

Focus On Details: Online Multi-object Tracking with Diverse Fine-grained Representation

Hao Ren¹, Shoudong Han^{1,*}, Huilin Ding, Ziwen Zhang, Hongwei Wang, Faquan Wang

National Key Laboratory of Science and Technology on Multispectral Information Processing,
School of Artificial Intelligence and Automation, Huazhong University of Science and Technology

{haoren2000, shoudonghan}@hust.edu.cn

Abstract

Discriminative representation is essential to keep a unique identifier for each target in Multiple object tracking (MOT). Some recent MOT methods extract features of the bounding box region or the center point as identity embeddings. However, when targets are occluded, these coarse-grained global representations become unreliable. To this end, we propose exploring diverse fine-grained representation, which describes appearance comprehensively from global and local perspectives. This fine-grained representation requires high feature resolution and precise semantic information. To effectively alleviate the semantic misalignment caused by indiscriminate contextual information aggregation, Flow Alignment FPN (FAFPN) is proposed for multi-scale feature alignment aggregation. It generates semantic flow among feature maps from different resolutions to transform their pixel positions. Furthermore, we present a Multi-head Part Mask Generator (MPMG) to extract fine-grained representation based on the aligned feature maps. Multiple parallel branches of MPMG allow it to focus on different parts of targets to generate local masks without label supervision. The diverse details in target masks facilitate fine-grained representation. Eventually, benefiting from a Shuffle-Group Sampling (SGS) training strategy with positive and negative samples balanced, we achieve state-of-the-art performance on MOT17 and MOT20 test sets. Even on DanceTrack, where the appearance of targets is extremely similar, our method significantly outperforms ByteTrack by 5.0% on HOTA and 5.6% on IDF1. Extensive experiments have proved that diverse fine-grained representation makes Re-ID great again in MOT.

1. Introduction

As a fundamental task in computer vision, multi-object tracking (MOT) is crucial for automatic driving, video

¹ Equal contribution

* Corresponding author




Figure 1. **Comparison with different methods of appearance representation:** (a) Bbox-based, (b) Center-based, (c) global-local fine-grained representation (ours).

surveillance, etc. MOT aims to localize targets and maintain their unique identities. Recent MOT methods [3, 4, 48, 50, 54] mainly follow the paradigm of tracking-by-detection, and divide tracking into two independent steps: detection and association. The detector detects targets in each frame first, and then appearance representation and position information are employed as the association basis to link targets with the corresponding trajectories. As the inherent attributes of the target, appearance and position complement each other in the association.

However, due to targets or camera motion, intra-class and inter-class occlusion are inevitable, which puts forward stricter requirements for appearance representation. As shown in Fig. 1a and Fig. 1b, these methods extract the features of the bounding box region or the center point as appearance embeddings. However, these coarse-grained global embeddings are extremely sensitive to noise, so that become unreliable once the signal-to-noise ratio is reduced. With the headway of the detector [5, 11, 27, 35, 58], appearance representation gradually cannot keep the same performance as detection. Some researchers [54] find that simply using position cues is enough to obtain satisfactory results, while appearance cues are unfavorable for further improvement.

To break this situation, we re-examine the recent appearance-based methods. The bbox-based methods [38]

in Fig. 1a adopt global average pooling, which converts features of the bounding box region into appearance embeddings. These methods equally treat target and interference features (background and other objects), which is unreasonable. As shown in Fig. 1b, researchers [49, 55, 57] notice this issue and utilize features at the target centers as their appearance embeddings, eliminating interference from the background as much as possible. Despite this, when the target is occluded, the feature at its center is still inevitably interfered with noise information from other objects. The fuzziness of global representation has become an impediment of these methods. On the contrary, our method focuses on different local details of targets, which is illustrated in Fig. 1c. Fine-grained global and local representations complement each other and jointly describe appearance. When the target is occluded, our method can still identify it according to visible parts, similar to human judgment.

As the basis of fine-grained representation, target feature maps require high-resolution and unambiguous semantic information. Shallow or deep outputs cannot meet these two requirements simultaneously. Therefore, it is feasible to enrich the semantic information of shallow features or improve the resolution of deep features. To reduce the burden, researchers usually adopt FPN [21] to aggregate multi-scale shallow feature maps indiscriminately, which causes semantic misalignment among features with different resolutions. Specifically, there is a spatial misalignment between the late feature maps after up-sampling and the early feature maps.

To solve this issue, we construct a *Flow Alignment FPN* (FAFPN) to learn the semantic flow among feature maps with different scales and effectively broadcast high-level features to high-resolution features. FAFPN aligns feature maps by semantic flow and aggregates context information to enrich semantic information while maintaining high resolution. Further, *Multi-head Part Mask Generator* (MPMG) is proposed to generate part masks for detailed representations without label supervision. Inspired by multi-head self-attention in Transformer [37], MPMG implements a multi-branch parallel structure, which can efficiently and comprehensively focus on different parts of the target. Combining FAFPN and MPMG, we can obtain a diverse fine-grained representation, including diverse local embeddings and background-filtered global embeddings.

In the training phase, some current MOT methods [20, 55] train Re-ID (re-identification) by following the video sequence or shuffling all video frames. The former does not disperse training data, while the latter is positive and negative samples imbalanced. To train Re-ID more reasonably, we propose a training strategy named *Shuffle-Group Sampling* (SGS). In this strategy, we group video frames into short segments in their order and then shuffle these segments. SGS disperses data and balances pos-

itive and negative samples. Our model incorporates all the above proposed techniques, named *fine-grained representation tracker* (FineTrack).

The main contributions of our work can be summarized as follows:

- We propose a *Flow Alignment FPN* (FAFPN). It learns the semantic flow among feature maps with different resolutions to correct spatial dislocation. Feature maps after FAFPN include high resolution and precise semantic information, which is the basis of fine-grained representation.
- We construct a *Multi-head Part Mask Generator* (MPMG) to focus on details of targets. MPMG employs self-attention to filtrate background noise and extract global-local embeddings to represent targets comprehensively.
- *Shuffle-Group Sampling* (SGS) is proposed to disperse training data and balance positive and negative samples. It reduces oscillation in model convergence to achieve better performance.

2. Related Work

Tracking-by-detection. With the development of detection, many works [1, 24, 26, 42] adopt the tracking-by-detection paradigm, and divide MOT into detection and association. The trackers following this paradigm first detect objects in video frames and then associate them dependent on their identity information. Some works [1, 10, 31, 34] model object movement for identity association. SORT [3] relies on the Kalman Filter [15] to predict future positions of the targets, calculates their overlap with detection, and utilizes the Hungarian algorithm [17] for identity association. When the target moves irregularly or is occluded, only using motion information is not enough to achieve better performance. Based on SORT [3], DeepSORT [43] introduces the appearance information of targets to enhance identity embedding.

Appearance and motion, as inherent properties of targets, are not independent but complementary while tracking. However, appearance does not always provide reliable clues for identification, particularly when targets are occluded. Our method focuses on diverse details of targets and constructs global-local fine-grained representation, which achieves better performance in occlusion scenes.

Re-ID in MOT. Some works [9, 50, 59] crop the image regions of detections and extract features with an extra Re-ID model. These works treat detection and Re-ID as independent tasks, imposing additional parameter burdens. To solve this problem, JDE [42] and FairMOT [55] implement a branch the same as detection for appearance representation. The structure of the multi-task branches achieves impressive performance but also raises an additional issue: com-

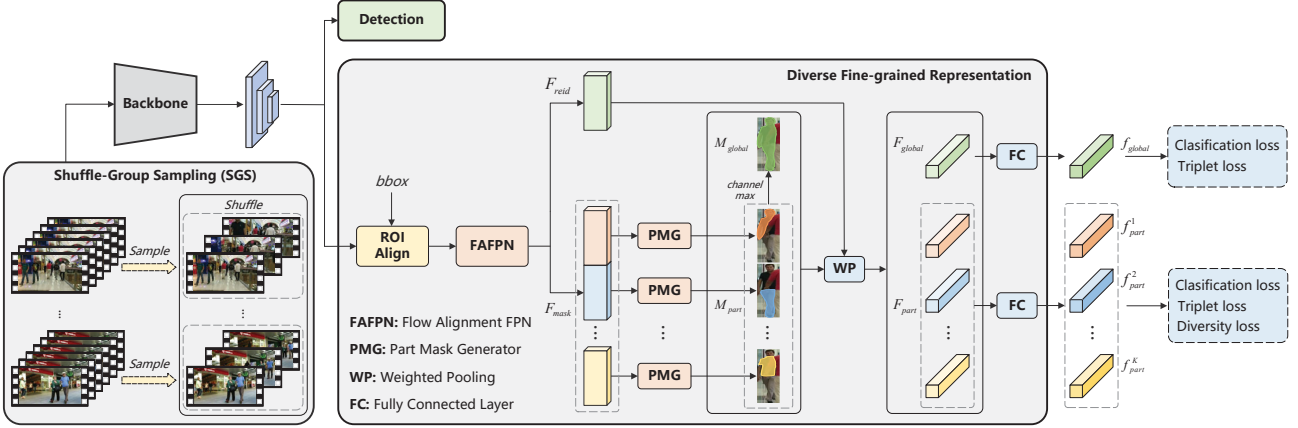


Figure 2. **Overall pipeline of FineTrack.** FineTrack comprises 5 steps: (1) SGS divides videos into groups and shuffles them. (2) Extract feature maps of the input frame for detection. Meanwhile, multi-scale shallow feature maps are obtained for ROIAlign to output target feature maps with different resolutions. (3) FAFPN aligns target feature maps and aggregates their context information for F_{mask} and F_{reid} . (4) Divide F_{mask} into K blocks along their channel and employ multi-head PMG to generate part masks M_{part} . Concatenate M_{part} and adopt Max of channel to form global masks M_{global} . (5) F_{global} and F_{part} can be obtained by a weighted pooling. Then two fully connected layers convert them into global and part embedding: f_{global} and f_{part} , respectively.

petition caused by different goals for detection and Re-ID. To alleviate this conflict, CStrack [20] constructs a CCN module to extract general and specific features more suitable for detection and Re-ID. However, these appearance-based methods extract the features of the boundary box region or the center point as appearance representation, which are inevitably interfered with irrelevant information. They focus on generating more discriminative global embedding instead of exploring more detailed appearance cues.

Although well-designed Re-ID modules can improve performance, global features are still prone to ambiguity. With a high-performance detector, the associator can link most targets only with motion cues. The global representation is also impotent for the remaining targets incorrectly tracked due to occlusion or other factors. Using global appearance cues at this point can even compromise tracking accuracy. Instead, our method focus on the details of targets. Following this idea, we construct the FAFPN to align and aggregate high-resolution shallow feature maps and feed them into the MPMG to obtain fine part masks. Further, part masks of targets are employed to extract the fine-grained global-local appearance embedding, which can more comprehensively represent the identity.

3. Methodology

3.1. Overview

In this work, we adopt the high-performance detector YOLOX [11] to detect targets. As shown in Fig. 2, during training, SGS groups video frames by sampling them sequentially, ensuring targets with positive and negative samples. After shuffling, these grouped data are fed into the

Backbone to obtain feature maps with different resolutions. Then, we adopt ROIAlign to crop and scale feature maps of the bbox regions to get the multi-scale feature maps of targets. Immediately, FAFPN aligns these target feature maps from different resolutions and aggregates them into fine-grained feature maps: F_{mask} and F_{reid} for mask generation and appearance embedding, respectively. After slicing F_{mask} along the channel dimension, the features are fed into separate *Part-Mask Generator* (PMG) to generate target masks. Further, F_{global} and F_{part} can be obtained by a weighted pooling. These embeddings are fed into two fully connected layers that do not share parameters to output f_{global} and f_{part} . For f_{global} , we calculate Classification loss and Triplet loss. For f_{part} , a diversity loss is added to expand the discrepancy among different part features.

3.2. Flow Aligned FPN

As a classical method for feature aggregation, FPN [21] has been extensively applied in different computer vision tasks. However, step-by-step downsampling and indiscriminate context aggregation cause semantic misalignment among feature maps of different scales. Fuzzy semantic information have a significant impact on visual tasks which require detailed descriptions. To obtain fine-grained features, we employ the *Flow Alignment Module* (FAM) [14, 19] to generate semantic flow among feature maps of different resolutions. The semantic flow can guide alignment and eliminate spatial dislocation among feature maps from different scales. Furthermore, we utilize the FAM to optimize the aggregation process of FPN and then construct a *Flow Alignment FPN* (FAFPN).

The overall structure of FAFPN is shown in Fig. 3. *Res-*

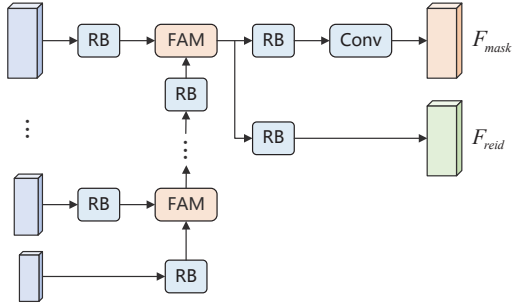


Figure 3. **Structure of FAFPN.** FAFPN consists of FAM and RB for aggregating feature maps from S different scales

Block (RB) replaces 1×1 Convolution in FPN to fuse shallow features more adequately. The *Sum* operation is replaced with the FAM, making semantic alignment possible. The multi-scale feature maps are modified to a unified channel dimension by RB and then fed into the FAM for alignment and aggregation. Finally, we use two RB to generate F_{mask} and F_{reid} respectively for meeting the different requirements of the mask and representation. Fig. 4 shows the specific structure of RB and FAM. The RB is a simple residual block. It combines 1×1 Convolution and BN layer to transform the input channel dimension and then applies the residual structure to get its output. The FAM takes two feature maps from different scales as input. f_l has the same size as f_h after up-sampling. Further, they are concatenated to generate semantic flow: $flow_{down}$ and $flow_{up}$, which represent the pixel offset of feature maps f_h and f_l , respectively. *Warp* operation aligns these feature maps based on semantic flow. In the end, the FAM employs an *Sum* operation to aggregate the context information of aligned feature maps.

After semantic alignment, multi-scale shallow feature maps are aggregated into high-resolution feature maps with precise semantic information, which are the basis for generating diverse fine-grained representations.

3.3. Multi-head Part Mask Generator

Some recent works rely on well-designed modules to optimize global embedding. These methods indiscriminately generalize the various details of the target into a coarse-grained global embedding, which causes semantic ambiguity. When the target is occluded, the features of the bounding box region or the center point are inevitably disturbed by noise.

To alleviate this problem, we propose exploring fine-grained appearance representation to focus on target diversity details. The direct approach is to adopt an attention module to enhance feature maps. There are many excellent attention modules in Convolutional Neural Networks, such as Non-Local [40], CBAM [44], etc. Simply employing attention modules can not entirely focus on the details

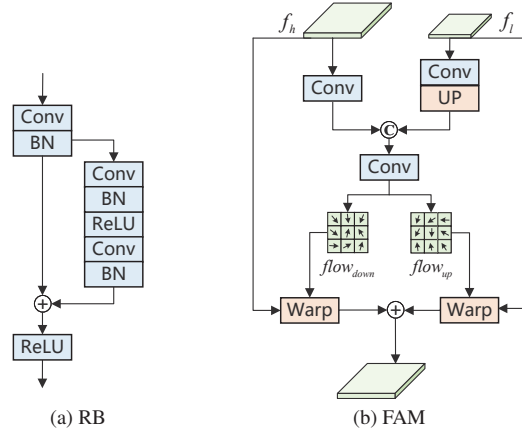


Figure 4. **Implementation of RB and FAM.** RB is a form of residual structure. FAM learns semantic flow to align feature maps from different scales.

of targets. Inspired by the multi-head self-attention module in Transformer [37], we establish multiple parallel attention branches to focus on different parts of targets. At the same time, we employ the Non-Local block as the branch attention module, followed by a Convolution to generate a mask of the target part. As shown in Fig. 5, we divide the input feature maps into blocks along their channel dimension, and each block can be considered as a different mapping of the target feature. These feature maps are fed into *Multi-head Part Mask Generator (MPMG)* to generate different part masks of the target without Ground Truth.

For any branch in MPMG namely PMG, three 1×1 Convolutions are respectively used to obtain query, key, and value. Query and key generate the attention matrix, which is multiplied by value and then added with input. After that, a layer combining 1×1 Convolution and Sigmoid is used to generate a part mask with values between 0 and 1. The same operation is performed for the remaining feature blocks, resulting in K different part masks. After concatenation, the channel of part masks is aggregated by *Max* to form a global mask containing different local information. It is worth noting that parameters are not shared among parallel branches, which ensures sufficient diversity of attention.

3.4. Train and Inference

Shuffle-Group Sampling. For better detection performance, researchers shuffle consecutive video frames, which is effective for training detection but inappropriate for training Re-ID. This training strategy makes a batch contain almost no targets with the same identity. Such imbalanced positive and negative samples cause the model only to discover the difference among the different targets but fail to capture the commonalities of the same targets. This is particularly detrimental to Re-ID, which aims to distinguish the

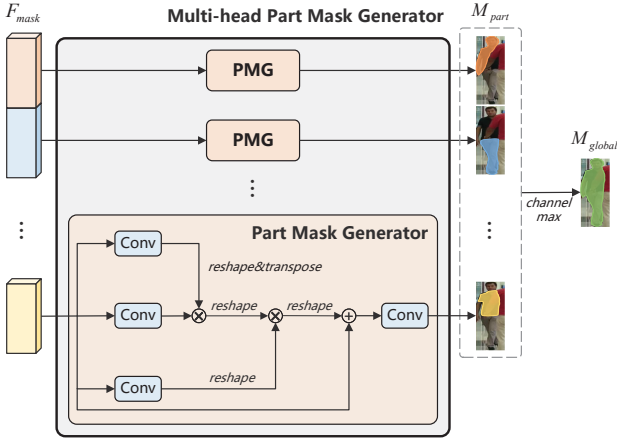


Figure 5. **Illustration of MPMG.** Multiple parallel PMG branches focus on different parts of the target and generate part masks.

identity of targets.

To solve the above problem, we construct the *Shuffle-Group Sampling* (SGS) training strategy. Different from random sampling, SGS adopts sequential sampling to group video frames. In this way, targets in the same batch hold positive samples with the same identity, thus alleviating the problem of imbalanced positive and negative samples caused by random sampling. In training, we disrupt the order of grouped data to reduce convergence fluctuations. This strategy allows contiguous batches to originate from different videos and form a sequence of track segments with significant appearance variations. Such a discrete distribution of training data optimizes the convergence direction of the model.

Training Loss. SGS enables a batch of targets to contain both positive and negative samples, so the model can be optimized using Triplet loss, which is extensively applied in Re-ID methods. The part and global features are denoted as $f_{part} = \{f_p^{n,k}, n \in [1, 2, \dots, N], k \in [1, 2, \dots, K]\}$ and $f_{global} = \{f_g^n, n \in [1, 2, \dots, N]\}$. Here N represents the number of targets in the image and each target is composed of K parts. Thus, $f_p^{n,k}$ means the k_{th} part feature of the n_{th} target, and f_g^n represents the global feature of the n_{th} target.

Then we can calculate the Triplet loss with a soft margin for part features based on Eq. (1):

$$L_{tri}^p(k) = \frac{1}{K} \sum_{k=1}^K \text{Triplet}(f_p^k), \quad (1)$$

where $\text{Triplet}(\cdot)$ is the same as [13] and f_p^k represents the k_{th} part features of n targets in the image.

Similarly, the Triplet loss of the global features is based on Eq. (2):

$$L_{tri}^g(k) = \text{Triplet}(f_{global}), \quad (2)$$

After obtaining K part features, K combinations of Linear layer and Softmax are respectively applied to get the classification result vectors $P = \{p_n^k, k \in [1, 2, \dots, K], n \in [1, 2, \dots, N]\}$. The meanings of K and N are the same as in f_{part} . p_n^k represents the classification vector of the k_{th} part feature of the n_{th} target. The dimension of p_n^k is M , which is also the number of all targets in the datasets. For target identity labels, we denote them as $Y = \{y_{n,m}, n \in [1, 2, \dots, N], m \in [1, 2, \dots, M]\}$, $y_{n,m}$ means whether the n_{th} target has the same identity as the m_{th} target in ID labels, with a value of 0 or 1. According to the above definition, we can calculate the classification loss of the k_{th} part feature of the n_{th} target with Eq. (3):

$$L_{n,k}^p(m) = y_{n,m} \log(p_n^k(m)), m \in [1, 2, \dots, M] \quad (3)$$

Further, we calculate the classification loss of part features based on with Eq. (4):

$$L_{cls}^p = -\frac{1}{K \cdot N} \sum_{k=1}^K \sum_{n=1}^N \sum_{m=1}^M L_{n,k}^p(m) \quad (4)$$

Similarly, the classification loss of global features is:

$$L_{cls}^g = -\frac{1}{N} \sum_{n=1}^N \sum_{m=1}^M y_{n,m} \log(g_n(m)), \quad (5)$$

where g_n is the classification vector of the n_{th} target's global feature.

Due to the peculiarities of multi-branch structures, using only classification loss and Triplet loss does not ensure that the model focuses on different parts of the target. To avoid multiple branches gazing at similar details, we employ the diversity loss L_{div} in Eq. (6) to distance different part features of the same target:

$$L_{div} = \frac{1}{N \cdot K(K-1)} \sum_{n=1}^N \sum_{k_i \neq k_j}^K \frac{\langle f_p^{n,k_i}, f_p^{n,k_j} \rangle}{\|f_p^{n,k_i}\|_2 \cdot \|f_p^{n,k_j}\|_2} \quad (6)$$

The purpose of diversity loss is intuitive, which is to keep the cosine similarity between different part features of the same target as low as possible.

We combine the above losses into a final training loss:

$$L = \alpha \cdot (L_{cls}^p + L_{tri}^p) + \beta \cdot (L_{cls}^g + L_{tri}^g) + \gamma \cdot L_{div}, \quad (7)$$

where α, β and γ are used to adjust the proportion of different losses and we set $\alpha = 0.5 * K, \beta = 0.3, \gamma = 2$.

Inference. Based on ByteTrack [54], we add a method similar to DeepSort [43] that calculates Re-ID features into feature distance matrix. It is worth mentioning that we concatenate part features of targets with global features as Re-ID features.

We adopt the exponential moving average (EMA) mechanism to update the features \tilde{f}_i^t of matched tracklets for the i_{th} tracklet at frame t based on Eq. (8), as in [55].

$$\tilde{f}_i^t = \lambda \tilde{f}_i^{t-1} + (1 - \lambda) f_i^t, \quad (8)$$

where f_i^t is the feature of the current matched detection, and $\lambda = 0.9$ is a momentum term.

In the tracking algorithm, the feature distance matrix d_{feat} is:

$$d_{feat} = 1 - \text{Similarity}(\tilde{f}^{t-1}, f^t), \quad (9)$$

where $\text{Similarity}(\cdot)$ outputs the cosine similarity matrix between tracklets features \tilde{f}^{t-1} and targets features f^t .

Meanwhile, we can calculate the IoU distance matrix d_{IoU} base on Eq. (10).

$$d_{IoU} = 1 - \text{IoU}(b_{det}, b_{pre}), \quad (10)$$

where $\text{IoU}(\cdot)$ outputs the IoU matrix between detection bboxes b_{det} and prediction bboxes b_{pre} .

To exclude the interference of distant targets, we only consider feature distances between pairs of targets with the IoU distance less than 1, which means there is bounding box overlap. The optimized feature distance matrix \tilde{d}_{feat} is:

$$\tilde{d}_{feat} = 1 - (1 - d_{feat}) \cdot (d_{IoU} < 1) \quad (11)$$

After squaring the product of the optimized feature distance matrix and IoU matrix, the final distance matrix is obtained for the association with Eq. (12):

$$d = \sqrt{\tilde{d}_{feat} \cdot d_{IoU}} \quad (12)$$

Finally, we set the association threshold to 0.5.

4. Experiments

4.1. Settings

Datasets. To verify the effectiveness of FineTrack, we test it on MOT17 [23], MOT20 [7], and DanceTrack [32] datasets. MOT17 and MOT20 datasets provide training sets but do not contain validation sets. The test metric only can be obtained by uploading the tracking results to the MOTChallenge website for evaluation. Therefore, in the ablation experiment phase, we divided the first half of each video of the MOT17 dataset into the training set and the second half as the validation set. Due to the difference between detection and Re-ID, we train them separately. For testing on MOT17, we train a detector using the same datasets as ByteTrack, including CrowdHuman [30], MOT17, Cityperson [52], and ETHZ [8], and then freeze the trained detector parameters and train Re-ID separately on MOT17.

DanceTrack is a large-scale multi-object tracking dataset for occlusions, frequent crossing, uniform appearance, and

Method	HOTA↑	IDF1↑	MOTA↑	FP↓	FN↓	IDs↓
<i>MOT17 private detection</i>						
DAN [34]	39.3	49.5	52.4	25423	234592	8431
TubeTK [24]	48.0	58.6	63.0	27060	177483	4137
MOTR [51]	-	66.4	65.1	45486	149307	2049
CTracker [26]	49.0	57.4	66.6	22284	160491	5529
MAT [12]	53.8	63.1	69.5	30660	138741	2844
QuasiDense [25]	53.9	66.3	68.7	26589	146643	3378
TransTrack [33]	54.1	63.5	75.2	50157	86442	3603
TransCenter [46]	54.5	62.2	73.2	23112	123738	4614
GSDT [41]	55.2	66.5	73.2	26397	120666	3891
PermaTrackPr [36]	55.5	68.9	73.8	28998	115104	3699
SOTMOT [56]	-	71.9	71.0	39537	118983	5184
FUFET [29]	57.9	68.0	76.2	32796	98475	3237
MTrack [48]	-	72.1	73.5	53361	101844	2028
FairMOT [55]	59.3	72.3	73.7	27507	117477	3303
CSTrack [20]	59.3	72.6	74.9	23847	114303	3567
SiamMOT [31]	-	72.3	76.3	-	-	-
ReMOT [47]	59.7	72.0	77.0	33204	93612	2853
Semi-TCL [18]	59.8	73.2	73.3	22944	124980	2790
CorrTracker [39]	60.7	73.6	76.5	29808	99510	3369
RelationTrack [49]	61.0	74.7	73.8	27999	118623	1374
TransMOT [6]	61.7	75.1	76.7	36231	93150	2346
ByteTrack [54]	63.1	77.3	80.3	25491	83721	2196
FineTrack	64.3	79.5	80.0	21750	90096	1272
<i>MOT20 private detection</i>						
MLT [53]	43.2	54.6	48.9	45660	216803	2187
TransTrack [33]	48.5	59.4	65.0	27197	150197	3608
FairMOT [55]	54.6	67.3	61.8	103440	88901	5243
Semi-TCL [18]	55.3	70.1	65.2	61209	114709	4139
CSTrack [20]	54.0	68.6	66.6	25404	144358	3196
GSDT [41]	53.6	67.5	67.1	31913	135409	3131
SiamMOT [31]	-	69.1	67.1	-	-	-
RelationTrack [49]	56.5	70.5	67.2	61134	104597	4243
SOTMOT [56]	-	71.4	68.6	57064	101154	4209
ByteTrack [54]	61.3	75.2	77.8	26249	87594	1223
FineTrack	63.6	79.0	77.9	24439	89012	980

Table 1. Comparison of the state-of-the-art methods under the private detection on the **MOT17** and **MOT20** test set. The best results are marked in **bold** and our method is highlighted in **pink**.

human body tracking with different body postures, with 100 videos. It uses 40, 25, and 35 videos as training, verification, and test set, respectively. The pedestrians in DanceTrack wear remarkably similar clothes, which is a massive challenge for Re-ID. Many existing methods that rely on Re-ID are at a disadvantage on this dataset.

Metrics. We use CLEAR-MOT Metrics [2], such as HOTA, MOTA, IDF1, IDs, FP, FN, etc., to evaluate the tracking performance. MOTA focuses on detection performance, and IDF1 [28] emphasizes association performance. Compared with them, HOTA [22] comprehensively balances detection, association, and localization effects. To accurately measure the performance of Re-ID and exclude the influence

Method	HOTA↑	IDF1↑	MOTA↑	AssA ↑	DetA ↑
CenterTrack [57]	41.8	35.7	86.8	22.6	78.1
FairMOT [55]	39.7	40.8	82.2	23.8	66.7
QuasiDense [25]	45.7	44.8	83.0	29.2	72.1
TransTrack [33]	45.5	45.2	88.4	27.5	75.9
TraDes [45]	43.3	41.2	86.2	25.4	74.5
ByteTrack [54]	47.7	53.9	89.6	32.1	71.0
FineTrack	52.7	59.8	89.9	38.5	72.4

Table 2. Comparison of the state-of-the-art methods on the **DanceTrack** test set. The best results are marked in **bold** and our method is highlighted in **pink**.

of detection and association, Rank-1 and mAP are utilized as metrics to measure the ability of Re-ID feature representation.

Implementation details. The partial feature and global feature dimensions are 128 and 256. The number of target masks K is set as 6 in MOT17 and 4 in MOT20. The detector training Settings are consistent with Bytetrack. During training, the batch size is set to 8. The model parameters are updated using the Adam optimizer [16] with an initial learning rate of 2×10^{-4} and 20 training epochs. The learning rate is reduced to 2×10^{-5} at the 10th epoch.

4.2. Comparison with the State-of-the-art Methods

In this part, we compare the performance of FineTrack with previous SOTA methods on three benchmarks, i.e., MOT17, MOT20 and DanceTrack. Notably, some MOT methods utilize extra Re-ID models and Re-ID datasets to improve their performance of identity embedding. For a fair comparison, FineTrack only uses MOT datasets for Re-ID training and does not use any additional labels for supervision, such as masks, etc.

MOT17: FineTrack shows significant advantages without introducing extra Re-ID models and datasets containing a large amount of identity information. As shown in Tab. 1, FineTrack method achieves the best tracking accuracy on the MOT17 datasets (i.e. 64.3% HOTA, 79.5% IDF1 and 80.0% MOTA, etc.). In cases where using detection simply can obtain excellent tracking accuracy, FineTrack achieves further improvement compared to ByteTrack (1.2% HOTA and 2.2% IDF1), which also adopts YOLOX as the detector. This demonstrates that our proposed fine-grained appearance representation dramatically improves the performance of identity embedding. Compared with other MOT methods, the advantages of FineTrack are more prominent.

MOT20: Compared with MOT17, MOT20 is more crowded, which causes frequent and dense occlusion. As shown in Tab. 1, FineTrack outperforms ByteTrack by 2.3% on HOTA, 3.8% on IDF1 with superior identity embeddings while reducing IDs from 1223 to 980. FineTrack has a stable and excellent performance in dense pedestrian scenes,

Method	Metrics	
	Rank-1↑	mAP↑
1 FineTrack(train by the order of videos)	89.7	56.7
2 FineTrack(shuffle all video frames)	90.9	60.4
3 FineTrack(w SGS)	92.3	61.8

Table 3. **Ablation study of SGS.** The first row does not disturb the video frames and trains in their order. The second row represents shuffling all the video frames. The last row is our SGS training strategy. (SGS: Shuffle-Group Sampling)

highlighting the role of fine-grained appearance embedding. **DanceTrack:** Pedestrians in DanceTrack dress very similarly and have complex body pose variations, making appearance-based methods perform poorly. FineTrack has excellent appearance modeling capabilities, as shown in Tab. 2, where we outperform ByteTrack entirely (5.0% on HOTA and 5.6% on IDF1, etc.). This further indicates the superiority of fine-grained appearance representations.

4.3. Ablation Studies

In this section, we verify the effectiveness of FineTrack through ablation studies. All experiments are conducted on the MOT17 dataset. Since the MOT Challenge does not provide validation sets, we split the MOT17 dataset into two parts. The first half of each video serves as the training set and the second half as the validation set. To fairly measure the performance of identity embedding, experiments in Tab. 3, Tab. 4 and Tab. 5 utilize the Ground Truth of bboxes as input, which can eliminate the interference caused by false detection. Tab. 3 verifies the effectiveness of the SGS training strategy, which is employed in the remaining ablation studies.

Analysis of SGS. To validate the advantages of our proposed SGS training strategy, we compare it with two other typical training methods, as reported in Tab. 3. Our proposed SGS is superior to the first training method by 2.6% on Rank-1 and 5.1% on mAP. While shuffling all video frames gets improvement (1.2% on Rank-1 and 3.7% on mAP) over training in the order of video frames, there is still a gap of 1.4% on both Rank-1 and mAP compared with SGS. This ablation experiment proves the effectiveness of FineTrack and reflects the importance of positive and negative samples balanced and discrete distribution of training data for training Re-ID.

Component-wise Analysis. In this part, we verified the effectiveness of FAFPN and MPMG through ablation experiments. As shown in Tab. 4, FineTrack effectively improves the performance of appearance embeddings. Because YOLOX [11] cannot extract Re-ID features, we use FPN [21] to aggregate multi-scale feature maps and then adopt Global Average Pooling (GAP) to obtain appearance embeddings. Finally, we combine FPN and GAP as our

Method	Components				Metrics	
	FPN	FAFPN	GAP	MPMG	Rank-1 \uparrow	mAP \uparrow
1 Baseline	✓		✓		88.3	50.4
2	✓			✓	89.8	55.7
3		✓	✓		91.8	58.2
4 FineTrack		✓		✓	92.3	61.8

Table 4. **Ablation study of the components in FineTrack.** The first row is our baseline, and the last row is the metrics that can be obtained by adopting our proposed FAFPN and MPMG. (GAP: Global Average Pooling, FAFPN: Flow Alignment FPN, MPMG: Multi-head Part Mask Generator)

Method	Components		Metrics	
	GAP	MPMG	Rank-1 \uparrow	mAP \uparrow
1 FineTrack(w/o FAM)	✓		91.7	57.9
2 FineTrack(w/o FAM)		✓	92.0	61.1
3 FineTrack(w FAM)	✓		91.8	58.2
4 FineTrack(w FAM)		✓	92.3	61.8

Table 5. **Ablation study of FAM.** Re-ID embedding can be extracted through GAP or MPMG. We verify the effectiveness of FAM in these two cases. (FAM: Flow Alignment Module)

baseline (row 1). Compared with baseline indicators, our method (row 4) outperforms the baseline by 4% on Rank-1 and 11.4% on mAP, demonstrating the value of exploring fine-grained appearance representation.

Then, we analyze the effectiveness of each module separately. For multi-scale feature maps aggregation, replacing FPN with FAFPN (row 3) can achieve improvement (3.5% on Rank-1 and 7.8% on mAP) when generating Re-ID embeddings with GAP. Furthermore, when MPMG replaces GAP, using FAFPN (row 4) is better than using FPN (row 2) by 2.5% on Rank-1 and 6.1% on mAP. This indicates that FAFPN captures feature context information more effectively than FPN. At the same time, employing MPMG to extract Re-ID embeddings (row 2) is better than the baseline (1.5% on Rank-1 and 5.3% on mAP). When FAFPN replaces FPN, MPMG is 0.5% and 3.6% higher than GAP on Rank-1 and mAP. MPMG focuses on the discriminative details of the target and can describe the appearance of targets more comprehensively than GAP.

Analyze of FAM. In this part, we conducted ablation experiments with or without FAM under the effect of GAP or MPMG, as reported in Tab. 5. Specifically, when GAP is used to generate Re-ID embeddings, employing FAM (row 1 and row 3) can increase 0.1% on Rank-1 and 0.3% on mAP. At the same time, if we adopt MPMG to extract Re-ID embeddings (row 2 and row 4), FAM can get the improvement (0.3% on Rank-1 and 0.7% on mAP), which is significantly better than adding FAM when using GAP. GAP is a global representation that aggregates all information indis-

Method	MOTA \uparrow	IDF1 \uparrow	IDs \downarrow
1 Baseline	75.5	77.8	253
2 Baseline+FAFPN	76.0	78.9	459
3 Baseline+MPMG	76.4	79.7	451
4 Baseline+FAFPN+MPMG	77.0	81.1	115

Table 6. **Ablation study of our components when tracking.** The same association strategy is adopted for our proposed module to obtain the tracking results.

criminally. On the contrary, MPMG generates part masks to distinguish features explicitly and can guide FAM to perform pixel-level alignment more reasonably, even without mask label supervision.

Analyze of our components when tracking. We adopt the same association strategy to prove the effectiveness of different components in FineTrack when tracking. As shown in Tab. 6, using FAFPN or MPMG alone also improves the tracking performance. When combining these two modules, there is a significant advantage over the baseline (1.5% on MOTA, 3.3% on IDF1, and the IDs decreases from 253 to 115), demonstrating the effectiveness of FineTrack.

5. Conclusion

In this work, we have argued that diverse fine-grained representation is essential for MOT. However, existing methods focus on coarse-grained global features, which are extremely sensitive to noise. To effectively filter irrelevant information, we propose to explore diverse fine-grained appearance representation to obtain more comprehensive embeddings. As reported in the ablation studies, our presented FAFPN has great advantages in terms of aligning semantic and aggregating contextual information. Meanwhile, our constructed MPMG can effectively focus on different parts of the target without label supervision. In the training phase, we propose the SGS training strategy to improve the model performance effectively. We have verified the effectiveness of our proposed FineTrack on three public benchmarks (MOT17, MOT20, and DanceTrack) and achieved state-of-the-art performance. The experiment results indicate that diverse fine-grained representation can significantly improve the performance of Re-ID in MOT. We hope this work can be a new solution for appearance representation to generate discriminative identity embeddings.

Acknowledgements

This work was supported by the National Key Laboratory Foundation of Science and Technology on Multispectral Information Processing under Grant No. 6142113220208, Science and Technology Plan Project of Jinan under Grant No. 202214002, and National Natural Science Foundation of China under Grant No. 61105006.

References

- [1] Philipp Bergmann, Tim Meinhardt, and Laura Leal-Taixe. Tracking without bells and whistles. In *ICCV*, pages 941–951, 2019. [2](#)
- [2] Keni Bernardin and Rainer Stiefelhagen. Evaluating multiple object tracking performance: the clear mot metrics. *EURASIP J Image Video Process*, 2008:1–10, 2008. [6](#)
- [3] Alex Bewley, Zongyuan Ge, Lionel Ott, Fabio Ramos, and Ben Upcroft. Simple online and realtime tracking. In *ICIP*, pages 3464–3468. IEEE, 2016. [1](#), [2](#)
- [4] Erik Bochinski, Volker Eiselein, and Thomas Sikora. High-speed tracking-by-detection without using image information. In *AVSS*, pages 1–6. IEEE, 2017. [1](#)
- [5] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, pages 213–229. Springer, 2020. [1](#)
- [6] Peng Chu, Jiang Wang, Quanzeng You, Haibin Ling, and Zicheng Liu. Transmot: Spatial-temporal graph transformer for multiple object tracking. *arXiv preprint arXiv:2104.00194*, 2021. [6](#)
- [7] 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. *arXiv preprint arXiv:2003.09003*, 2020. [6](#)
- [8] Andreas Ess, Bastian Leibe, Konrad Schindler, and Luc Van Gool. A mobile vision system for robust multi-person tracking. In *CVPR*, pages 1–8. IEEE, 2008. [6](#)
- [9] Kuan Fang, Yu Xiang, Xiao Cheng Li, and Silvio Savarese. Recurrent autoregressive networks for online multi-object tracking. In *WACV*, pages 466–475. IEEE, 2018. [2](#)
- [10] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Detect to track and track to detect. In *ICCV*, pages 3038–3046, 2017. [2](#)
- [11] Zheng Ge, Songtao Liu, Feng Wang, Zeming Li, and Jian Sun. Yolox: Exceeding yolo series in 2021. *arXiv preprint arXiv:2107.08430*, 2021. [1](#), [3](#), [7](#)
- [12] Shoudong Han, Piao Huang, Hongwei Wang, En Yu, Donghaisheng Liu, and Xiaofeng Pan. Mat: Motion-aware multi-object tracking. *Neurocomputing*, 476:75–86, 2022. [6](#)
- [13] Shuting He, Hao Luo, Pichao Wang, Fan Wang, Hao Li, and Wei Jiang. Transreid: Transformer-based object re-identification. In *ICCV*, pages 15013–15022, 2021. [5](#)
- [14] Zilong Huang, Yunchao Wei, Xinggang Wang, Wenyu Liu, Thomas S Huang, and Humphrey Shi. Alignseg: Feature-aligned segmentation networks. *TPAMI*, 44(1):550–557, 2021. [3](#)
- [15] Rudolph Emil Kalman. A new approach to linear filtering and prediction problems. 1960. [2](#)
- [16] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [7](#)
- [17] Harold W Kuhn. The hungarian method for the assignment problem. *Nav. Res. Logist.*, 2(1-2):83–97, 1955. [2](#)
- [18] Wei Li, Yuanjun Xiong, Shuo Yang, Mingze Xu, Yongxin Wang, and Wei Xia. Semi-tcl: Semi-supervised track contrastive representation learning. *arXiv preprint arXiv:2107.02396*, 2021. [6](#)
- [19] Xiangtai Li, Ansheng You, Zhen Zhu, Houlong Zhao, Maoke Yang, Kuiyuan Yang, Shaohua Tan, and Yunhai Tong. Semantic flow for fast and accurate scene parsing. In *ECCV*, pages 775–793. Springer, 2020. [3](#)
- [20] Chao Liang, Zhipeng Zhang, Xue Zhou, Bing Li, Shuyuan Zhu, and Weiming Hu. Rethinking the competition between detection and reid in multiobject tracking. *TIP*, 31:3182–3196, 2022. [2](#), [3](#), [6](#)
- [21] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, pages 2117–2125, 2017. [2](#), [3](#), [7](#)
- [22] Jonathon Luiten, Aljosa Osep, Patrick Dendorfer, Philip Torr, Andreas Geiger, Laura Leal-Taixé, and Bastian Leibe. Hota: A higher order metric for evaluating multi-object tracking. *IJCV*, 129(2):548–578, 2021. [6](#)
- [23] Anton Milan, Laura Leal-Taixé, Ian Reid, Stefan Roth, and Konrad Schindler. Mot16: A benchmark for multi-object tracking. *arXiv preprint arXiv:1603.00831*, 2016. [6](#)
- [24] 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*, pages 6308–6318, 2020. [2](#), [6](#)
- [25] 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*, pages 164–173, 2021. [6](#), [7](#)
- [26] 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*, pages 145–161. Springer, 2020. [2](#), [6](#)
- [27] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *NeurIPS*, 28, 2015. [1](#)
- [28] Ergys Ristani, Francesco Solera, Roger Zou, Rita Cucchiara, and Carlo Tomasi. Performance measures and a data set for multi-target, multi-camera tracking. In *ECCV*, pages 17–35. Springer, 2016. [6](#)
- [29] Chaobing Shan, Chunbo Wei, Bing Deng, Jianqiang Huang, Xian-Sheng Hua, Xiaoliang Cheng, and Kewei Liang. Tracklets predicting based adaptive graph tracking. *arXiv preprint arXiv:2010.09015*, 2020. [6](#)
- [30] Shuai Shao, Zijian Zhao, Boxun Li, Tete Xiao, Gang Yu, Xiangyu Zhang, and Jian Sun. Crowdhuman: A benchmark for detecting human in a crowd. *arXiv preprint arXiv:1805.00123*, 2018. [6](#)
- [31] Bing Shuai, Andrew Berneshawi, Xinyu Li, Davide Modolo, and Joseph Tighe. Siammot: Siamese multi-object tracking. In *CVPR*, pages 12372–12382, 2021. [2](#), [6](#)
- [32] Peize Sun, Jinkun Cao, Yi Jiang, Zehuan Yuan, Song Bai, Kris Kitani, and Ping Luo. Dancetrack: Multi-object track-

- ing in uniform appearance and diverse motion. In *CVPR*, pages 20993–21002, 2022. 6
- [33] 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, 7
- [34] ShiJie Sun, Naveed Akhtar, HuanSheng Song, Ajmal Mian, and Mubarak Shah. Deep affinity network for multiple object tracking. *TPAMI*, 43(1):104–119, 2019. 2, 6
- [35] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *ICCV*, pages 9627–9636, 2019. 1
- [36] Pavel Tokmakov, Jie Li, Wolfram Burgard, and Adrien Gaidon. Learning to track with object permanence. In *ICCV*, pages 10860–10869, 2021. 6
- [37] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 30, 2017. 2, 4
- [38] Paul Voigtlaender, Michael Krause, Aljosa Osep, Jonathon Luiten, Berin Balachandar Gnana Sekar, Andreas Geiger, and Bastian Leibe. Mots: Multi-object tracking and segmentation. In *CVPR*, pages 7942–7951, 2019. 1
- [39] Qiang Wang, Yun Zheng, Pan Pan, and Yinghui Xu. Multiple object tracking with correlation learning. In *CVPR*, pages 3876–3886, 2021. 6
- [40] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, pages 7794–7803, 2018. 4
- [41] Yongxin Wang, Kris Kitani, and Xinshuo Weng. Joint object detection and multi-object tracking with graph neural networks. In *ICRA*, pages 13708–13715. IEEE, 2021. 6
- [42] Zhongdao Wang, Liang Zheng, Yixuan Liu, Yali Li, and Shengjin Wang. Towards real-time multi-object tracking. In *ECCV*, pages 107–122. Springer, 2020. 2
- [43] Nicolai Wojke, Alex Bewley, and Dietrich Paulus. Simple online and realtime tracking with a deep association metric. In *ICIP*, pages 3645–3649. IEEE, 2017. 2, 5
- [44] Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. Cbam: Convolutional block attention module. In *ECCV*, pages 3–19, 2018. 4
- [45] 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*, pages 12352–12361, 2021. 7
- [46] Yihong Xu, Yutong Ban, Guillaume Delorme, Chuang Gan, Daniela Rus, and Xavier Alameda-Pineda. Transcenter: Transformers with dense queries for multiple-object tracking. *arXiv preprint arXiv:2103.15145*, 2021. 6
- [47] Fan Yang, Xin Chang, Sakriani Sakti, Yang Wu, and Satoshi Nakamura. Remot: A model-agnostic refinement for multiple object tracking. *IMAVIS*, 106:104091, 2021. 6
- [48] En Yu, Zhuoling Li, and Shoudong Han. Towards discriminative representation: Multi-view trajectory contrastive learning for online multi-object tracking. In *CVPR*, pages 8834–8843, 2022. 1, 6
- [49] En Yu, Zhuoling Li, Shoudong Han, and Hongwei Wang. Relationtrack: Relation-aware multiple object tracking with decoupled representation. *TMM*, 2022. 2, 6
- [50] 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*, pages 36–42. Springer, 2016. 1, 2
- [51] Fangao Zeng, Bin Dong, Yuang Zhang, Tiancai Wang, Xiangyu Zhang, and Yichen Wei. Motr: End-to-end multiple-object tracking with transformer. In *ECCV*, pages 659–675. Springer, 2022. 6
- [52] Shanshan Zhang, Rodrigo Benenson, and Bernt Schiele. Citypersons: A diverse dataset for pedestrian detection. In *CVPR*, pages 3213–3221, 2017. 6
- [53] Yang Zhang, Hao Sheng, Yubin Wu, Shuai Wang, Wei Ke, and Zhang Xiong. Multiplex labeling graph for near-online tracking in crowded scenes. *IEEE Internet Things J.*, 7(9):7892–7902, 2020. 6
- [54] Yifu Zhang, Peize Sun, Yi Jiang, Dongdong Yu, Fucheng Weng, Zehuan Yuan, Ping Luo, Wenyu Liu, and Xinggang Wang. Bytetrack: Multi-object tracking by associating every detection box. In *ECCV*, pages 1–21. Springer, 2022. 1, 5, 6, 7
- [55] 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. 2, 6, 7
- [56] Linyu Zheng, Ming Tang, Yingying Chen, Guibo Zhu, Jinqiao Wang, and Hanqing Lu. Improving multiple object tracking with single object tracking. In *CVPR*, pages 2453–2462, 2021. 6
- [57] Xingyi Zhou, Vladlen Koltun, and Philipp Krähenbühl. Tracking objects as points. In *ECCV*, pages 474–490. Springer, 2020. 2, 7
- [58] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *arXiv preprint arXiv:1904.07850*, 2019. 1
- [59] Zongwei Zhou, Junliang Xing, Mengdan Zhang, and Weiming Hu. Online multi-target tracking with tensor-based high-order graph matching. In *ICPR*, pages 1809–1814. IEEE, 2018. 2