, showing impressive results. In this paper, we ask ourselves whether simple good old TbD methods are also capable of achieving the performance of end-to-end models. To this end, we propose two key ingredients that allow a standard re-identification network to excel at appearance-based tracking. We extensively analyse its failure cases, and show that a combination of our appearance features with a simple motion model leads to strong tracking results. Our tracker generalizes to four public datasets, namely MOT17, MOT20, BDD100k, and DanceTrack, achieving state-of-the-art performance. <https://github.com/dvl-tum/GHOST>.

1. Introduction

Multi-Object Tracking (MOT) aims at finding the trajectories of all moving objects in a video. The dominant paradigm in the field has long been tracking-by-detection, which divides tracking into two steps: (i) frame-wise object detection, (ii) data association to link the detections and form trajectories. One of the simplest forms of data association for online trackers is frame-by-frame matching using the Hungarian algorithm [26]. Matching is often driven by cues such as appearance, e.g., re-identification (reID) features [11, 20, 35, 45, 57, 59, 75], or motion cues [4, 42, 47, 65, 74]. Even recent trackers propose data-driven motion priors [4, 54, 64, 79] or appearance cues, which may include external reID models [4, 64].

Most recent trackers based on Transformers [36, 56, 71] learn all necessary cues from data through self- and cross-attention between frames and tracked objects. While this implicitly gets rid of any heuristic typically embedded in

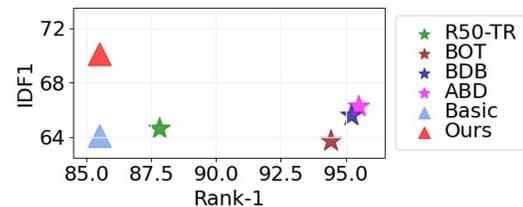


Figure 1. IDF1/Rank-1 of different state-of-the-art re-ID approaches. R50-TR [78], BOT [35], BDB [14], ABD [10]. *Basic* is our baseline, *Ours* is our appearance model.

the handcrafted appearance and motion cues, and could be the path to more general trackers, the training strategies are highly complex and the amount of data needed to train such models is very large, to the point where MOT datasets [15] are not enough and methods rely on pre-training on detection datasets such as CrowdHuman [52].

While interesting and challenging from a research point of view, it is questionable whether we should also follow the path of *learning everything* in multi-object tracking, when there are strong priors that we know how to define and leverage, such as the good old appearance and motion cues. As we show in this paper, there are key observations that need to be made in order to properly leverage such cues. These observations might seem simple and obvious but have been largely overlooked by the community. If we spend as much time in properly understanding and implementing such cues as we do in training Transformers, we will be rewarded with a simple Hungarian tracker with appearance and motion cues that still dominates state-of-the-art on multiple benchmarks, and does not even need to be trained on any tracking data.

Our first observation is that simply using state-of-the-art re-identification (reID) networks for appearance matching is not enough for the real scenarios of MOT. In Figure 1, we visualize the performance of several state-of-the-art reID approaches on Market-1501 dataset [77] (x-axis), as well as the model’s performance when used in a simple matching-based tracker (y-axis). It shows that the reID performance does not necessarily translate to MOT performance. We identify two problems causing the weak performance of reID models on MOT: (i)

*Correspondance to j.seidenschwarz@tum.de.

†Currently at NVIDIA.

reID models need to account for the different challenges expected at different time horizons, *i.e.*, while in nearby frames appearance of objects will vary minimally, in longer time gaps more severe changes are expected, *e.g.*, caused by (partial-) occlusions and (ii) reID performance tends to be inconsistent across MOT sequences because of their varying image statistics, which in turn differs from the relatively stable conditions of the corresponding reID training dataset. We propose two simple but key design choices to overcome the aforementioned problems, *i.e.*, on-the-fly domain adaptation, and different policies for **active** and **inactive** tracks. Moreover, we conduct an extensive analysis under different conditions of visibility, occlusion time, and camera movement, to determine in which situations reID is not enough and we are in need of a motion model. We combine our reID with a simple linear motion model using a weighted sum so that each cue can be given more weight when needed for different datasets.

Our findings culminate in our proposed **Good Old Hungarian Simple Tracker** or **GHOST** (the order of the letters of the acronym does not change the product) that generalizes to four different datasets, remarkably outperforming the state-of-the-art while, most notably, *never being trained on any tracking dataset*.

In summary, we make the following contributions:

- We provide key design choices that significantly boost the performance of reID models for the MOT task.
- We extensively analyze in which underlying situations appearance is not sufficient and when it can be backed up by motion.
- We generalize to four datasets achieving state-of-the-art performance by combining appearance and motion in our simple and general TbD online tracker GHOST.

With this paper, we hope to show the importance of domain-specific knowledge and the impact it can have, even on the simple good old models. Our observations, *i.e.*, the importance of domain adaptation, the different handling of short- and long-term associations as well as the interplay between motion and appearance are straightforward, almost embedded into the subconscious of the tracking community, and yet they have largely been overlooked by recent methods. Introducing our simple but strong tracker, we hope our observations will inspire future work to integrate such observations into sophisticated models further improving the state of the art, in a solution where data and priors will gladly meet.

2. Related Work

In the last years, TbD was the most common paradigm used in MOT [4, 7, 21, 40, 41, 58, 62, 65, 74]. Pedestrians

are first detected using object detectors [46, 47, 66]. Then, detections are associated across frames to form trajectories corresponding to a certain identity utilizing motion, location, appearance cues, or a combination of them. The association can either be solved frame-by-frame for online applications or offline in a track-wise manner over the sequence.

Graph-Based Approaches. One common formalism to perform data association following the TbD paradigm is viewing each detection as a node in a graph with edges linking several nodes over the temporal domain to form trajectories. Determining which nodes are connected can then be solved using maximum flow [3] or minimum cost approaches [22, 44, 72] by, *e.g.*, taking motion models into account [28]. Recent advances combined track-wise graph-based models with neural networks [7]. We challenge those recent advances by showing that we can obtain strong TbD trackers without using a complex graph model combining our strong but simple cues.

Motion-Based Association. Different from the graph-based approaches, many TbD approaches perform frame-by-frame association directly using motion and location cues from detections and existing trajectories [5, 6, 40, 42, 79]. For short term preservation, those trackers exploit that given two nearby frames object displacements tend to be small. This allows them to utilize spatial proximity for matching by exploiting, *e.g.*, Kalman filters [5]. Taking this idea further, approaches following the tracking-by-regression paradigm utilize object detectors [4, 79] to regress bounding box positions. Recent advances introduced Transformer-based approaches [36, 56, 71] that perform tracking following the tracking-by-attention paradigm. Using sophisticated motion models, those approaches reach outstanding performance on several datasets, especially in short-term associations. However, especially the Transformer-based approaches require sophisticated training strategies. Contrary to all those approaches, we show that a simple linear motion model suffices to model short-term associations in most scenarios. In scenes with moving cameras or scenarios requiring non-trivial long-term associations, *e.g.*, scenarios with many occlusions, purely motion-based trackers struggle which calls for a combination with appearance-based cues.

Appearance-Based Association. To achieve better performance in long-term association scenarios, numerous approaches use additional appearance-based re-identification networks that encode appearance cues to re-identify persons after occlusions [4, 7, 25, 27, 49, 53, 63, 70]. Further exploring this direction, a recent work [41] proposed to train a detection network solely utilizing embedding information during the training process. Enhancing MOT towards real-time, several works proposed

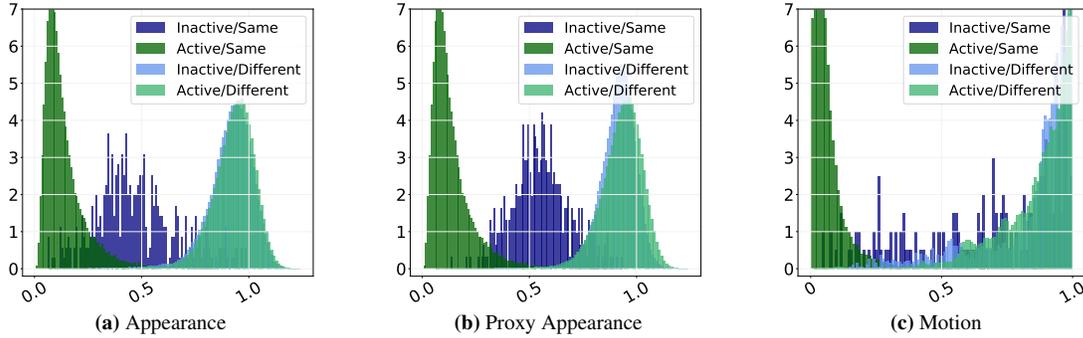


Figure 2. Distance histograms when utilizing (a) the appearance features of the last detection for active and inactive trajectories, (b) the appearance features of the last detection for active and the proxy distance for appearance features of inactive trajectories, (d) the motion distance using IoU measure for active and inactive trajectories.

to jointly compute detections and embeddings in a multi-task setting [18, 33, 62, 65, 74]. Some of them introduce more balanced training schemes [18, 65] to better leverage the synergies between both cues. While promising, approaches using appearance additionally to motion cues require rather complex association schemes with several steps [5, 74]. Also, complex and highly differing training schedules or differing inference strategies make it hard to draw conclusions about what really is driving the progress in the field. In contrast, GHOST does not rely on complex procedures but combines lightweight motion and spiced-up appearance cues in a simple yet strong TbD tracker that only requires little training data.

Person Re-Identification and Domain Adaptation. In contrast to the tracking domain, the goal of person reID is to retrieve person bounding boxes from a large gallery set that show the same person as a given query image based on appearance cues. However, state-of-the-art reID models tend to significantly drop in performance when evaluated on out-of-domain samples, *i.e.*, samples coming from other datasets [12, 76, 81]. As during application, person reID models are applied to different cameras, several approaches on cross-dataset evaluation emerged that transfer the knowledge from a given source, *i.e.*, training to a given target, *i.e.*, test domain utilizing domain adaptation (DA) [12, 76, 81]. DA often relies on adapting Batch Normalization (BN) statistics to account for distribution shifts between different domains. The weight matrix and the BN statistics store label and domain-related knowledge, respectively [30]. To update the latter, the statistics of BN layers can be updated, *e.g.*, by taking the mean and variance of all target domain images [30], by re-training using pseudo labels [9] or combining train and test dataset statistics [51]. Apart from the statistics, the learned parameters β and γ can also be updated [60]. An approach similar to ours but for classification [38] updates BN statistics during test time in a batch-wise manner. Inspired by those recent advances, we enhance our appearance model to be better suited for MOT using a simple on-the-fly domain adaptation approach. This directly adapts the

model’s learned training dataset statistics (source) to the sequences (targets).

3. Methodology

Based on the good old Hungarian TbD paradigm, GHOST combines design choices that have been overlooked by the community until now. We start by giving the general pipeline of GHOST (see Sec. 3.1) and then build our strong appearance model (see Sec. 3.2).

3.1. A simple tracking-by-detection tracker

Our tracker takes as input a set of detections $\mathcal{O} = \{o_1, \dots, o_M\}$, each represented by $o_i = (f_i, p_i)$. f_i are appearance feature vectors generated from the raw detection pixels using a Convolutional Neural Network (CNN) and p_i is the bounding box position in image coordinates. A trajectory or track is defined as a set of time-ordered detections $T_j = \{o_{j_1}, \dots, o_{j_{N_j}}\}$ where N_j is the number of detections in trajectory j . Moreover, each trajectory has a corresponding predicted position \hat{p}_j^t at time step t , produced by our linear motion model. During the tracking process, detections are assigned to trajectories. If no new detection is added to a trajectory at a given frame, we set its status to **inactive** whereas it remains **active** otherwise. We use a memory bank to keep inactive trajectories of up to 50 frames. The goal is to find the trajectories $\mathcal{T} = \{T_1, \dots, T_M\}$ that best match the detections to the underlying ground truth trajectories.

Towards that end, we associate existing detections over consecutive frames utilizing bipartite matching via the Hungarian algorithm as commonly done [5, 32, 62, 63, 74]. The assignment is driven by a cost matrix that compares new detections with the tracks already obtained in previous frames. To populate the cost matrix, we use appearance features, motion cues, or both. Our final tracker utilizes a simple weighted sum of both. We filter detection-trajectory pairs (i, j) after the matching using matching thresholds τ_i .

3.2. Strong appearance model for MOT

Our appearance model is based on ResNet50 [19] with one additional fully-connected layer at the end for downsampling, and trained on a common person reID dataset [77]. It is important to note that we *do not train any part of our reID model on any MOT dataset*. As we will show in experiments, this basic reID model does not perform well on the MOT task. We, therefore, propose two design choices to make our appearance model stronger: (i) we handle active and inactive tracks differently; (ii) we add on-the-fly domain adaptation. For this, we analyze distances between detections and tracks in given MOT sequences.

Appearance distance histograms. In Fig 2 we analyze the histograms of distances between new detections and **active** or **inactive** tracks on MOT17 validation set (please refer to Sec. 8 supplementary for details) utilizing different distance measures. In dark and light colors we show the distance of a track to a new detection of the same (positive match) and different (negative match) identity, respectively.

Different handling of active and inactive tracks. While the appearance embeddings of one identity barely vary between two consecutive frames, the embeddings of the same identity before and after occlusion can show larger distances due to, *e.g.*, partial occlusion or varying poses. This pattern can be observed in Fig 2(a) where we visualize the distance between new detections and the last detection of **active** or **inactive** tracks. The two dark-coloured histograms vary significantly, which suggests that a different treatment of **active** or **inactive** tracks is necessary. Furthermore, we can see the overlap between both negative and positive matches for **inactive** tracks showing the inherent difficulty of matching after occlusion.

Hence, for **active** tracks, we leverage the appearance features of the detection assigned to track j at frame $t - 1$ f_j^{t-1} for the distance computation to detection i at frame t $d_{i,j} = d(f_i, f_j^{t-1})$. For the **inactive** tracks we compute the distance between the appearance feature vectors of all N_k detections in the **inactive** track k and the new detection i and utilize the mean of those distances as a *proxy distance*:

$$d_{i,k} = \frac{1}{N_k} \sum_{n=1}^{N_k} d(f_i, f_k^n) \quad (1)$$

The resulting distance histograms are as visualized in Fig 2(b). This proxy distance leads to a more robust estimate of the true underlying distance between a detection and an **inactive** track. Hence, in contrast to when using a single feature vector of the **inactive** track (see Fig 2(a)) utilizing the proxy distance leads to better-separated histograms (see Fig 2(b)).

Moreover, the different histograms of **active** and **inactive** trajectories call for different handling during the bipartite matching. To be specific, thresholds typically determine

up to which cost a matching should be allowed. Looking at Fig 2(b), different thresholds τ_i divide the histograms of distances from **active** and **inactive** trajectories to detections of the same (dark colors) and different (pale colors) identities. Utilizing *different matching thresholds* τ_{act} and τ_{inact} for active and inactive trajectories allows us to keep one single matching. In contrast to cascaded matching [5], our assignment is simpler and avoids applying bipartite matching several times at each frame.

On-the-fly Domain Adaptation. As introduced in Section 2, recent developments in the field of person reID propose to apply domain adaptation (DA) techniques as the source dataset statistics may not match the target ones [12, 76, 81]. For MOT this is even more severe since each sequence follows different statistics and represents a new target domain. We, therefore, propose to apply an on-the-fly DA in order to prevent performance degradation of reID models when applied to varied MOT sequences. This allows us to capitalize on a strong reID over all sequences.

Recently, several works on person reID introduced approaches utilizing ideas from DA to achieve cross-dataset generalization by adapting normalization layers to Instance-Batch, Meta Batch, or Camera-Batch Normalization layers [12, 76, 81]. Contrary to the above-mentioned approaches, we utilize the mean and variance of the features of the current batch, which corresponds to the detections in one frame during test time, in the BN layers of our architecture:

$$\hat{x}_i = \gamma \frac{x_i - \mu_b}{\sqrt{\sigma_b + \epsilon}} + \beta \quad (2)$$

where x_i are features of sample i , μ_b and σ_b are the mean and variance of the current batch, ϵ is a small value that ensures numerical stability, and γ and β are learned during training. While not requiring any sophisticated training procedure or test time adaptations nor several BN layers, this approximates the statistics of the sequences reasonably well as all images of one sequence have highly similar underlying distributions and leads to more similar distance histograms across tracking sequences. This in turn allows us to define matching thresholds τ_i that are well suited for all sequences, *i.e.*, that separate all histograms well. For a more detailed analysis please refer to the supplementary material.

We empirically show that applying these design choices to our appearance model, makes it more robust towards occlusions and better suited for different sequences.

4. Experiments

4.1. Implementation Details

Our appearance model follows common practice [4, 7, 13], with a ResNet50 [19] model with one additional fully-connected layer at the end to downsample the

diff τ	IP	DA	MOT 17			BDD		
			HOTA \uparrow	IDF1 \uparrow	MOTA \uparrow	HOTA \uparrow	IDF1 \uparrow	MOTA \uparrow
			61.6	72.2	69.5	41.3	47.7	42.0
✓			61.8	72.7	69.6	41.9	48.6	41.9
✓	✓		62.4	73.6	69.6	42.9	50.4	43.7
✓	✓	✓	63.3	75.3	69.6	43.7	51.5	43.9

Table 1. Ablation of the single parts of our method on the validation set of CenterTrack pre-processed bounding boxes. IP=inactive proxies, DA=domain adaptation.

feature vectors. We train our model on the Market-1501 dataset [77] using label-smoothed cross-entropy loss with temperature for 70 epochs, with an initial learning rate of 0.0001, and decay the learning rate by 10 after 30 and 50 epochs. For optimization, we utilize the RAdam optimizer [31]. Moreover, we add a BN layer before the final classification layer during training and utilize class balanced sampling as in [35]. We resize the input images to 384×128 and apply random cropping as well as horizontal flipping during training [35]. Evaluated on Market-1501 dataset this model achieves 85.2 rank-1, which is far below the current state-of-the-art performance [11, 20, 35, 45, 57, 59, 75]. For tracking, we define the appearance distance between i and j as the cosine distance between appearance embeddings $d_a(i, j) = 1 - \frac{f_i \cdot f_j}{\|f_i\| \|f_j\|}$. As motion distance we use the intersection over union (IoU) between two bounding boxes $d_m(i, j) = IoU(p_i, p_j) = \frac{|p_i \cap p_j|}{|p_i \cup p_j|}$.

4.2. Datasets and Metrics

In this section we introduce the datasets we evaluate GHOST on. MOT17 and MOT20 can be evaluated in public and private detection setting. For the private detection settings, BDD and DanceTrack we use detections generated by YOLOX-X [17] following the training procedure of [73]. **MOT17.** The dataset [37] consists of seven train and test sequences of moving and static cameras. As common practice, for public detections we utilize bounding boxes refined by CenterTrack [23, 50, 79] as well as Tracktor [4, 25, 32, 43, 50, 65, 67] for MOT17. For our ablation studies, we split MOT17 train sequences along the temporal dimension and use the first half of each sequence as train and the second half as evaluation set [64, 79].

MOT20. Different from MOT17, MOT20 [16] consists of four train and test sequences being heavily crowded with over 100 pedestrians per frame. For the public setting, we utilize bounding boxes pre-processed by Tracktor [4, 7, 50].

DanceTrack. The dataset [55] significantly differs from MOT17 and MOT20 datasets in only containing videos of dancing humans having highly similar appearance, diverse motion, and extreme articulation. It contains 40, 25 and 35 videos for training, validation and testing.

BDD. The MOT datasets of BDD 100k [69] consists of 1400, 200 and 400 train, validation and test sequences with eight different classes with highly differing frequencies of the different classes. Note that our appearance model was

	HOTA \uparrow	IDF1 \uparrow	IDF1 \uparrow
BN update random patches	59.1	68.6	68.4
BN update 10 frames before	62.0	73.0	68.6
BN update whole sequence first	62.4	73.8	69.1
fine tuning on MOT17	63.0	74.3	69.6
ResNet50 IBN	63.2	74.6	69.7
our domain adaptation	63.2	75.4	69.7

Table 2. Ablation of different domain adaptation approaches.

never trained on classes other than pedestrians.

Metrics. The benchmarks provide several evaluation metrics among which HOTA metric [34], IDF1 score [48] and MOTA [24] are the most common. While MOTA metric mainly is determined by object coverage and IDF1 mostly focus on identity preservation, HOTA balances both.

4.3. Appearance Ablation

In this section, we investigate the impact of the design choices of our appearance model on tracking performance. To this end, we do not utilize motion. In Table 1, we report our results on public detection bounding boxes of MOT17 as well as on BDD. The first row shows the performance of our basic appearance model.

Different Handling of Active and Inactive Tracks. As introduced in subsection 3.2, the distance histograms for active and inactive tracks differ significantly. In the second row in Table 1, we show that utilizing different thresholds for active and inactive tracks (diff τ) improves our tracking performance by 0.5 percentage points (pp) (0.9pp) in IDF1 and 0.2pp (0.6pp) in HOTA on MOT17 (BDD). Moreover, utilizing our proxy distance (IP) computation instead of the last detection for inactive tracks further adds 0.9pp (1.8pp) in IDF1 and 0.6pp (1.0pp) in HOTA.

On-the-fly domain adaptation. Additionally, we leverage our on-the-fly domain adaptation (DA) introduced in Subsection 3.2 that accounts for differences among the sequences, allowing us to have a well-suited threshold over all sequences. We gain another 1.7pp (1.1pp) in IDF1 and 0.9pp (0.8pp) in HOTA metrics. We also compare our on-the-fly DA to various different other domain adaptation approaches (see Table 2). First, we ablate different versions of GHOST, *i.e.*, utilizing random patches of a given frame instead of the pedestrian bounding boxes, utilizing the bounding boxes of the 10 frames before the current frame as well as feeding the whole sequence first to update the parameters. Except for the last, none of them leads to a performance improvement. We argue, that random patches do not represent the statistics of pedestrians well. Also, we fine-tune our reID network on MOT17. For this, we split the train sequences into three cross-validation splits and fine-tune one model for each split. This is necessary as the same identities are given if a sequence is split along the temporal domain. While the model is fine-tuned on tracking sequences, the sequences still differ in their distributions among each other. Hence, despite fine-

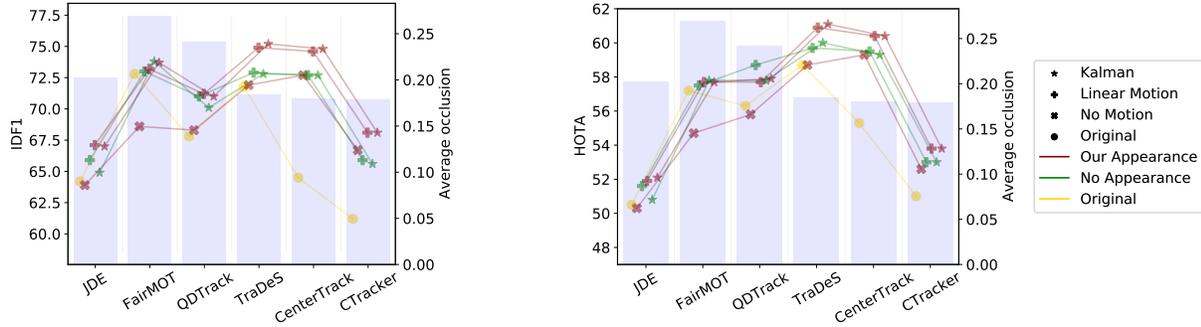


Figure 3. Analysis of contribution of appearance, motion and their combination upon private tracker bounding boxes [41, 42, 62, 64, 74, 79].

tuning also significantly improving the performance it does not surpass DA. The same holds for using Instance-Batch Norm (IBN) [39] instead of BN layers, which combine the advantages of Instance and Batch Normalization layers.

4.4. Strengths of Motion and Appearance

In this subsection, we analyze the performance of our appearance cues as introduced in Subsection 3.2 to find their strengths with respect to given tracking conditions, namely, visibility level of the detection, occlusion time of the track, and camera movement. We also analyze the complementary performance of our linear motion model that we introduce in the following. Here and in the following Subsection 4.5, we apply GHOST on the bounding boxes produced by several private trackers, treating them as raw object detections. We note that this is *not* a state-of-the-art comparison, and emphasize that we solely use those experiments for analysis and to show the potential of our insights.

Linear Motion Model. While many works apply more complex, motion models, *e.g.*, Kalman Filters [62, 74], social motion models [28], or utilize detectors as motion model [4, 79], we choose on purpose a simple linear model for our experiments. Although the world does not move with constant velocity, many short-term movements, as in the case of two consecutive frames, can be approximated with linear models and assuming constant velocity. Given detections of a track j , we compute the mean velocity v_j between the last k consecutive detections and predict the current position of a track by

$$\hat{p}_j^t = p_j^{t-1} + v_j \cdot \Delta t, \quad (3)$$

where Δt is the time difference from one frame to another and p_j^{t-1} is the position of the previous detection for active and the last predicted position for inactive tracks.

To obtain the motion distance between the new detections and the tracks, we compute the IoU between the position of a new detection p_i and \hat{p}_j^t . We also visualize the corresponding distance histogram in Fig 2(c), showing that distance histograms between detections and tracks of the same and different identities are well separated for active

tracks. In the following, we underline this observation by showing that this simple motion model is able to solve most situations. We set k to use all previous positions in Subsections 4.4 and 4.5.

Analysis Setup. For our analysis, we investigate the rate of correct associations (RCA) on MOT17 validation set [15], which we define as:

$$RCA = \frac{TP-Ass}{FP-Ass + TP-Ass}, \quad (4)$$

where TP-Ass and FP-Ass represents true positive and false positive association, respectively. We average RCA over the sets of pre-processed private detections from several trackers (see Section 4.5) to get less noisy statistics.

Observations. We visualize RCA between detections and trajectories with increasing occlusion time for different visibility levels (Fig 4) as well as the performance of motion and appearance for static and moving sequences with respect to occlusion time and visibility (Fig 5 and Fig 6). For highly visible bounding boxes (Fig 4(c)) appearance performs better than motion with respect to long- and short-term associations. While intuitively motion should perform especially well on short-term associations independent of the visibility, Fig 6 reveals that it struggles in moving sequences. This is due to the combination of the camera movement and the bounding box movement which turns motion into being more non-linear. On the other hand, in static sequences, the linear motion model performs better with respect to long-term associations than appearance (see Fig 5). This is caused by the fact that the lower the visibility (see Fig 4(a)), the higher the tendency of motion to perform better for long-term occlusions since motion is a strong cue in low visibility, *i.e.*, occluded scenarios, (see Fig 6). We show a more detailed analysis in the supplementary.

Conclusion. In conclusion, the interplay of three factors mainly influences the performance of motion and appearance: visibility, occlusion time, and camera motion. However, we saw that appearance and motion complement each other well with respect to those factors. Hence, we now move on to creating a strong tracker that combines our appearance and a simple linear motion model.

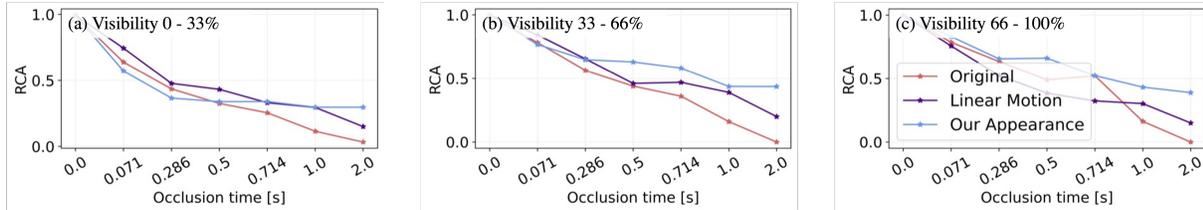


Figure 4. RCA with respect to different visibility levels for short-term vs. long-term associations.

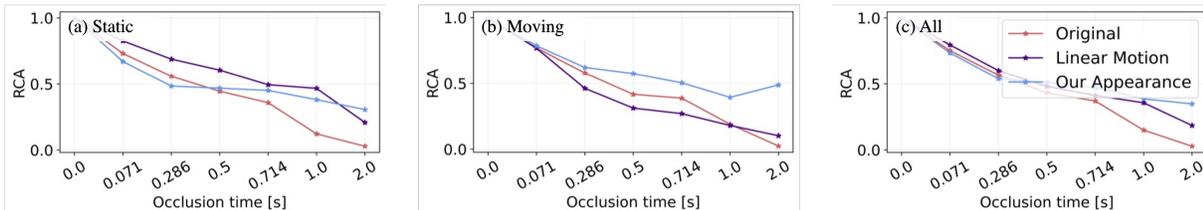


Figure 5. RCA for static, moving, and all sequences with respect to short-term vs. long-term associations.

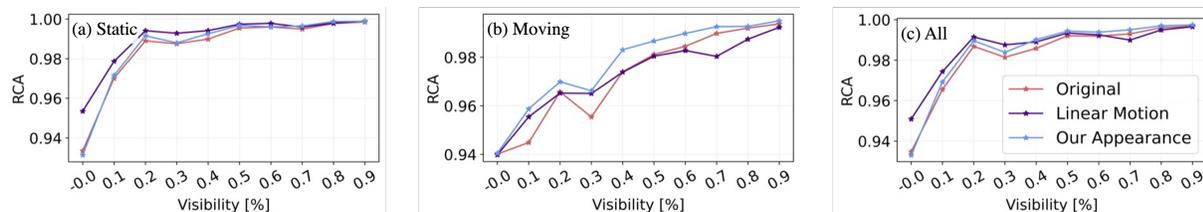


Figure 6. RCA for static, moving, and all sequences with respect to different visibility levels.

4.5. Simple cues lead to a strong tracker

In this subsection, we ablate the combination of appearance and motion into a simple Hungarian-based online tracker. We use the same setting as in Subsection 4.4, *i.e.*, we apply GHOST upon other trackers. Hence, we only report metrics related to the association performance, *i.e.*, IDF1, and HOTA in Subsection 4.5, as our goal is not to improve on the MOTA metric, which is heavily dependent on detection performance. We visualize the results in Fig 3, where markers and colors define which motion and appearance model is used, respectively. The blue bars represent the average occlusion level of detection bounding boxes of the different trackers.

Appearance. Compared to the performance of the original trackers, our appearance model improves the performance by up to 8.2pp in IDF1 and 4pp in HOTA for detection sets with lower average occlusion levels. In detection sets with high occlusion, pure appearance struggles, confirming that it is not suited for associations in those scenarios.

Motion. Interestingly, applying only the simple linear motion model without appearance always improves or performs on par with the original trackers. Utilizing a Kalman filter instead of the linear motion model does not impact the performance significantly. This further highlights the strength of the simple linear motion model. While setting the number of previous positions to use to approximate a tracks current velocity k to use all previous

positions in this section, we generally found that moving camera sequences profit from a lower k value due to the combination of the camera movement and the bounding box movement leading to less stable motion. The same holds for extreme movements, *e.g.* as in the DanceTrack dataset.

Combination. We also visualize our appearance combined with linear motion or a Kalman Filter. Although we find that sequences with a moving camera profit from a lower motion weight while detections with high occlusion level profit from a higher motion weight, we fix the motion weight to 0.5 in this experiment. Moreover, we visualize the performance of our appearance model combined with a Kalman filter. Fig 3 shows that using a Kalman filter instead of the linear motion model does not impact the performance notably. However, both combinations improve significantly over using motion or our appearance alone.

4.6. Comparison to State of the Art

We compare GHOST to current state-of-the-art approaches. In all tables **bold** represents the best results, **red** the second best, and **blue** the third best.

MOT17 Dataset. On *public detections*, GHOST improves over the best previous methods, *e.g.*, we improve over ArTIST-C [50] 1.8pp in HOTA (see Table 3). As expected, we do not improve in MOTA, as it is mostly dependent on the detection performance. In the *private detection setting* we perform on par with ByteTrack [73] in HOTA and IDF1 outperforming second-best approaches by 3.7pp and 4.6pp.

	HOTA \uparrow	IDF1 \uparrow	MOTA \uparrow	IDSW \downarrow
<i>Public MOT17</i>				
Tracktor v2 [†] [4]	44.8	55.1	56.3	1987
UNSt [2]	46.4	58.3	56.8	1914
TrackPool [†] [25]	-	60.5	55.9	1188
CenterTrack [‡] [79]	48.2	59.6	61.5	3039
ArTIST-C [‡] [50]	48.9	59.7	62.3	2062
GHOST [†]	47.4	60.6	56.5	1144
GHOST [‡]	50.7	63.5	61.6	1715
<i>Private MOT17</i>				
CenterTrack [79]	52.2	64.7	67.8	3039
TraDeS [64]	52.7	63.9	69.1	3555
QDTrack [41]	53.9	66.3	68.7	3378
FairMOT [74]	59.3	72.3	73.7	3303
MeMOT [8]	56.9	72.5	69.0	2724
GTR [80]	59.1	71.5	75.3	2859
MOTR [71]	57.8	68.6	73.4	2439
ByteTrack* [73]	62.8	77.1	78.9	2363
ByteTrack* [73]	63.1	77.3	80.3	2196
GHOST	62.8	77.1	78.7	2325
<i>Public MOT20</i>				
GMPHD [1]	35.6	43.5	44.7	7492
SORT [5]	36.1	45.1	42.7	4470
ArTIST [†] [50]	41.6	51.0	53.6	1531
Tracktor v2 [†] [4]	42.1	52.7	52.6	1648
GHOST [†]	43.4	55.3	52.7	1437
<i>Private MOT20</i>				
GSDT [61]	53.6	67.5	67.1	3230
FairMOT [74]	54.6	67.3	61.8	5243
MeMOT [8]	54.1	66.1	63.7	1938
MTrack [68]	-	69.2	63.5	6031
ByteTrack* [73]	60.4	74.5	74.2	925
ByteTrack* [73]	61.3	75.2	77.8	1223
GHOST	61.2	75.2	73.7	1264

Table 3. Comparison to state-of-the-art for public and private detections on MOT17 and MOT20. [†] and [‡] indicate bounding boxes refined by Tracktor [4] and CenterTrack [79]. * since ByteTrack uses different thresholds for different sequences of the test set and interpolation we recomputed their results without both (recomputed black / original gray).

MOT20 Dataset. Despite the complexity of the sequences being heavily crowded and exhibiting difficult lightning conditions, GHOST improves state-of-the-art [4] on *public detection setting* of MOT20 by 2.6 pp in IDF1 and 1.3 pp in HOTA metric (see Table 3). This shows that the combination of our strong appearance cues and the linear motion adapts to the underlying conditions of higher occlusion levels. In the *private detection setting* of MOT20, we outperform all other methods by up to 0.8 pp on HOTA and 0.7 pp on IDF1. GHOST is able to fully leverage the additional available bounding boxes in the private detection setting, substantially increasing identity preservation compared to the private setting.

DanceTrack Dataset. Despite the claim of the corresponding paper [55] that reID and linear motions models are not suitable for DanceTrack, GHOST outperforms all other methods on the test set (see Table 4). Surprisingly, Transformer-based methods like GTR [80] and MOTR [71] that are said to be general cannot transfer their performance from MOT17 and MOT20. We improve over prior state-of-the-art by 2.5 pp on HOTA, 3.8 pp on IDF1 and even 0.7 pp on MOTA.

	HOTA \uparrow	IDF1 \uparrow	MOTA \uparrow	DetA \uparrow	AssA \uparrow
CenterTrack [79]	41.8	35.7	86.8	78.1	22.6
FairMOT [74]	39.7	40.8	82.2	66.7	23.8
QDTrack [41]	54.2	50.4	87.7	80.1	36.8
TraDeS [64]	43.3	41.2	86.2	74.5	25.4
MOTR [71]	54.2	51.5	79.7	73.5	40.2
GTR [80]	48.0	50.3	84.7	72.5	31.9
ByteTrack [73]	47.7	53.9	89.6	71.0	32.1
GHOST	56.7	57.7	91.3	81.1	39.8

Table 4. Comparison to state-of-the-art on DanceTrack.

	mHOTA \uparrow	mIDF1 \uparrow	mMOTA \uparrow	HOTA \uparrow	IDF1 \uparrow	MOTA \uparrow
<i>validation</i>						
Yu et. al. [69]	-	44.5	25.9	-	66.8	56.9
ByteTrack [73]	45.4	54.6	45.2	61.6	70.2	68.7
QDTrack [41]	41.7	51.5	36.3	60.9	71.4	63.7
MOTR [71]	-	43.5	32.0	-	-	-
TEter [29]	-	53.3	39.1	-	-	-
GHOST	45.7	55.6	44.9	61.7	70.9	68.1
<i>test</i>						
Yu et. al. [69]	-	44.7	26.3	-	68.2	58.3
ByteTrack [73]	-	55.8	40.1	-	71.3	69.6
QDTrack [41]	-	52.3	35.5	-	72.3	64.3
GHOST	46.8	57.0	39.5	62.2	72.0	68.9

Table 5. Comparison to state-of-the-art on BDD100k.

BDD100k Dataset. Despite our appearance model *never* being trained on more than pedestrian images, it is able to generalize well to the novel classes. GHOST outperforms state-of-the-art in mHOTA and mIDF1 on the validation set by 0.3 pp and 2 pp and in mIDF1 on the test set by 1.2 pp (see Table 5). In IDF1 metrics, QDTrack [41] outperforms GHOST. Since mHOTA and mIDF1 are obtained by averaging per-class IDF1 and HOTA while IDF1 and HOTA are achieved by averaging over detections, the results show that GHOST generalizes well to less frequent classes while other approaches like QDTrack [41] overfit to more frequent classes, *e.g.*, car (see also per-class validation set results in the supplementary).

5. Conclusion

In this paper, we show that good old TbD trackers are able to generalize to various highly differing datasets incorporating domain-specific knowledge. For our general simple Hungarian tracker **GHOST**, we introduce a spiced-up appearance model that handles active and inactive trajectories differently. Moreover, it adapts itself to the test sequences by applying an on-the-fly domain adaptation. We analyze where our appearance and simple linear motion model struggle with respect to visibility, occlusion time, and camera movement. Based on this analysis, we decide to use a weighted sum that gives more weight to either cue when needed, depending on the situations in the datasets. Despite being straightforward, our insights have been largely overlooked by the tracking community. We hope to inspire future research to further investigate on, extend, and integrate these ideas into novel and more sophisticated trackers.

Acknowledgements This work was partially funded by the Sofja Kovalevskaja Award of the Humboldt Foundation.

References

- [1] Nathanael L. Baisa. Occlusion-robust online multi-object visual tracking using a GM-PHD filter with cnn-based re-identification. *J. Vis. Commun. Image Represent.*, 80:103279, 2021. 8
- [2] Favyen Bastani, Songtao He, and Sam Madden. Self-supervised multi-object tracking with cross-input consistency. *NeurIPS*, 2021. 8
- [3] Jérôme Berclaz, François Fleuret, Engin Türetken, and Pascal Fua. Multiple object tracking using k-shortest paths optimization. *TPAMI*, 33(9):1806–1819, 2011. 2
- [4] Philipp Bergmann, Tim Meinhardt, and Laura Leal-Taixé. Tracking without bells and whistles. In *ICCV*, 2019. 1, 2, 4, 5, 6, 8
- [5] Alex Bewley, ZongYuan Ge, Lionel Ott, Fabio Tozeto Ramos, and Ben Upcroft. Simple online and realtime tracking. In *ICIP*, 2016. 2, 3, 4, 8
- [6] Erik Bochinski, Volker Eiselein, and Thomas Sikora. High-speed tracking-by-detection without using image information. In *AVSS*, 2017. 2
- [7] Guillem Brasó and Laura Leal-Taixé. Learning a neural solver for multiple object tracking. In *CVPR*, 2020. 2, 4, 5
- [8] Jiarui Cai, Mingze Xu, Wei Li, Yuanjun Xiong, Wei Xia, Zhuowen Tu, and Stefano Soatto. Memot: Multi-object tracking with memory. In *CVPR*, 2022. 8
- [9] Woong-Gi Chang, Tackgeun You, Seonguk Seo, Suha Kwak, and Bohyung Han. Domain-specific batch normalization for unsupervised domain adaptation. In *CVPR*, 2019. 3
- [10] Tianlong Chen, Shaojin Ding, Jingyi Xie, Ye Yuan, Wuyang Chen, Yang Yang, Zhou Ren, and Zhangyang Wang. Abdnnet: Attentive but diverse person re-identification. In *ICCV*, 2019. 1
- [11] De Cheng, Yihong Gong, Sanping Zhou, Jinjun Wang, and Nanning Zheng. Person re-identification by multi-channel parts-based CNN with improved triplet loss function. In *CVPR*, 2016. 1, 5
- [12] Seokeon Choi, Taekyung Kim, Minki Jeong, Hyoungseob Park, and Changick Kim. Meta batch-instance normalization for generalizable person re-identification. In *CVPR*, 2021. 3, 4
- [13] Peng Dai, Renliang Weng, Wongun Choi, Changshui Zhang, Zhangping He, and Wei Ding. Learning a proposal classifier for multiple object tracking. In *CVPR*, 2021. 4
- [14] Zuozhuo Dai, Mingqiang Chen, Xiaodong Gu, Siyu Zhu, and Ping Tan. Batch dropout network for person re-identification and beyond. In *ICCV*, 2019. 1
- [15] Patrick Dendorfer, Aljosa Osep, Anton Milan, Konrad Schindler, Daniel Cremers, Ian Reid, Stefan Roth, and Laura Leal-Taixé. Motchallenge: A benchmark for single-camera multiple target tracking. *IJCV*, 129(4):845–881, 2021. 1, 6
- [16] Patrick Dendorfer, Hamid Rezaatofghi, Anton Milan, Javen Shi, Daniel Cremers, Ian D. Reid, Stefan Roth, Konrad Schindler, and Laura Leal-Taixé. MOT20: A benchmark for multi object tracking in crowded scenes. *arXiv*, abs/2003.09003, 2020. 5
- [17] Zheng Ge, Songtao Liu, Feng Wang, Zeming Li, and Jian Sun. YOLOX: exceeding YOLO series in 2021. *arXiv*, abs/2107.08430, 2021. 5
- [18] Song Guo, Jingya Wang, Xinchao Wang, and Dacheng Tao. Online multiple object tracking with cross-task synergy. In *CVPR*, 2021. 3
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 4
- [20] Alexander Hermans, Lucas Beyer, and Bastian Leibe. In defense of the triplet loss for person re-identification. *arXiv*, abs/1703.07737, 2017. 1, 5
- [21] Andrea Hornáková, Roberto Henschel, Bodo Rosenhahn, and Paul Swoboda. Lifted disjoint paths with application in multiple object tracking. In *ICML*, 2020. 2
- [22] Hao Jiang, Sidney S. Fels, and James J. Little. A linear programming approach for multiple object tracking. In *CVPR*, 2007. 2
- [23] Shyamgopal Karthik, Ameya Prabhu, and Vineet Gandhi. Simple unsupervised multi-object tracking. *arXiv*, abs/2006.02609, 2020. 5
- [24] Rangachar Kasturi, Dmitry B. Goldgof, Padmanabhan Soundararajan, Vasant Manohar, John S. Garofolo, Rachel Bowers, Matthew Boonstra, Valentina N. Korzhova, and Jing Zhang. Framework for performance evaluation of face, text, and vehicle detection and tracking in video: Data, metrics, and protocol. *TPAMI*, 2009. 5
- [25] Chanhoo Kim, Fuxin Li, Mazen Alotaibi, and James M. Rehg. Discriminative appearance modeling with multi-track pooling for real-time multi-object tracking. In *CVPR*, 2021. 2, 5, 8
- [26] H. W. Kuhn and Bryn Yaw. The hungarian method for the assignment problem. *Naval Res. Logist. Quart.*, pages 83–97, 1955. 1
- [27] Laura Leal-Taixé, Cristian Canton-Ferrer, and Konrad Schindler. Learning by tracking: Siamese CNN for robust target association. In *CVPR*, 2016. 2
- [28] Laura Leal-Taixé, Gerard Pons-Moll, and Bodo Rosenhahn. Everybody needs somebody: Modeling social and grouping behavior on a linear programming multiple people tracker. In *ICCV*, 2011. 2, 6
- [29] Siyuan Li, Martin Danelljan, Henghui Ding, Thomas E. Huang, and Fisher Yu. Tracking every thing in the wild. In *ECCV*, 2022. 8
- [30] Yanghao Li, Naiyan Wang, Jianping Shi, Jiaying Liu, and Xiaodi Hou. Revisiting batch normalization for practical domain adaptation. In *ICLR*, 2017. 3
- [31] Liyuan Liu, Haoming Jiang, Pengcheng He, Weizhu Chen, Xiaodong Liu, Jianfeng Gao, and Jiawei Han. On the variance of the adaptive learning rate and beyond. In *ICLR*, 2020. 5
- [32] Qiankun Liu, Qi Chu, Bin Liu, and Nenghai Yu. GSM: graph similarity model for multi-object tracking. In *IJCAI*, 2020. 3, 5
- [33] Zhichao Lu, Vivek Rathod, Ronny Votel, and Jonathan Huang. Retinatrack: Online single stage joint detection and tracking. In *CVPR*, 2020. 3

- [34] Jonathon Luiten, Aljosa Osep, Patrick Dendorfer, Philip H. S. Torr, Andreas Geiger, Laura Leal-Taixé, and Bastian Leibe. HOTA: A higher order metric for evaluating multi-object tracking. *IJCV*, 2021. 5
- [35] Hao Luo, Youzhi Gu, Xingyu Liao, Shenqi Lai, and Wei Jiang. Bag of tricks and a strong baseline for deep person re-identification. In *CVPR*, 2019. 1, 5
- [36] Tim Meinhardt, Alexander Kirillov, Laura Leal-Taixé, and Christoph Feichtenhofer. Trackformer: Multi-object tracking with transformers. *arXiv*, abs/2101.02702, 2021. 1, 2
- [37] Anton Milan, Laura Leal-Taixé, Ian D. Reid, Stefan Roth, and Konrad Schindler. MOT16: A benchmark for multi-object tracking. *arXiv*, abs/1603.00831, 2016. 5
- [38] Zachary Nado, Shreyas Padhy, D. Sculley, Alexander D’Amour, Balaji Lakshminarayanan, and Jasper Snoek. Evaluating prediction-time batch normalization for robustness under covariate shift. *CoRR*, 2020. 3
- [39] Xingang Pan, Ping Luo, Jianping Shi, and Xiaoou Tang. Two at once: Enhancing learning and generalization capacities via ibn-net. In *ECCV*, 2018. 6
- [40] 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 *CVPR*, 2020. 2
- [41] Jiangmiao Pang, Linlu Qiu, Xia Li, Haofeng Chen, Qi Li, Trevor Darrell, and Fisher Yu. Quasi-dense similarity learning for multiple object tracking. In *CVPR*, 2021. 2, 6, 8
- [42] 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 *ECCV*, 2020. 1, 2, 6
- [43] Jinlong Peng, Tao Wang, Weiyao Lin, Jian Wang, John See, Shilei Wen, and Errui Ding. TPM: multiple object tracking with tracklet-plane matching. *Pattern Recognition*, 2020. 5
- [44] Hamed Pirsiavash, Deva Ramanan, and Charless C. Fowlkes. Globally-optimal greedy algorithms for tracking a variable number of objects. In *CVPR*, 2011. 2
- [45] Ruijie Quan, Xuanyi Dong, Yu Wu, Linchao Zhu, and Yi Yang. Auto-reid: Searching for a part-aware convnet for person re-identification. In *ICCV*, 2019. 1, 5
- [46] Joseph Redmon and Ali Farhadi. YOLO9000: better, faster, stronger. In *CVPR*, 2017. 2
- [47] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. Faster R-CNN: towards real-time object detection with region proposal networks. In *NeurIPS*, 2015. 1, 2
- [48] Ergys Ristani, Francesco Solera, Roger S. Zou, Rita Cucchiara, and Carlo Tomasi. Performance measures and a data set for multi-target, multi-camera tracking. In *ECCV Workshops*, 2016. 5
- [49] Amir Sadeghian, Alexandre Alahi, and Silvio Savarese. Tracking the untrackable: Learning to track multiple cues with long-term dependencies. In *ICCV*, 2017. 2
- [50] Fatemeh Sadat Saleh, Sadegh Aliakbarian, Hamid Rezatofighi, Mathieu Salzmann, and Stephen Gould. Probabilistic tracklet scoring and inpainting for multiple object tracking. In *CVPR*, 2021. 5, 7, 8
- [51] Steffen Schneider, Evgenia Rusak, Luisa Eck, Oliver Bringmann, Wieland Brendel, and Matthias Bethge. Improving robustness against common corruptions by covariate shift adaptation. In *NeurIPS*, 2020. 3
- [52] 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*, abs/1805.00123, 2018. 1
- [53] Jeany Son, Mooyeol Baek, Minsu Cho, and Bohyung Han. Multi-object tracking with quadruplet convolutional neural networks. In *CVPR*, 2017. 2
- [54] Daniel Stadler and Jürgen Beyerer. Improving multiple pedestrian tracking by track management and occlusion handling. In *CVPR*, 2021. 1
- [55] Peize Sun, Jinkun Cao, Yi Jiang, Zehuan Yuan, Song Bai, Kris Kitani, and Ping Luo. Dancetrack: Multi-object tracking in uniform appearance and diverse motion. In *CVPR*, 2022. 5, 8
- [56] Peize Sun, Yi Jiang, Rufeng Zhang, Enze Xie, Jinkun Cao, Xinting Hu, Tao Kong, Zehuan Yuan, Changhu Wang, and Ping Luo. Transtrack: Multiple-object tracking with transformer. *arXiv*, abs/2012.15460, 2020. 1, 2
- [57] 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 *ECCV*, 2018. 1, 5
- [58] Pavel Tokmakov, Jie Li, Wolfram Burgard, and Adrien Gaidon. Learning to track with object permanence. *ICCV*, 2021. 2
- [59] Rahul Rama Varior, Mrinal Haloi, and Gang Wang. Gated siamese convolutional neural network architecture for human re-identification. In *ECCV*, 2016. 1, 5
- [60] Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno A. Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. In *ICLR*, 2021. 3
- [61] Yongxin Wang, Kris Kitani, and Xinshuo Weng. Joint object detection and multi-object tracking with graph neural networks. In *ICRA*, 2021. 8
- [62] Zhongdao Wang, Liang Zheng, Yixuan Liu, Yali Li, and Shengjin Wang. Towards real-time multi-object tracking. In *ECCV*, 2020. 2, 3, 6
- [63] Nicolai Wojke, Alex Bewley, and Dietrich Paulus. Simple online and realtime tracking with a deep association metric. In *ICIP*, 2017. 2, 3
- [64] Jialian Wu, Jiale Cao, Liangchen Song, Yu Wang, Ming Yang, and Junsong Yuan. Track to detect and segment: An online multi-object tracker. In *CVPR*, 2021. 1, 5, 6, 8
- [65] Yihong Xu, Aljosa Osep, Yutong Ban, Radu Horaud, Laura Leal-Taixé, and Xavier Alameda-Pineda. How to train your deep multi-object tracker. In *CVPR*, 2020. 1, 2, 3, 5
- [66] 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 *CVPR*, 2016. 2
- [67] Junbo Yin, Wenguan Wang, Qinghao Meng, Ruigang Yang, and Jianbing Shen. A unified object motion and affinity model for online multi-object tracking. In *CVPR*, 2020. 5

- [68] En Yu, Zhuoling Li, and Shoudong Han. Towards discriminative representation: Multi-view trajectory contrastive learning for online multi-object tracking. In *CVPR*, 2022. 8
- [69] Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, and Trevor Darrell. BDD100K: A diverse driving dataset for heterogeneous multitask learning. In *CVPR*, 2020. 5, 8
- [70] 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 *ECCV*, 2016. 2
- [71] Fangao Zeng, Bin Dong, Tiancai Wang, Cheng Chen, Xiangyu Zhang, and Yichen Wei. MOTR: end-to-end multiple-object tracking with transformer. In *ECCV*, 2022. 1, 2, 8
- [72] Li Zhang, Yuan Li, and Ramakant Nevatia. Global data association for multi-object tracking using network flows. In *CVPR*, 2008. 2
- [73] Yifu Zhang, Peize Sun, Yi Jiang, Dongdong Yu, Zehuan Yuan, Ping Luo, Wenyu Liu, and Xinggang Wang. Bytetrack: Multi-object tracking by associating every detection box. In *ECCV*, 2022. 5, 7, 8
- [74] Yifu Zhang, Chunyu Wang, Xinggang Wang, Wenjun Zeng, and Wenyu Liu. Fairmot: On the fairness of detection and re-identification in multiple object tracking. *IJCV*, 129(11):3069–3087, 2021. 1, 2, 3, 6, 8
- [75] Zhizheng Zhang, Cuiling Lan, Wenjun Zeng, and Zhibo Chen. Densely semantically aligned person re-identification. In *CVPR*, 2019. 1, 5
- [76] Yuyang Zhao, Zhun Zhong, Fengxiang Yang, Zhiming Luo, Yaojin Lin, Shaozi Li, and Nicu Sebe. Learning to generalize unseen domains via memory-based multi-source meta-learning for person re-identification. In *CVPR*, 2021. 3, 4
- [77] Liang Zheng, Liyue Shen, Lu Tian, Shengjin Wang, Jingdong Wang, and Qi Tian. Scalable person re-identification: A benchmark. In *ICCV*, 2015. 1, 4, 5
- [78] Kaiyang Zhou and Tao Xiang. Torchreid: A library for deep learning person re-identification in pytorch. *arXiv*, arXiv:1910.10093, 2019. 1
- [79] Xingyi Zhou, Vladlen Koltun, and Philipp Krähenbühl. Tracking objects as points. In *ECCV*, 2020. 1, 2, 5, 6, 8
- [80] Xingyi Zhou, Tianwei Yin, Vladlen Koltun, and Philipp Krähenbühl. Global tracking transformers. In *CVPR*, 2022. 8
- [81] Zijie Zhuang, Longhui Wei, Lingxi Xie, Tianyu Zhang, Hengheng Zhang, Haozhe Wu, Haizhou Ai, and Qi Tian. Rethinking the distribution gap of person re-identification with camera-based batch normalization. In *ECCV*, 2020. 3, 4