

Self-Supervised Multi-Object Tracking with Path Consistency

Zijia Lu*, Bing Shuai†, Yanbei Chen, Zhenlin Xu, Davide Modolo
AWS AI Labs

lu.zij@northeastern.edu, {bshuai, yanbec, xzhenlin, dmodolo}@amazon.com

Abstract

In this paper, we propose a novel concept of path consistency to learn robust object matching without using manual object identity supervision. Our key idea is that, to track an object through frames, we can obtain multiple different association results from a model by varying the frames it can observe, i.e., skipping frames in observation. As the differences in observations do not alter the identities of objects, the obtained association results should be consistent. Based on this rationale, we generate multiple observation paths, each specifying a different set of frames to be skipped, and formulate the Path Consistency Loss that enforces the association results are consistent across different observation paths. We use the proposed loss to train our object matching model with only self-supervision. By extensive experiments on three tracking datasets (MOT17, PersonPath22, KITTI), we demonstrate that our method outperforms existing unsupervised methods with consistent margins on various evaluation metrics, and even achieves performance close to supervised methods.

1. Introduction

Multi-Object Tracking (MOT) is the task of identifying and tracking all object instances present in a video. It benefits a wide spectrum of applications, such as monitoring vehicle traffic, wildlife activity, pedestrians, etc. While there have been several advanced supervised methods developed for MOT [5, 8, 29, 33, 34, 46, 52, 55, 58], the overall advancement in MOT is slower compared to other vision tasks, mostly due to the prohibitively expensive cost of acquiring high-quality object identities annotations for MOT training. As of today, MOT17 [31], a dataset of 14 videos, remains one of the most widely used datasets in MOT.

To overcome this problem, some recent works [1, 16, 21, 22, 24, 30] started shifting their interest on unsupervised learning for MOT. These approaches leverage off-the-shelf object detectors to localize all the object instances in each

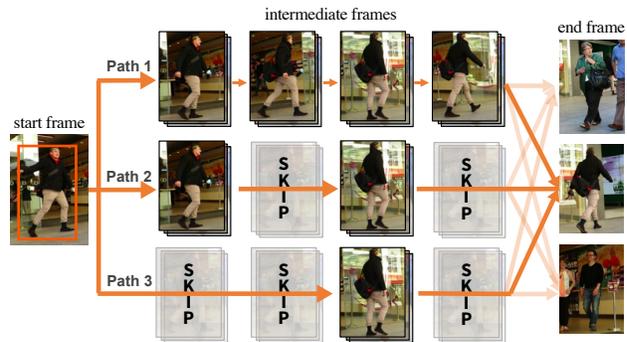


Figure 1. We propose a novel concept of path consistency for self-supervised MOT. We define an observation path for an object as a temporal list of observed frames from the start to end frame. As such, for the same object, we can generate multiple paths by skipping intermediate frames. As different observations of the same object does not alter its identity, the association results should be consistent across different paths.

frame of a video and then learn the temporal object association across frames in an unsupervised manner. Multiple works [16, 21, 22] propose to learn the object associations by creating pseudo-tracklets to estimate object IDs. However, they are vulnerable to the noises in tracklets, especially when the groundtruth tracklet of an object is incorrectly broken into several pseudo-tracklets, each being given a different object ID from the others. To avoid relying on pseudo-labels, other works [24, 30] explore self-supervision with temporal consistency objective, by tracking objects from the current frame to a future frame then back to the current frame while requiring each object is associated back to itself. However, these works only consider future frames in very short temporal distances (e.g., 1-2 frames) where objects lack substantial appearance difference, thus limiting the quality of their object embeddings. Meanwhile, UNS [1] leverages a motion model and an appearance model to create cross-supervision between their tracking outputs, yet it also only considers matching objects among adjacent frames.

In fact, the ability of robust object matching over long

*Currently at Northeastern University. Work conducted during an internship with AWS AI Labs.

†Corresponding author.

temporal distances is critical for MOT models as objects often encounter occlusion during tracking, where their bounding boxes are lost for a potentially long period of time, thus cannot be tracked by mere frame-by-frame matching. It requires long-distance matching to re-associate their bounding boxes after occlusion. However, very few works have managed to propose reliable solutions that tackle occlusion without using groundtruth supervision.

In this paper, we introduce the novel concept of path consistency, which can be used as a reliable self-supervised signal to learn a robust object matching model. To elaborate, when tracking an object, we define an observation path of it as a temporal list of observed frames between a start frame and an end frame. As shown in Figure 1, we are able to generate multiple different observation paths for the same object by randomly skipping intermediate frames. Based on an observation path, we can compute the temporal association between the object and all other objects from the start frame to the end frame. As such, path consistency dictates that the temporal associations should be consistent across paths, since the identities of objects do not change in different paths. Based on this rationale, we propose the *Path Consistency Loss* (PCL) that estimates the associations of objects along different paths and trains the object matching model by minimizing the discrepancy among association probability derived from different paths.

Distinct from prior self-supervised methods [1, 24, 30], our formulation learns object matching over both *short and long distances*. As paths contain frame skipping of varied lengths, the model trained with PCL is able to learn not only short-distance matching (from paths with minimal frame skipping), but also long-distance matching (from paths with consecutive frame skipping) which is essential to handle occlusion during inference. In addition, we sample challenging training examples where paths span long temporal horizons, hence enabling our path consistency formulation to learn robust object matching models and handle occlusion.

We evaluate our method on three benchmark tracking datasets, MOT17 [31], PersonPath22 [39] and KITTI [13], and show it successfully achieves new state-of-the-art results over prior unsupervised methods and performs even close to recent supervised methods. We also provide extensive ablation studies to validate the design of our framework and improvement in handling occlusion. Our contributions are summarized as follows:

- We propose the novel concept of path consistency that can be used as a reliable self-supervised signal to train a robust object matching model. Based on the concept, we formulate the Path Consistency Loss (PCL) that is capable of learning robust long-distance matching, which is essential in handling occlusion and large appearance changes in tracking.
- We establish new state-of-the-art performance over prior

unsupervised methods on three challenging tracking benchmarks. Extensive ablation studies also validate our improvement in long-distance matching and tracking over occlusion.

2. Related Work

Supervised multi-object tracking. Most supervised MOT methods [2, 5, 15, 17, 19, 25–29, 36–38, 43, 44, 44, 50, 51, 53, 55, 58, 60] follow the *tracking-by-detection paradigm* where all object instances are first detected in each frame, then they are linked based on their similarities across frames to form tracklets. This paradigm benefits from the advances in object detection [6, 14, 35, 45, 57, 59] and has two threads of works: offline and online MOT. Offline MOT methods [2, 15, 17, 36, 37, 44, 51, 53] focus on finding global optimal associations of objects and are computationally heavy in general. Online MOT methods [3, 5, 29, 32, 38, 48, 52, 55, 56, 58, 60] focus on optimizing local optimum on associating detected objects and have real-time speed, e.g., Tracktor [3] uses bounding box regression to estimate new positions of objects. Meanwhile, the joint learning of the detection and object matching has made rapid progress in performance [48, 55], e.g., [5, 29, 52] introduce transformers and object tokens to associate objects with existing tracklets or to identify new ones. Despite the efficacy of supervised methods, they require costly human-annotated labels [10, 13, 31, 39]. In this work, we learn an online object matching model without labels, which we believe is an important step towards mitigating the limited size of tracking datasets.

Unsupervised multi-object tracking. This task aims to alleviate the high annotation cost of tracking labels. While existing motion trackers [4, 7, 56] do not require labels for training, they cannot leverage the appearance information of objects in tracking, thus often lose objects in the event of movement change or long occlusion. To leverage appearance information, recent unsupervised methods focused on developing new object matching models [1, 22, 24, 30]. [20] proposes a setting of open-vocabulary MOT.

Among recent works, pseudo-label methods [16, 22] generate pseudo-tracklets to train their object matching models. SimpleReID [16] and UEANet [21] obtain pseudo-tracklets from motion trackers without detecting tracklet errors, thus those errors propagate to their models. U2MOT [22] improves the consistency of object IDs within tracklets, yet is still vulnerable to errors in the early stage of training and, more importantly, errors between tracklets. Specifically, the groundtruth tracklet of an object is often broken into several pseudo-tracklets if it experiences occlusions, where each pseudo-tracklet is then given a different object ID, treated as negative samples of each other. On the other hand, self-supervised methods [1, 24, 30] derive new

unsupervised training objective to replace pseudo-labels. OUTrack [24] and UCSL [30] utilize temporal consistency, i.e., an object tracked from current frame to a future frame then back to current frame should associate back to itself. However, this objective does not hold if objects disappear among the frames, which limits those methods to only allow future frames at near distances (1-2 frames) to ensure objects likely remain present. Unfortunately, objects lack substantial appearance changes in such short distances, causing their failure to handle occlusion by matching objects over long temporal distances. We show that object matching learned with short distances cannot generalize to long distances (Table 5). UNS [1] creates cross-supervision between motion and appearance trackers yet also learn short-distance matching. It discusses occlusion-based supervision yet has failed and discarded it, likely due to its focus on short-distance matching. Meanwhile, PointId [41] learns a person embedding model by data augmentation, but it requires pretraining on ReID datasets with groundtruth labels.

Unlike existing methods, we propose a novel self-supervised path consistency loss that simultaneously learns object matching over short to long distances without requiring pseudo or groundtruth labels. Thus we are able to consistently track objects in the event of long occlusion.

3. Methodology

3.1. Framework Overview

We tackle the unsupervised MOT task with the concept of Path Consistency, which trains a robust object matching model capable of associating objects over long temporal distances, enabling consistent tracking in the event of occlusion. Following existing unsupervised works [1, 37, 44], we assume the bounding boxes of objects in each video frame are provided by an off-the-shelf detector. Our object matching model takes the detected objects and estimate object embeddings, which encode the appearance and spatial location of each object. Specifically, let o_i^t denote the i -th object in frame t and h_i^t its embedding. We define the probability that o_i^t matches to an object o_j^r in another frame r as

$$p(o_i^t \rightarrow o_j^r) = \frac{\exp(h_i^t \cdot h_j^r)}{\sum_{u=1}^{N_r} \exp(h_i^t \cdot h_u^r)}, \quad (1)$$

where N_r is the total number of objects in frame r . Similarly, we also compute $p(o_j^r \rightarrow o_i^t)$, the matching probability from o_j^r to o_i^t . Importantly, to deal with occlusions and object exit in object association, we add a special *null* object, ϕ , to the list of localized objects in each frame. If a real object o_i^t is no longer visible in the following frame r , it gets matched to ϕ .

During inference, our method uses the learned probability p to associate objects to existing tracklets. Given a

tracklet $\tau = \{o_m^{t_m}\}_{m=1}^M$, consisting of the latest M objects, it estimates the similarity between τ and a new object o_i^t as

$$\frac{1}{M} \sum_m \text{gm} [p(o_m^{t_m} \rightarrow o_i^t), p(o_i^t \rightarrow o_m^{t_m})], \quad (2)$$

in which, gm denotes the geometric mean operator. Similar to all other works [22, 24, 30, 55, 56], we perform bipartite matching based on the tracklet-object similarities to either extend / terminate existing tracklets, or initiate new ones.

In the absence of tracking annotations, we propose a novel path consistency loss to learn the object matching model over *varied temporal distances with only self-supervision* (Section 3.2) and two regularization losses to improve β convergence (Section 3.3). As our method is agnostic to model structure and inference procedure, we provide their details in the supplementary materials.

3.2. Self Supervision with Path Consistency

In this section, we elaborate the concept of path consistency as a reliable self-supervision signal to train a robust object matching model. Technically, suppose that we have a query object $o_i^{t_s}$ in the start frame t_s , we aim to associating it to its correspondent object, $o_j^{t_e}$, in the end frame t_e ($t_e > t_s$). Note that the corresponding object $o_j^{t_e}$ is the *null* object in the case that the query object is not visible in frame t_e (due to being occluded or exiting the scene).

We define an *observation path* as the list of frames a model can observe to compute the association probability distribution between $o_i^{t_s}$ and all visible objects in frame t_e . We are able to construct multiple paths by randomly skipping frames, thus creating a different visual observation for each path. Take the example in Figure 2 (right) as illustration, all frames are observed in Path1 whereas only a subset of frames are observed in Path2 and Path3. The core principle of the path consistency is, despite the different observations in the paths, the identities of objects remain unchanged, thus the association probability between $o_i^{t_s}$ and $o_j^{t_e}$ should always be higher than all other alternatives. That is, the association probability distribution between $o_i^{t_s}$ and all visible objects in frame t_e should be consistent across paths. Based on this principle, we introduce our path consistency loss in Section 3.2.1, and then present our method to sample frames (t_s, t_e) in Section 3.2.2.

3.2.1 Path Consistency Loss

Path Consistency Loss (PCL) computes the association probability distributions from all possible paths, $\{q_\pi | \pi \sim \Pi\}$, and enforces consistency among them¹, where Π is the set of all paths and q_π is the association probability distribution obtained using a specific path

¹When the number of paths is large, we conduct Monte Carlo estimation by computing the loss with a smaller set of sampled paths.

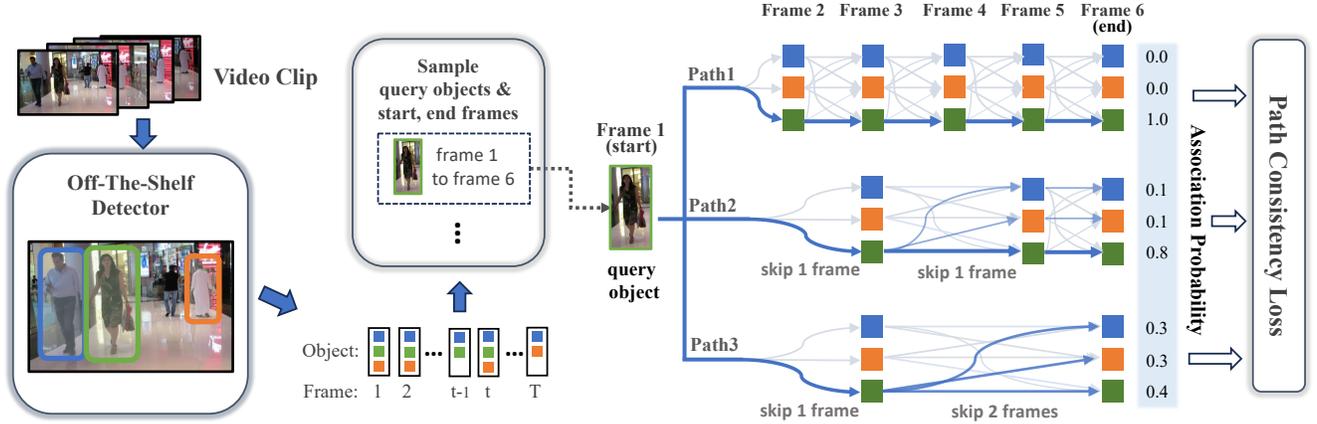


Figure 2. Overview of Path Consistency Loss (PCL). Our method takes a video clip as input, where objects are localized by an off-the-shelf detector, and uses a selection strategy to choose suitable query objects and their correspondent end frames, then computes PCL to learn association between query objects and objects in end frames. Association probabilities obtained from different paths provide cross-supervision among them and enables self-supervised model learning.

$\pi = \{t_s, t_1, \dots, t_K, t_e\}$. t_k is an intermediate frame in the path. We use $q_\pi^{t_k}(j)$ to denote the probability that query object $o_i^{t_s}$ matches with the j -th object at frame t_k .

Computing Association. To obtain association probability q_π , we aggregate the matching probability p defined in Eq. 1 according to the path π . We have

$$q_\pi^{t_k}(j) = \begin{cases} p(o_i^{t_s} \rightarrow o_j^{t_1}), & \text{if } t_k = t_1 \\ \sum_u q_\pi^{t_{k-1}}(u) \cdot p(o_u^{t_{k-1}} \rightarrow o_j^{t_k}), & \text{otherwise} \end{cases} \quad (3)$$

where the first case defines the association from query object to objects in the first intermediate frame t_1 , which is the matching probability, p , between the objects. The second case extends the association from a frame t_{k-1} to t_k by jointly considering the matching probability from *all objects* in t_{k-1} to $o_j^{t_k}$. In contrast, pseudo-label methods [16, 21, 22] only consider the most likely matching, leading to error accumulation. Overall, we have the association probability between query object and objects in end frame as

$$q_\pi^{t_e}(j) = \sum_u p(o_i^{t_s} \rightarrow o_u^{t_1}) \cdots \sum_v p(o_v^{t_K} \rightarrow o_j^{t_e}). \quad (4)$$

which cumulatively multiplies the matching probabilities according to the frames in path.

To reduce the space of possible associations and facilitate learning, we constrain that only two spatially close objects between two neighboring frames can be matched. We define a binary spatial constraint mask $C_{t_k, v}^{t_{k-1}, u}$, which equals to zero if $o_v^{t_k}$ is one of the S furthest objects to $o_u^{t_{k-1}}$.

Thus, we update $q_\pi^{t_e}$ as

$$\frac{1}{Z} \sum_u (C_{t_1, u}^{t_s, i}) p(o_i^{t_s} \rightarrow o_u^{t_1}) \cdots \sum_v (C_{t_e, j}^{t_K, v}) p(o_v^{t_K} \rightarrow o_j^{t_e}), \quad (5)$$

where Z is a normalization factor to ensure the probability $q_\pi^{t_e}$ sums to 1.

Computing Loss. After computing the association probabilities for all paths $\{q_\pi^{t_e} | \pi \sim \Pi\}$, We define the path consistency loss to enforce consistency among them to learn our model with self-supervision,

$$\mathcal{L}_{PC}(o_i^{t_s}, t_e) = \frac{1}{|\Pi|} \sum_\pi KL(q_\pi^{t_e} || \hat{q}) + H(q_\pi^{t_e}), \quad (6)$$

which includes two terms. The first term is the KL Divergence between each association probability $q_\pi^{t_e}$ and the average of all probabilities, $\hat{q} = \sum_\pi q_\pi^{t_e} / |\Pi|$. Thus, it is minimized only when *all probabilities follow the same distribution*. The second term minimizes the entropy, $H(\cdot)$, of each probability to require only one object in frame t_e can be associated to the query object, thus avoids the trivial solution where all $q_\pi^{t_e}$ are consistent yet being uniform distributions.

Our path consistency loss has the key advantages of (i) linear complexity w.r.t the number of paths as we compare each $q_\pi^{t_e}$ to the averaged probability \hat{q} , hence avoid computing pairwise discrepancy between all $q_\pi^{t_e}$. (ii) \hat{q} reflects association consistency and modulates the loss. When the consistency is low, multiple objects can have high probabilities in \hat{q} , indicating different association results. Unlike pseudo-label methods [16, 21, 22] that will enforce a specific association, PCL allows exploring all potential associations by using \hat{q} in $KL(\cdot)$, while \hat{q} also gives higher weights to the more likely ones.

3.2.2 Selecting Frame Pairs to compute PCL

In principle, based on our PCL formulation we can sample any two random frames (t_s, t_e) from a video clip, and then use all objects in t_s as query objects. However, we observe such unconstrained frame sampling leads to most query objects not visible in the end frames. The unbalance between visible and disappeared query objects causes the model to always matching objects to the *null* objects and makes the training hard to converge.

To constrain the frame sampling, we pick frame pairs (t_s, t_e) under the condition that one object in t_s should remain visible in intermediate and end frames to serve as the query object while t_s and t_e should be temporally disclose to support learning long-distance matching. As such, we develop a simple heuristics that uses the temporal chain of the IOU similarity to reject most frame pairs. With the existing solution, we observe that the query objects are present in end frames 98% of all sampled pairs. We include the details of our approach in the supplementary details. Notice that we can also employ more advanced approaches to include more challenging training examples and reject more noisy frame pairs, which we will leave for future works.

3.3. Regularization Losses

We further introduce the following two losses to regularize the model training together with our path consistency loss.

One-to-one matching loss. To avoid that the path for two different query object cross each other, we constrain that a real, non-*null* object in frame r , o_j^r , can not be matched by more than one real objects in frame t . As elaborated earlier, the probability $p(o_i^t \rightarrow o_j^r)$ will be close to 1 between two matched real objects o_i^t and o_j^r . As such, we compute $\sum_i p(o_i^t \rightarrow o_j^r)$ to approximate the number of real objects in frame t that are matched to object o_j^r and penalize the value if it exceeds 1. Mathematically, our loss is defined as:

$$\mathcal{L}_{OM} = \frac{1}{T^2} \sum_{t,r} \frac{1}{N^r} \sum_j \max \left(1, \sum_i p(o_i^t \rightarrow o_j^r) \right), \quad (7)$$

which penalizes any many-to-one real object matching between every two frames (t, r) in the input video clip of length T . N^r is the number of real objects in frame r .

Bidirectional consistency loss. In our model, we compute matching probability from frame t to a later frame r (forward in time) but also probability from frame r to frame t (backward in time). To ensure the matching probabilities are invariant to the temporal directions, we design a bidirectional consistency loss as:

$$\mathcal{L}_{BC} = \frac{1}{T^2 N_t N_r} \sum_{t,r} \sum_{i,j} ||p(o_i^t \rightarrow o_j^r) - p(o_j^r \rightarrow o_i^t)||^2. \quad (8)$$

Given Eq. 7 and Eq. 8, our overall model objective is the sum of path consistency loss and two regularization losses: $\mathcal{L} = \mathcal{L}_{PC} + \mathcal{L}_{OM} + \mathcal{L}_{BC}$, where we combine all loss terms with a balance weight of 1.0.

4. Experiments

4.1. Experimental setup

Datasets. We conduct experiments on three tracking datasets: **MOT17** [31] is a widely adopted person tracking dataset containing 14 video sequences, featuring crowded scenes in indoor shopping malls or outdoor streets. The videos are split into 7 training sequences and 7 test sequences. **PersonPath22** [39] is a large-scale pedestrian dataset with 236 videos, which are collected from varied environment settings that contain *longer occlusions and more crowded scenes*. It is split into 138 training videos and 98 test videos. **KITTI** [13] is person and car tracking dataset, featuring videos with lower frame rates. We follow UNS [1] to report performance on car tracking.

Evaluation metrics. Following the evaluation protocols of the datasets, we report our results with several metrics: HOTA (DetA, AssA, LocA), MOTA (IDsw, FP and FN) and IDF1. *We emphasize on HOTA, AssA, IDsw and IDF1 as they measure association accuracy. Other metrics are more influenced by the detection accuracy of detectors.*

Implementation details. Our model input is a video clip of $T = 48$ frames, which we found balances between performance and efficiency. For PCL, we sample at most $G = 25$ different paths for each query object and end frame and set $S = \sqrt{N}$ in spatial mask constraint, where N is average number of objects per frame. We implement our model with Pytorch and trained with Adam Optimizer on Tesla V100 and a learning rate of 0.0001. In inference, we use the latest $M = 4$ bounding boxes of each tracklet to estimate its association with new objects and maintain unmatched tracklets for 30 frames. We provide more implementation details in the supplementary materials.

4.2. Comparison to the state-of-the-art

We compare our method with recent unsupervised object matching methods, where SimpleReID [16], U2MOT [22] are *pseudo-label methods* and OTrack [24], UNS [1], UCSL [30] are *self-supervised methods*. For complete comparison, we include the results of motion-based methods [4, 7, 56], as they do not require object ID annotation, and also results of recent supervised methods.

Comparison on MOT17. In Table 1, we report the results on MOT17 dataset with both private and public detections. For fair comparison, we use YoloX detections from [56] for private detection and follow the detection pre-processing of UNS [1] for public detection. Our method outperforms

	Sup.	HOTA \uparrow	IDF1 \uparrow	MOTA \uparrow	IDsw \downarrow
Private Detection					
TransTrack [42]	✓	54.1	63.9	74.5	3663
TrackFormer [29]	✓	57.3	68.0	74.1	2829
MeMOT [5]	✓	56.9	69.0	72.5	2724
MOTR [52]	✓	57.8	68.6	73.4	2439
MeMOTR [12]	✓	58.8	71.5	72.8	1902
FairMOT [55]	✓	59.3	72.3	73.7	1074
FineTrack [34]	✓	64.3	79.5	80.0	1270
UTM [49]	✓	64.0	78.7	81.8	1431
MotionTrack [33]	✓	65.1	80.1	81.1	1140
SimpleReID [16]	×	50.4	60.7	69.0	-
OUTrack [24]	×	-	70.2	73.5	4110
PointID [40]	×	-	72.4	74.2	2748
ByteTrack [56] †	×	63.1	77.3	80.3	2196
OC-SORT [7] †	×	63.2	77.5	78.0	1950
UCSL [30]	×	58.4	70.4	73.0	-
U2MOT [22] †	×	64.2	78.2	79.9	1506
<i>Ours</i> †	×	65.0	79.6	80.9	1749
Public Detection					
Tractor++ [3] ‡	✓	42.1	52.3	53.5	2072
MHT-BLSTM [18]	✓	41.0	51.9	47.5	2069
FAMNet [9]	✓	-	48.7	52.0	3072
LSST [11]	✓	47.1	62.3	54.7	1243
GSM [23]	✓	45.7	57.8	56.4	1485
CenterTrack [58]	✓	-	60.5	55.7	2540
TrackFormer [29]	✓	-	62.3	57.6	4018
SORT [4] ‡	×	34.0	39.8	43.1	4852
UNS [1] ‡	×	46.4	58.3	56.8	1320
<i>Ours</i> ‡	×	49.0	61.2	58.8	1219

Table 1. Comparison on MOT17 [31] dataset. \uparrow/\downarrow : higher/lower is better. “Sup.” means requiring supervision. “†” denotes using YoloX detection. “‡” denotes using Tractor++ detection preprocessing. The best results of unsupervised methods are in bold.

all unsupervised works on both private and public detections, and even exceeds recent fully-supervised methods, e.g., UTM [49], FineTrack [34].

Comparing with *pseudo-label methods*, we surpass the best competitor, U2MOT [22], by 0.8% in HOTA, 1.4% in IDF1 and 1% in MOTA with the same YoloX detection, as it is vulnerable to errors in pseudo labels. In contrast, our PCL naturally measures model confidence with \hat{q} and avoids enforcing incorrect, unconfident associations. Meanwhile, we also exceed *self-supervised methods* [1, 24, 30] as they only learn short-distance matching among adjacent frames. Our method learns object matching over both short and long distances thus is robust to appearance change and occlusion. Using the same public detection as UNS [1], we surpass it by 2.6% in HOTA, 2.9% in IDF1 and 2% in MOTA. Some works [16, 24] use better detectors to improve public detection, hence their performance are not comparable to us.

Comparison on PersonPath22. In Table 2, we report results on PersonPath22 dataset, which is more challenging than MOT17 as *its videos contain more crowded scenes with long occlusions*. We compare with both public and private detections and successfully outperform ByteTrack

	Sup.	IDF1 \uparrow	MOTA \uparrow	IDsw \downarrow	FP \downarrow	FN \downarrow
CenterTrack [58]	✓	46.4	59.3	10319	24K	72K
SiamMOT [38]	✓	53.7	67.5	8942	13K	63K
FairMOT [55]	✓	61.1	61.8	5095	15K	80K
TrackFormer [29]	✓	57.1	69.7	8633	23K	47K
SORT [4]	×	39.6	53.4	9927	2K	110K
UNS [1]	×	41.8	55.5	13877	17K	86K
PointID [40]	×	63.1	68.6	6148	9K	66K
ByteTrack [56] *	×	57.2	64.6	5530	9k	71k
ByteTrack [56] †	×	66.8	75.4	5931	17k	40k
OC-SORT [7] †	×	68.2	74.0	4931	17k	41k
<i>Ours</i> *	×	60.5	64.8	5008	7k	73k
<i>Ours</i> †	×	69.4	75.8	4663	15k	38k

Table 2. Comparison on PersonPath22 [39] dataset. \uparrow/\downarrow : higher/lower is better. The best results of unsupervised methods are in bold. “*” denotes using the public detection of Personpath22 provided by the authors. “†” denotes using YoloX detections.

	Sup.	HOTA \uparrow	AssA \uparrow	DetA \uparrow	MOTA \uparrow
FAMNet [9]	✓	52.6	45.5	61.0	75.9
mmMOT [54]	✓	62.1	54.0	72.3	83.2
CenterTrack [58]	✓	73.0	71.2	75.6	88.8
LGM [47]	✓	73.1	72.3	74.6	87.6
TuSimple [8]	✓	71.6	71.1	72.6	86.3
PermaTr [46]	✓	78.0	78.4	78.3	91.3
SORT [4]	×	42.5	41.3	44.0	53.2
UNS [1]	×	62.5	65.3	61.1	-
OC-SORT [7]	×	76.5	76.4	77.2	90.8
<i>Ours</i>	×	78.8	80.3	77.9	91.0

Table 3. Comparison on KITTI [13] dataset (Tracking Car). \uparrow/\downarrow : higher/lower is better. The best results of unsupervised methods are in bold. We use the same detections as OC-SORT.

by 3.3% and 2.6% in IDF1 respectively, also reducing IDsw by large margins. The larger improvement in IDF1 demonstrates the better ability of our model in matching over long temporal distances and re-associate objects over occlusion to deliver more consistent tracking results. With even public detection, our model has achieved better performance than CenterTrack [58], TrackFormer [29] and surpasses FairMOT [55] in MOTA, despite being unsupervised. These results show that our model effectively benefits from more training data, exposing the possibility of large-scale self-supervised learning with enormous tracking videos available online.

More importantly, in Table 4, we study model performance under challenging scenarios, including different occlusion length (measured by frames), object moving speeds and people density. Our model consistently outperform the best competitor, ByteTrack, by large margins, thanks to our PCL learning more robust object matching to occlusion and appearance changes.

	Occlusion Length (frames)				
	0	1-20	20-40	40-60	60+
UNS [1]	49.2	36.8	32.8	27.0	26.1
OC-SORT [7]	77.9	61.9	50.1	39.3	39.2
Ours	79.2	64.3	50.7	40.7	39.8

	Movement Speed			People Density	
	slow	medium	fast	low	high
UNS [1]	42.7	40.3	38.3	44.7	42.4
OC-SORT [7]	65.9	63.6	61.7	69.7	64.8
Ours	68.3	65.1	62.6	71.4	66.3

Table 4. Model Performance (IDF1) under different challenging scenarios on PersonPath22.

$ t - r $	1-4	5-8	9-16	17-32	32-48
PCL ($s = T$)	99.5	99.1	98.1	96.3	94.5
$s=32$	99.5	99.2	98.2	96.3	93.0
$s=24$	99.5	99.1	98.0	94.8	88.0
$s=16$	99.6	99.1	98.0	91.4	81.7
$s=8$	99.5	99.1	95.7	84.6	71.3
$s=4$	99.5	98.0	91.8	78.3	64.1

Table 5. Accuracy of the matching probability $p(o_i^t \rightarrow o_j^r)$. Our method has learned accurate object matching over short to long temporal distances.

Comparison on KITTI. In Table 3, we further report our performance on KITTI dataset for tracking car, following the evaluation protocol of UNS [1]. KITTI features varied moving velocity of objects and lower frame rate on average, thus objects have large appearance and location changes in adjacent frames [7]. We use the shared detections from OC-SORT [7].

Our model exceeds the best competitor, OC-SORT [7], by 2.3% in HOTA, 3.9% in AssA, 0.7% in DetA and 0.2% in MOTA. This is because cars have changes in their moving speeds and bounding box aspect ratios during driving (e.g., when making turn), which leads to low overlap between bounding boxes of the same object in consecutive frames and failure of pure motion trackers [4, 7]. On the other hand, our method learns object embeddings to encode object appearance thus successfully links cars in various driving scenarios, further validating our learned object matching is robust to appearance and location changes.

4.3. Ablation Study

We conduct ablation study to examine our key design choices. We train our unsupervised models on MOT17 test set and evaluate on MOT17 train set, since the annotations of the former set are not publicly available. Finally, we use the publicly released SDP detections.

Has the model learned long-distance object matching?

We exam the accuracy of our learned matching probability,

	L=0	L=10	L=20	L=30	L=40	L=50	L=60
UNS	64.0	62.5	58.0	55.0	53.5	52.0	50.0
OC-SORT	67.7	66.9	65.8	65.1	63.9	63.5	62.5
Ours	68.9	68.3	67.4	66.9	66.1	65.9	65.3
Gain	+1.2	+1.4	+1.6	+1.9	+2.2	+2.4	+2.8

Table 6. Comparison of IDF1 for tracking over occlusion. When increasing the minimal occlusion length (number of frames), our model has more robust performance and exceeds the best competitor OC-SORT by large margins (as shown in last row).

$p(o_i^t \rightarrow o_j^r)$, where p is correct if o_i^t is the correct matching of o_i^t and has the largest matching probability. In the first two rows of Table 5, we show the accuracy and break it down by the temporal distances between objects $|t - r|$. It shows our matching probability is robust to temporal distances. As $|t - r|$ increases, the accuracy drops only by small margins and remains 94.5% even at the distance of 32-48 frames, validating our model has learned object matching from short to long distances.

Can the model only learn short-distance matching?

Next, we further exam if learning long-distance matching is necessary and if a model only learned short-distance matching can generalize to long-distance matching? In rows 2-6 of Table 5, we limit the maximal number of continuous frames that can be skipped in PCL, denoted by s . A smaller s means the model is only trained to match objects over short temporal distances.

As s decreases from 32 to 4, while accuracy for short distances (1-8) remains high, it deteriorates quickly for longer distances. Comparing to the first row where we do not limit s , accuracy drops by 30.4% on the distance of 32-48. It shows that *model trained for short-distance matching cannot generalize to longer distances*, which is the drawback of previous self-supervised methods [1, 24, 30]. Our PCL has successfully learned matching probability over varied temporal distances while not requiring groundtruth annotations. We also observe that 50% of frame skipping in paths are longer than 8 frames, which boosts our model robustness for long distance matching and the appearance changes caused by it.

Can the model handle long occlusion?

After validating the accuracy of probability matching, we exam our model ability for handling occlusion at inference. We notice that most occlusions in MOT17 videos are of short lengths, e.g., less than 5 frames, which cannot reflect a model’s ability to handle long-distance matching. Hence, we increase the occlusion length in a video by keep dropping the bounding boxes of objects after occlusion, until the occlusion lengths reach L frames. If the original length of an occlusion is longer than L , we do not increase it. Thus, $L = 0$ means no occlusion is extended while $L = 60$ means all occlusions last for at least 60 frames.



Figure 3. Qualitative comparison between our model and UNS [1]. We visualize the tracking on three frames. UNS cannot track the person in green bounding box once he is occluded and assigns him a new ID (purple) on frame 80. It also fails to track the person in pink bounding box when she is only partially visible in frame 80. We can consistently track with both people with the same IDs (green, pink).

	\mathcal{L}_{PC}	\mathcal{L}_{OM}	\mathcal{L}_{BC}	IDF1	HOTA	MOTA	IDsw
1	×	✓	✓	3.4	8.3	1.4	70k
2	(a)	✓	✓	26.9	26.8	37.8	29k
3	(b)	✓	✓	66.6	59.4	63.9	329
4	✓	×	✓	62.7	57.1	63.4	591
5	✓	✓	×	66.3	59.8	63.6	466
6	✓	✓	✓	68.9	60.9	63.7	257

Table 7. Ablation study on proposed losses. In row 1-3, we demonstrate the effect of our path consistency loss. In row 4-5, we show the effect of our regularization losses. “(a)” denotes randomly sampling start and end frames for paths and “(b)” denotes disabling spatial constraint mask in Eq 5.

In Table 6, we report IDF1 with different occlusion length L and compare with UNS and one of the state-of-the-art trackers, OC-SORT, as a strong baseline. Since other related methods [16, 22, 24, 30] do not release their codes or models, we are unable to compare with them. As can be seen, our model consistently provides higher IDF1 for all ranges of occlusion lengths. More importantly, when increasing occlusion lengths, the performance of UNS and OC-SORT decreases quickly, while our performance is more robust and decreases slowly, which echos our results in Table 4. Our IDF1 at $L = 50$ is still higher than the IDF1 of OC-SORT at $L = 20$, showing our model can learn to match the same objects over varied temporal distances and is more robust to occlusion.

What is the effect of path consistency loss? In Table 7 (row 1-3), we exam the results of removing PCL and also two modifications of it: (a) randomly sampling the starting and end frames for paths; (b) removing the spatial constraint mask in Eq.5; Firstly, without PCL (row 1), the model cannot learn meaningful object matching using only the regularization losses thus the tracking accuracy is close to zeros. When randomly sampling the starting and end frames (row 2), it causes most query objects are not visible in the end frames, causing imbalance between visible and disappeared objects. Thus, model matches most objects to *null* objects, also leading to low IDF1. Next, when removing the spatial constraint mask (row 3), the model may incorrectly match two spatial disclose objects, decreasing IDF1 by 2.3%.

What is the effect of regularization losses? In Table 7 (row 4-5), we ablate the effect of the regularization losses in Section 3.3. As can be seen, while PCL contributes most to the model performance, the two regularization losses \mathcal{L}_{OM} , \mathcal{L}_{BC} are also helpful for achieving better performance. Since \mathcal{L}_{OM} prevents multiple objects matching to the same objects and \mathcal{L}_{BC} ensures the consistency between forward and backward matching, removing either of these loss terms leads to around 6.2% and 2.6% drop in IDF1, respectively. These results show the collective benefits of using different losses to optimize our unsupervised object tracker.

4.4. Qualitative results

In Figure 3, we visualize the tracking results for one video from MOT17 using our model and UNS [1], where the color indicates their IDs. Both our model and UNS can track two non-occluded front people from frame 1 to frame 80. However, UNS lose the man in green bounding box after he is occluded and assign a new ID (purple) to him on frame 80. It also cannot associate the woman in the pink bounding box when she is only partially visible in frame 80. Our model correctly associate two people with the same IDs (green and pink). This qualitative comparison suggests that our model can better track the object over occlusion.

5. Conclusion

To tackle unsupervised MOT task, we proposed a novel concept of path consistency. We defined a path to track an object as a temporal list of observed frames from the start to the end frame. As such, to track the same object, we can generate multiple paths by skipping intermediate frames while the association results should be consistent across different paths. We leveraged this consistency in our self-supervised Path Consistency Loss to learn object matching model without groundtruth ID supervision. Our PCL especially learns robust long distance matching to tackle occlusion at inference time. With experiments on three popular datasets, our approach has exceeded existing unsupervised methods and performed on par with recent state-of-the-art supervised methods. We provided a comprehensive ablation study to validate the model has learned long-distance matching and is robust under challenging scenarios.

References

- [1] Favyen Bastani, Songtao He, and Samuel Madden. Self-supervised multi-object tracking with cross-input consistency. In *Advances in Neural Information Processing Systems*, pages 13695–13706. Curran Associates, Inc., 2021. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [2] Jerome Berclaz, Francois Fleuret, Engin Turetken, and Pascal Fua. Multiple object tracking using k-shortest paths optimization. *TPAMI*, 33(9):1806–1819, 2011. [2](#)
- [3] Philipp Bergmann, Tim Meinhardt, and Laura Leal-Taixé. Tracking without bells and whistles. In *The IEEE International Conference on Computer Vision (ICCV)*, 2019. [2](#), [6](#)
- [4] Alex Bewley, Zongyuan Ge, Lionel Ott, Fabio Ramos, and Ben Uprocft. Simple online and realtime tracking. In *2016 IEEE International Conference on Image Processing (ICIP)*, pages 3464–3468, 2016. [2](#), [5](#), [6](#), [7](#)
- [5] Jiarui Cai, Mingze Xu, Wei Li, Yuanjun Xiong, Wei Xia, Zhuowen Tu, and Stefano Soatto. Memot: Multi-object tracking with memory. In *CVPR 2022*, 2022. [1](#), [2](#), [6](#)
- [6] Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: Delving into high quality object detection. In *CVPR*, 2018. [2](#)
- [7] Jinkun Cao, Jiangmiao Pang, Xinshuo Weng, Rawal Khirrodkar, and Kris Kitani. Observation-centric sort: Rethinking sort for robust multi-object tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9686–9696, 2023. [2](#), [5](#), [6](#), [7](#)
- [8] Wongun Choi. Near-online multi-target tracking with aggregated local flow descriptor. In *Proceedings of the IEEE international conference on computer vision*, pages 3029–3037, 2015. [1](#), [6](#)
- [9] Peng Chu and Haibin Ling. Famnet: Joint learning of feature, affinity and multi-dimensional assignment for online multiple object tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6172–6181, 2019. [6](#)
- [10] Patrick Dendorfer, Hamid Rezatofighi, Anton Milan, Javen Shi, Daniel Cremers, Ian Reid, Stefan Roth, Konrad Schindler, and Laura Leal-Taixé. Mot20: A benchmark for multi object tracking in crowded scenes, 2020. [2](#)
- [11] Weitao Feng, Zhihao Hu, Wei Wu, Junjie Yan, and Wanli Ouyang. Multi-object tracking with multiple cues and switcher-aware classification, 2019. [6](#)
- [12] Ruopeng Gao and Limin Wang. MeMOTR: Long-term memory-augmented transformer for multi-object tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9901–9910, 2023. [6](#)
- [13] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *International Journal of Robotics Research (IJRR)*, 2013. [2](#), [5](#), [6](#)
- [14] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, 2017. [2](#)
- [15] Roberto Henschel, Laura Leal-Taixé, Daniel Cremers, and Bodo Rosenhahn. Improvements to frank-wolfe optimization for multi-detector multi-object tracking. *arXiv preprint arXiv:1705.08314*, 2017. [2](#)
- [16] Shyamgopal Karthik, Ameya Prabhu, and Vineet Gandhi. Simple unsupervised multi-object tracking. *CoRR*, abs/2006.02609, 2020. [1](#), [2](#), [4](#), [5](#), [6](#), [8](#)
- [17] Chanh Kim, Fuxin Li, Arridhana Ciptadi, and James M Rehg. Multiple hypothesis tracking revisited. In *ICCV*, 2015. [2](#)
- [18] Chanh Kim, Fuxin Li, and James M. Rehg. Multi-object tracking with neural gating using bilinear lstm. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. [6](#)
- [19] Shuai Li, Yu Kong, and Hamid Rezatofighi. Learning of global objective for network flow in multi-object tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8855–8865, 2022. [2](#)
- [20] Siyuan Li, Tobias Fischer, Lei Ke, Henghui Ding, Martin Danelljan, and Fisher Yu. Ovtrack: Open-vocabulary multiple object tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5567–5577, 2023. [2](#)
- [21] Yu-Lei Li. Unsupervised embedding and association network for multi-object tracking. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*. International Joint Conferences on Artificial Intelligence Organization, 2022. [1](#), [2](#), [4](#)
- [22] Kai Liu, Sheng Jin, Zhihang Fu, Ze Chen, Rongxin Jiang, and Jieping Ye. Uncertainty-aware unsupervised multi-object tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9996–10005, 2023. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [8](#)
- [23] Qiankun Liu, Qi Chu, Bin Liu, and Nenghai Yu. Gsm: Graph similarity model for multi-object tracking. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 530–536. International Joint Conferences on Artificial Intelligence Organization, 2020. Main track. [6](#)
- [24] Qiankun Liu, Dongdong Chen, Qi Chu, Lu Yuan, Bin Liu, Lei Zhang, and Nenghai Yu. Online multi-object tracking with unsupervised re-identification learning and occlusion estimation. *Neurocomputing*, 483:333–347, 2022. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [25] Shuai Liu, Xin Li, Huchuan Lu, and You He. Multi-object tracking meets moving uav. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8876–8885, 2022. [2](#)
- [26] Yang Liu, Idil Esen Zulfikar, Jonathon Luiten, Achal Dave, Deva Ramanan, Bastian Leibe, Aljoša Ošep, and Laura Leal-Taixé. Opening up open world tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19045–19055, 2022.
- [27] Fan Ma, Mike Zheng Shou, Linchao Zhu, Haoqi Fan, Yilei Xu, Yi Yang, and Zhicheng Yan. Unified transformer tracker for object tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8781–8790, 2022.
- [28] Christoph Mayer, Martin Danelljan, Goutam Bhat, Matthieu Paul, Danda Pani Paudel, Fisher Yu, and Luc Van Gool.

- Transforming model prediction for tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8731–8740, 2022.
- [29] Tim Meinhardt, Alexander Kirillov, Laura Leal-Taixe, and Christoph Feichtenhofer. Trackformer: Multi-object tracking with transformers. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 2, 6
- [30] Sha Meng, Dian Shao, Jiacheng Guo, and Shan Gao. Tracking without label: Unsupervised multiple object tracking via contrastive similarity learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 16264–16273, 2023. 1, 2, 3, 5, 6, 7, 8
- [31] Anton Milan, Laura Leal-Taixé, Ian D. Reid, Stefan Roth, and Konrad Schindler. MOT16: A benchmark for multi-object tracking. *CoRR*, abs/1603.00831, 2016. 1, 2, 5, 6
- [32] Jiangmiao Pang, Linlu Qiu, Xia Li, Haofeng Chen, Qi Li, Trevor Darrell, and Fisher Yu. Quasi-dense similarity learning for multiple object tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 164–173, 2021. 2
- [33] Z. Qin, S. Zhou, L. Wang, J. Duan, G. Hua, and W. Tang. Motiontrack: Learning robust short-term and long-term motions for multi-object tracking. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 1, 6
- [34] Hao Ren, Shoudong Han, Huilin Ding, Ziwen Zhang, Hongwei Wang, and Faquan Wang. Focus on details: Online multi-object tracking with diverse fine-grained representation. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 1, 6
- [35] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NeurIPS*, 2015. 2
- [36] Ergys Ristani and Carlo Tomasi. Features for multi-target multi-camera tracking and re-identification. In *CVPR*, 2018. 2
- [37] Hao Sheng, Yang Zhang, Jiahui Chen, Zhang Xiong, and Jun Zhang. Heterogeneous association graph fusion for target association in multiple object tracking. *TCSVT*, 29(11): 3269–3280, 2018. 2, 3
- [38] Bing Shuai, Andrew Berneshawi, Xinyu Li, Davide Modolo, and Joseph Tighe. Siammot: Siamese multi-object tracking. In *CVPR*, 2021. 2, 6
- [39] Bing Shuai, Alessandro Bergamo, Uta Buechler, Andrew Berneshawi, Alyssa Boden, and Joe Tighe. Large scale real-world multi person tracking. In *European Conference on Computer Vision*. Springer, 2022. 2, 5, 6
- [40] Bing Shuai, Xinyu Li, Kaustav Kundu, and Joseph Tighe. Id-free person similarity learning. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14669–14679, 2022. 6
- [41] Bing Shuai, Xinyu Li, Kaustav Kundu, and Joseph Tighe. Id-free person similarity learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 14689–14699, 2022. 3
- [42] Peize Sun, Jinkun Cao, Yi Jiang, Rufeng Zhang, Enze Xie, Zehuan Yuan, Changhu Wang, and Ping Luo. Transtrack: Multiple-object tracking with transformer. *arXiv preprint arXiv: 2012.15460*, 2020. 6
- [43] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20993–21002, 2022. 2
- [44] Siyu Tang, Mykhaylo Andriluka, Bjoern Andres, and Bernt Schiele. Multiple people tracking by lifted multicut and person re-identification. In *CVPR*, 2017. 2, 3
- [45] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *ICCV*, 2019. 2
- [46] Pavel Tokmakov, Jie Li, Wolfram Burgard, and Adrien Gaidon. Learning to track with object permanence. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10840–10849, 2021. 1, 6
- [47] Gaoang Wang, Renshu Gu, Zuozhu Liu, Weijie Hu, Mingli Song, and Jenq-Neng Hwang. Track without appearance: Learn box and tracklet embedding with local and global motion patterns for vehicle tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9876–9886, 2021. 6
- [48] Zhongdao Wang, Liang Zheng, Yixuan Liu, Yali Li, and Shengjin Wang. Towards real-time multi-object tracking. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XI 16*, pages 107–122. Springer, 2020. 2
- [49] S. You, H. Yao, B. Bao, and C. Xu. Utm: A unified multiple object tracking model with identity-aware feature enhancement. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 6
- [50] En Yu, Zhuoling Li, and Shoudong Han. Towards discriminative representation: Multi-view trajectory contrastive learning for online multi-object tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8834–8843, 2022. 2
- [51] Amir Roshan Zamir, Afshin Dehghan, and Mubarak Shah. Gmcp-tracker: Global multi-object tracking using generalized minimum clique graphs. In *ECCV*, 2012. 2
- [52] Fangao Zeng, Bin Dong, Yuang Zhang, Tiancai Wang, Xianguyu Zhang, and Yichen Wei. Motr: End-to-end multiple-object tracking with transformer. In *European Conference on Computer Vision (ECCV)*, 2022. 1, 2, 6
- [53] Li Zhang, Yuan Li, and Ramakant Nevatia. Global data association for multi-object tracking using network flows. In *CVPR*, 2008. 2
- [54] Wenwei Zhang, Hui Zhou, Shuyang Sun, Zhe Wang, Jianping Shi, and Chen Change Loy. Robust multi-modality multi-object tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2365–2374, 2019. 6
- [55] Yifu Zhang, Chunyu Wang, Xinggong Wang, Wenjun Zeng, and Wenyu Liu. Fairmot: On the fairness of detection and re-identification in multiple object tracking. *International Journal of Computer Vision*, 129:3069–3087, 2021. 1, 2, 3, 6

- [56] Yifu Zhang, Peize Sun, Yi Jiang, Dongdong Yu, Fucheng Weng, Zehuan Yuan, Ping Luo, Wenyu Liu, and Xinggang Wang. Bytetrack: Multi-object tracking by associating every detection box. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. [2](#), [3](#), [5](#), [6](#)
- [57] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *arXiv preprint arXiv:1904.07850*, 2019. [2](#)
- [58] Xingyi Zhou, Vladlen Koltun, and Philipp Krähenbühl. Tracking objects as points. *ECCV*, 2020. [1](#), [2](#), [6](#)
- [59] Xingyi Zhou, Vladlen Koltun, and Philipp Krähenbühl. Probabilistic two-stage detection. *arXiv preprint arXiv:2103.07461*, 2021. [2](#)
- [60] Xingyi Zhou, Tianwei Yin, Vladlen Koltun, and Philipp Krähenbühl. Global tracking transformers. In *CVPR*, 2022. [2](#)