

GIAOTracker: A comprehensive framework for MCMOT with global information and optimizing strategies in VisDrone 2021

Yunhao Du¹, Junfeng Wan¹, Yanyun Zhao^{1,2}, Binyu Zhang¹, Zhihang Tong¹, Junhao Dong¹

¹Beijing University of Posts and Telecommunications

²Beijing Key Laboratory of Network System and Network Culture, China

{dyh.bupt, wanjunfeng, zyy, zhangbinyu, tongzh, djh1999}@bupt.edu.cn

Abstract

*In recent years, algorithms for multiple object tracking tasks have benefited from great progresses in deep models and video quality. However, in challenging scenarios like drone videos, they still suffer from problems, such as small objects, camera movements and view changes. In this paper, we propose a new multiple object tracker, which employs **Global Information And some Optimizing strategies**, named GIAOTracker. It consists of three stages, i.e., on-line tracking, global link and post-processing. Given detections in every frame, the first stage generates reliable tracklets using information of camera motion, object motion and object appearance. Then they are associated into trajectories by exploiting global clues and refined through four post-processing methods. With the effectiveness of the three stages, GIAOTracker achieves state-of-the-art performance on the VisDrone MOT dataset and wins the 2nd place in the VisDrone2021 MOT Challenge.*

Keywords Multiple object tracking · Drone videos · Multi-stages

1. Introduction

Drones (or general UAVs) equipped with cameras have been widely applied to various fields, e.g., agriculture, meteorology, aerial photography, fast delivery and surveillance [1]. Consequently, drone video understanding is receiving increasingly attention. Multiple Object Tracking (MOT), which aims to identify and track one or multiple categories of objects, is the key component in autonomous drone systems. However, it suffers from problems, including large number of small objects by aerial capturing, irregular object motion, camera movements, variant views, occlusion from trees and bridges, etc. It is more complex and difficult to solve these problems than those in camera fixed scenarios, e.g., surveillance videos, which makes MOT in drone videos still a challenging task.

In this paper, we present a comprehensive framework with global information and optimizing strategies for Multi-Class Multi-Object Tracking (MCMOT) in drone videos, named GIAOTracker. To alleviate detection noises, a new feature storage and update strategy EMA Bank is proposed to maintain both variant feature states and information of feature changes simultaneously. As for object motion modeling, linear Kalman filter algorithm is widely used [7, 47, 69, 71, 79], which simply sets a uniform measurement noise scale to all objects without considering detection quality. To obtain more accurate motion state, we propose a Noise Scale Adaptive Kalman algorithm (NSA Kalman) which adaptively modulates the noise scale according to the quality of object detection.

Trajectory global information plays an important role in solving the fracture problem. Global information is not well exploited in many recent MOT works [7–9, 69, 71, 73, 79]. Instead, we introduce a global link stage to associate tracklets into trajectories. Specifically, to reduce noises caused by occlusion and view changes, we propose a novel tracklet appearance feature extractor GIModel (Global Information model), which extracts both global and part spatial features in each frame and fuses them with a self-attention based temporal modeling block for more robust representation.

For MCMOT task in complex scenes, e.g., UAV videos, reasonable post-processing strategies could greatly improve the tracking performance. However, only a few works focus on the post-processing procedure [2, 16, 51]. To refine tracking results in a more comprehensive way, we propose to use four post-processing methods. To remove redundant trajectories caused by duplicate detections, a temporal-IoU based NMS (Non-Maximum Suppression) between two trajectories is used. Then, missing detections are interpolated linearly into the trajectory gaps as in TPM [51]. Considering that longer trajectories tend to be more accurate, we use a length-dependent coefficient to rescore trajectories frame by frame. Last but not least, to the best of our knowledge, few works explore the fusion strategy for the MOT task. Inspired by SoftNMS [11], we introduce TrackNMS to

fuse different tracking results, which significantly improves tracking performance.

On the VisDrone MOT test-challenge dataset, GIAOTracker achieves **52.55** mAP with detections generated by DetectoRS [53]. After fusing two tracking results (46.66 mAP & 52.55 mAP), we achieve **54.18** mAP and win the 2nd place in the VisDrone2021 MOT Challenge. Furthermore, with annotation detections as the input, our tracking performance improves by 92.4% (from 43.12 to 82.97 mAP) on the test-dev dataset, which proves the effectiveness of our tracking framework.

The main contributions of this article are summarized as follows:

- For MCMOT in drone videos, we present a comprehensive framework with global information and some optimizing strategies (GIAOTracker), which consists of three stages, i.e., online tracking, global link and post-processing.
- We propose EMA Bank strategy and NSA Kalman algorithm, which aim at more accurate and robust association.
- We introduce a tracklet feature extractor GIModel, which extracts frame-level global-part features and then fuses them with a self-attention based temporal modeling block.
- We explore a series of reasonable and effective post-processing strategies, including trajectory denoising, detection box interpolation, trajectory re-scoring and tracking model fusion.

2. Related works

2.1. SDE and JDE

Most recent MOT methods could be classified into two categories: “Separate Detection and Embedding (SDE)” and “Joint Detection and Embedding (JDE)” [69]. SDE, which is also termed as tracking-by-detection [7, 20, 37, 45, 52, 68, 71, 73], consists of the following two steps: 1) detection, in which all objects are localized and classified in sequences [3, 10, 13, 27, 28, 54–57, 65]; 2) association, where detections belonging to the same object are associated by assigning the same ID [12, 14, 15, 17, 22, 34, 38, 44, 46, 74, 76, 78]. SDE strategy optimizes detection and embedding separately, which is more flexible and suitable for complex scenarios. However, it tends to cost much time in inference. Instead, JDE incorporates detection model and embedding into a unified framework [6, 23, 39, 48, 50, 60, 63, 69, 79, 82]. It typically modifies detectors, e.g., Faster R-CNN [57], CenterNet [83], YOLOv3 [56] by adding a predictor [6, 29, 77, 82] or an embedding branch [69, 79] and leverages them to implement detection and tracking jointly. Generally, JDE meth-

ods usually behave better and faster than SDE in common applications. Whereas, they would fail when applied to more sophisticated scenarios.

In this paper, in light of the complexity of drone videos, our GIAOTracker follows the SDE paradigm. It allows us to train the detector independently, which could generate more accurate localization and classification results than JDE paradigm.

2.2. Online and Offline

We could also divide MOT methods into online tracking and offline tracking methods on whether using global information. Online methods perform association on the fly without knowledge of future information [7–9, 69, 71, 73, 79]. Most recent MOT methods are online and achieve state-of-the-art performance. Besides, compared with offline methods, online tracking has more application scenarios, e.g., real-time tracking system. Offline methods, on the other hand, are allowed to employ future frames and tend to result in better tracking quality [26, 33, 42, 49, 66, 67].

Inspired by the hierarchical data association strategy [34, 75], our GIAOTracker includes three stages: online tracking, global link and post-processing. The first stage (GIAOTracker-Online) performs online tracking to generate reliable tracklets and the second stage associates them into trajectories with global information. This hierarchical framework gives a tradeoff between accuracy and flexibility, as one could only apply GIAOTracker-Online for those online tracking scenarios, or use the full GIAOTracker method for better performance. Moreover, the global link and post-processing stages are both plug-and-play, which can be plugged into any existing MOT frameworks easily.

2.3. MOT in Drone

MOT in drone videos is a challenging task due to small objects, camera movements, variant views, etc. VisDrone2018 dataset is proposed in [84], which focuses on core problems in computer vision. VisDrone-VDT2018 [85], VisDrone-MOT2019 [70] and VisDrone-MOT2020 [21] propose abundant methods which greatly improve the ability of intelligent system to understand drone videos.

V-IOU Tracker [9] improves IOU Tracker [8] by visual tracking [31, 35] to continue a track if no detection is available, which achieves state-of-the-art performance at high processing speeds in VisDrone-VDT2018 Challenge. However, it doesn’t take camera movements into account and global information is not well exploited. Thus, we use ORB [58] and RANSAC [24] to deal with camera movements and leverage a novel tracklet feature extractor GIModel to implement associations among tracklets. HMTT [47] provides a hierarchical multi-target tracker based on detection for drone vision, where four stages are proposed to deal with multiple problems like variant views, unreliable detec-

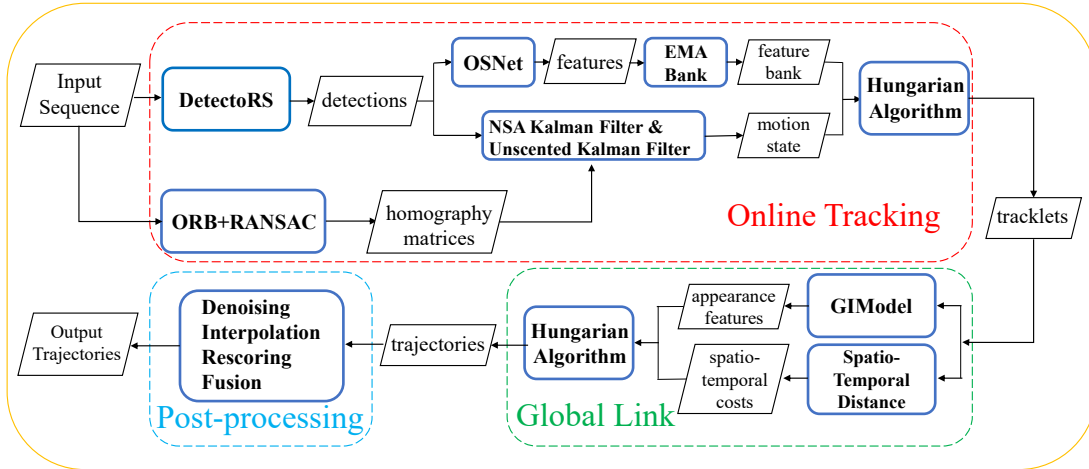


Figure 1. Overview of our proposed GIAOTracker pipeline for MCMOT

tions and missing detections. But it pays little attention on post-processing. In contrast, we highlight the importance of post-processing procedure and apply four methods to refine tracking results. COFE [2] proposes a coarse-to-fine tracking framework to reduce the classification noises and wins the first place in VisDrone-MOT2020 Challenge. We improve it by replacing the “hard-vote” mechanism with “soft-vote” and employing global information to further improve tracking performance.

3. Method

We aim at MCMOT in drone videos in a hierarchical data association way. Figure 1 illustrates our GIAOTracker framework built upon SDE paradigm, which consists of three stages, i.e., online tracking, global link and post-processing.

3.1. Online Tracking

DeepSORT [71] is a typical and strong MOT algorithm following tracking-by-detection paradigm, which generates detections first and then associates them with object motion information and appearance features (top in Figure 2). Particularly, all modules in DeepSORT, i.e. detector, Kalman and feature extractor, are plug-and-play, which allows flexibility to improve it. Consequently, the online tracking stage of our GIAOTracker applies DeepSORT as baseline. Figure 2 shows the comparison between DeepSORT and our GIAOTracker-Online. Taking sequences as input, we use DetectoRS [53] to generate detections $\{b^t\}_{t=1}^T$ frame by frame and then link them into tracklets $\{tl_n\}_{n=1}^N$. To deal with camera movement, we utilize ORB [58] and RANSAC [24] to fastly align inter-frame images. Then we improve the baseline from two aspects, i.e., object appearance and object motion. In order to obtain more robust appearance

features, we replace the simple feature extractor in DeepSORT with OSNet [81] and train it on VisDrone dataset. A new feature storage and update strategy EMA Bank is proposed to achieve more accurate association between tracklets and detections.

EMA Bank. There exist two mainstream feature storage and update approaches. DeepSORT [71] implements a feature bank to store raw features of previous L_b detections and use them to calculate the minimum cosine distance with detection features. For tracklet tl_i , its feature bank is $FB_i = \{\mathbf{f}_i^t\}_{t=1}^T$, where \mathbf{f}_i^t is the raw detection feature. Such mechanism maintains variant feature states of one tracklet and is robust to sudden changes of object appearance. However, simply using raw features is sensitive to detection noises. Instead, JDE/FairMOT [69, 79] use an EMA (Exponential Moving Average) feature update strategy, in which only one feature is maintained for every tracklet. For tracklet tl_i , its feature state e_i is updated by:

$$e_i^t = \alpha e_i^{t-1} + (1 - \alpha) \mathbf{f}_i^t \quad (1)$$

where \mathbf{f}_i^t is the appearance embedding of the detection in frame t and α is the momentum term. This incremental feature update strategy leverages the information of inter-frame feature changes and could depress detection noises. To integrate the advantages of both approaches above, we explore an intuitive method, named EMA Bank. For tracklet tl_i , its bank is $EB_i = \{e_i^t\}_{t=1}^T$, where e_i^t is calculated by equation (1). The proposed EMA Bank simultaneously takes multi-frame information and inter-frame change information into account, which is more suitable for complex scenarios.

In online tracking framework, motion prediction is another key module, in which Kalman filter [36] is commonly used. For vehicle objects, we apply Unscented Kalman Filter (UKF) algorithm [64] which is more robust for nonlinear

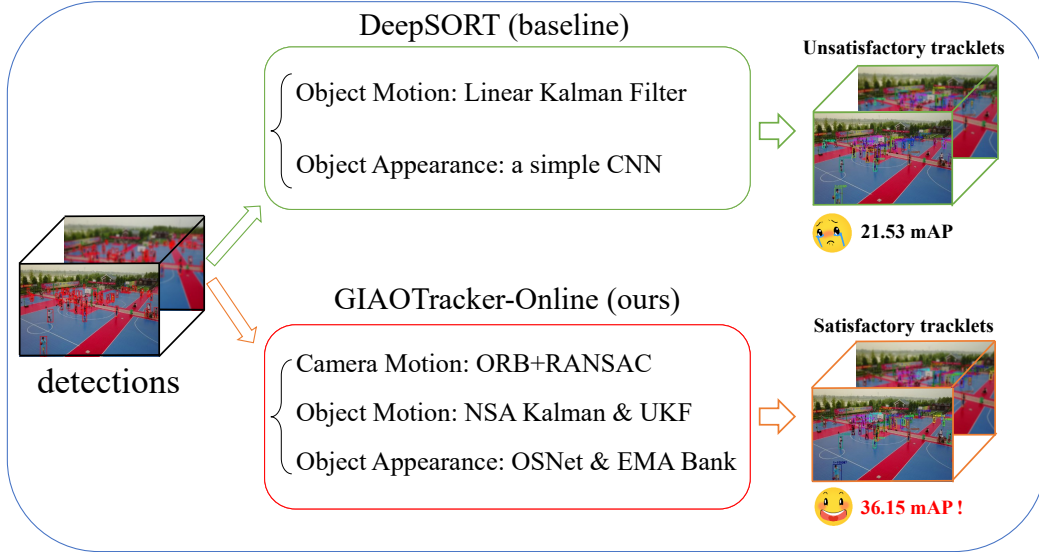


Figure 2. Comparison between DeepSORT (baseline) and GIAOTracker-Online (ours).

motion. We also propose a modified Kalman filter algorithm named NSA Kalman Filter, which could adaptively modulate the noise scale during the state update procedure.

Algorithm 1: NSA Kalman filter (state update step at state k)

Input: Measurement z_k ; Measurement confidence c_k ;

Predicted state estimate $\hat{x}_{k|k-1}$;

Predicted estimate covariance $P_{k|k-1}$;

The observation model H_k ;

The the measurement noise covariance R_k ;

Output: Updated state estimate $\hat{x}_{k|k}$;

Updated estimate covariance $P_{k|k}$;

1: $\tilde{y}_k = z_k - H_k \hat{x}_{k|k-1}$; (measurement pre-fit residual)

2: $\tilde{R}_k = (1 - c_k)R_k$; (NSA covariance)

3: $S_k = H_k P_{k|k-1} H_k^T + \tilde{R}_k$; (pre-fit residual covariance)

4: $K_k = P_{k|k-1} H_k^T S_k^{-1}$; (optimal Kalman gain)

5: $\hat{x}_{k|k} = \hat{x}_{k|k-1} + K_k \tilde{y}_k$; (updated state estimate)

6: $P_{k|k} = (I - K_k H_k) P_{k|k-1}$; (updated estimate covariance)

NSA Kalman. In DeepSORT, Kalman filter based on linear motion hypothesis is used to model objects motion. It consists of state estimation step and state update step. In the first step, Kalman filter produces estimates of current state variables, along with their uncertainties. Then these estimates are updated with a weighted average of the estimated state and the measurement. Specifically, it uses the measurement noise covariance R_* to represent the measurement (i.e., detections in the current frame) noise scale. A

larger noise scale means a smaller weight of the measurement during state update step, since its larger uncertainty. In Kalman algorithm [36], the noise scale is a constant matrix. However, intuitively different measurement contains different scales of noise. In substance, measurement noise scale should vary with detection confidence. Therefore, we propose a formula to adaptively calculate the noise covariance, named NSA noise covariance \tilde{R}_k :

$$\tilde{R}_k = (1 - c_k)R_k \quad (2)$$

where R_k is the preset constant measurement noise covariance and c_k is the detection confidence score at state k . The whole state update of our NSA Kalman filter is shown in the Algorithm 1, where the NAS step is marked with a red dotted box. Experimental results show that it significantly improves the tracking performance, though our NAS Kalman is simple (Table 1).

Rough2Fine. We adopt soft voting method to classify tracklets for GIAOTracker-Online. It's well known that some object categories are difficult to distinguish in Vis-Drone, e.g., car & van. Therefore, instead of tracking different categories independently, we follow the Coarse-to-Fine pipeline in COFE [2] and improve it with a "soft-vote" mechanism (named Rough2Fine). As in COFE, we first implement rough-class tracking and then determine fine classes of trajectories with voting mechanism. Different from the "hard-vote" in COFE, our "soft-vote" mechanism assigns multiple classes to a single trajectory, in which the voting weight is positively correlated with confidence score. Experiments show that "soft-vote" is more robust to classification errors than "hard-vote" (Table 1).

Based on the above camera correction, object motion

prediction and appearance feature processing strategies, we associate object detections to form reliable tracklets in an online tracking way. Though it has improved the baseline by a large margin, we argue that future information is not exploited by our GIAOTracker-Online, so its performance is still limited. In the next section, we'll introduce a global link algorithm which links tracklets into trajectories by employing global information.

3.2. Global link

Global link stage links short tracklets to long trajectories based on Hungarian algorithm [40]. In order to make full use of global information for tracklet association, we integrate appearance and spatio-temporal distances of tracklets into a single matching cost for Hungarian algorithm. In this section, we first introduce our appearance feature extractor GIModel for tracklets, and then present the tracklet association algorithm.

GIModel. We build our GIModel based on ResNet50-TP [25] and improve it by adding part-level features and self-attention based temporal modeling [62]. Figure 3 shows the framework of our GIModel. Taking consecutive N frames of a tracklet clip as input, GIModel extracts frame-level features first and outputs the clip features by temporal modeling. Different from the baseline who only extracts spatial global features (Figure 3 (a)), we add part-level features supervised by additional triplet loss [32] as shown in Figure 3 (b). This enables GIModel to focus on detailed features of different parts of objects and be more robust to occlusion. In inference, global and part features are directly concatenated. As for temporal modeling, instead of simply fusing frame-level features by average pooling (Fig. 2 (c)), we implement inter-frame information interaction with a Transformer encoder layer [62] showed by Figure 3 (d) in order to aggregate features of multi-frames and suppress noises for one clip. Given N frames of clip j , the self-attention based features $\{\hat{\mathbf{f}}_j^t\}_{n=1}^N$ are obtained by a Transformer encoder layer:

$$\{\hat{\mathbf{f}}_j^t\}_{n=1}^N = TEL(\{\mathbf{f}_j^t\}_{n=1}^N) \quad (3)$$

where \mathbf{f}_j^t represents the raw frame-level features and $TEL(\cdot)$ represents the Transformer encoder layer. Then the 3D feature $\tilde{\mathbf{f}}_j$ for clip j is calculated by:

$$\tilde{\mathbf{f}}_j = \frac{1}{N} \sum_{n=1}^N \hat{\mathbf{f}}_j^n \quad (4)$$

The feature bank for tracklet tl_i is the set of its clip features $\tilde{F}_i = \{\tilde{\mathbf{f}}_j^i\}_{j=1}^M$.

To conclude, part features focus on detailed spatial information and self-attention based temporal modeling aggregates temporal context information more effectively. Experimental results (Table 2) show that our GIModel performs better than baseline by a large margin.

Association. We calculate the matching cost matrix used by Hungarian algorithm with appearance cost and spatio-temporal cost for tracklet association. Specifically, given tl_i and tl_j , their appearance cost is the smallest cosine distance between their feature banks \tilde{F}_i and \tilde{F}_j :

$$C_a(i, j) = \min\{1 - \tilde{\mathbf{f}}_i^n \tilde{\mathbf{f}}_j^m | \tilde{\mathbf{f}}_i^n \in \tilde{F}_i, \tilde{\mathbf{f}}_j^m \in \tilde{F}_j\} \quad (5)$$

Additionally, the spatio-temporal distance costs $C_s(i, j)$ and $C_t(i, j)$ measure the time interval and space distances for two tracklets respectively. If the cost meets the threshold constraint,

$$C_a(i, j) < Th_a \text{ and } C_t(i, j) < Th_t \text{ and } C_s(i, j) < Th_s \quad (6)$$

then matching cost is calculated as follows:

$$C(i, j) = \lambda_a C_a(i, j) + \lambda_t C_t(i, j) + \lambda_s C_s(i, j) \quad (7)$$

where Th_a, Th_t, Th_s are the preset thresholds of appearance feature cost, time cost and space cost respectively, and $\lambda_a, \lambda_t, \lambda_s$ are weight coefficients. Cosine distance is an effective metric for measuring appearance features. Time interval and space distances of tracklets are directly used to avoid errors in velocity estimation for object motion. In this way, based on Hungarian algorithm, we associate the tracklets to form long trajectories well.

In this section, we introduce a tracklet association algorithm based on tracklets appearance features and spatio-temporal distances. Particularly, GIModel is proposed to extract representative 3D features for tracklets, which is robust to detection noises and abrupt changes of object appearance. Next, we'll present some post-processing methods to further improve tracking accuracy.

3.3. Post-processing

In this section, we explore four post-processing procedures to refine the trajectories, i.e., denoising, interpolation, rescoring and fusion.

Denoising. There exist some duplicate detections which would result in redundant trajectories. STGT [16] uses detection-wise spatio-IoU based matching procedure to remove unmatched detection candidates. Instead, our denoising algorithm uses trajectory-wise temporal-IoU to implement SoftNMS [11] between trajectories.

Interpolation. Within one trajectory, missing detections would also decrease tracking accuracy. As in TPM [51], detections are interpolated linearly into gaps of the trajectory. Considering that larger gaps will bring more noises, only missing frames less than 60 are filled in.

Rescoring. While evaluating, the average score is used to measure the quality of trajectories. According to our observation, however, longer trajectories tend to be more accurate. Therefore, we use a length-dependent coefficient to

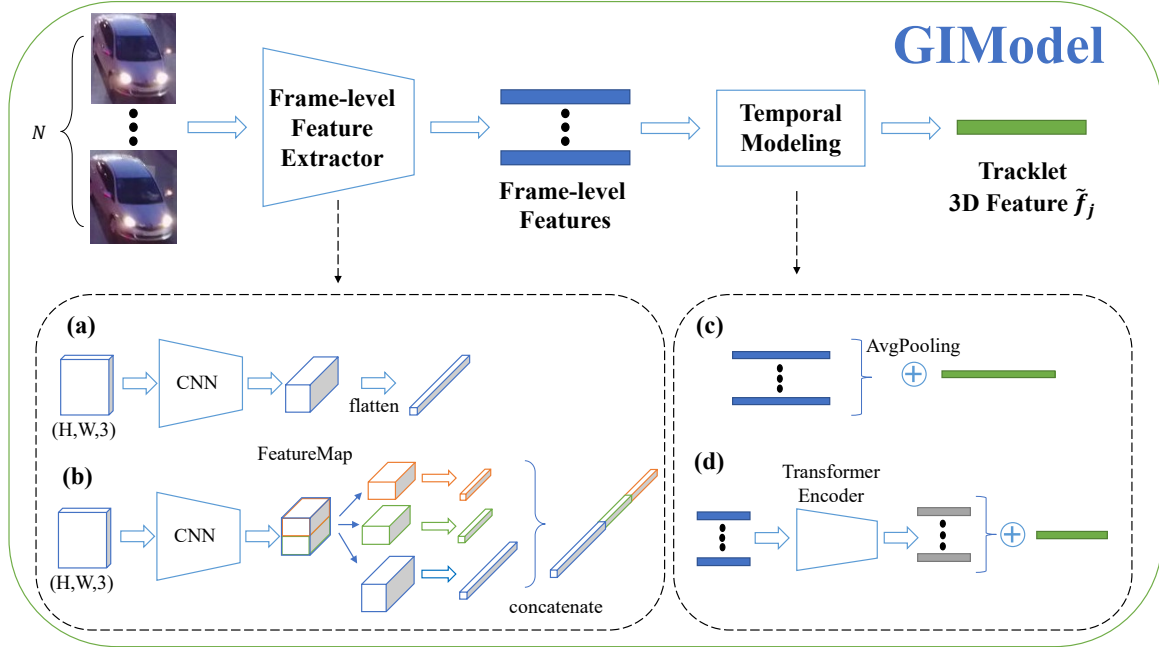


Figure 3. The framework of our GIModel. (a) Simply extracting global features based on a CNN. (b) Given CNN feature maps, extracting both global and part features. (c) Simply averaging frame-level features for temporal modeling. (d) Temporal modeling with a Transformer encoder layer. We take ResNet50-TP [25] as baseline and improve it by adding part-level features ((b) vs. (a)) and self-attention based temporal modeling ((d) vs. (c)).

rescore trajectories per frame. For trajectory i whose length is l_i , the rescoring weight is calculated as:

$$\omega_i = \frac{1 - e^{-l_i/\tau}}{1 + e^{-l_i/\tau}} \quad (8)$$

where the temperature factor $\tau = 25$. For frame j , its confidence s_i^j is rescored as $\hat{s}_i^j = \omega_i \cdot s_i^j$. In this way, the confidence of long trajectories is relatively increased while decreased of short trajectories.

Fusion. Model fusion is a common strategy to improve performance in some computer vision tasks, such as image classification [19] and object detection [59]. To the best of our knowledge, few works focus on fusion strategies for MOT task. Inspired by the success of NMS-based model fusion on object detection task, we propose TrackNMS to fuse different tracking results. In short, TrackNMS is based on the idea of SoftNMS [11] with two main differences:

1. IoU: SoftNMS uses spatio-IoU of detections to determine the degree of suppression. Instead, TrackNMS uses temporal-IoU between trajectories.
2. Sort: Unlike SoftNMS who sorts detections by their scores, TrackNMS applies the sum (instead of “mean”) of trajectory frame-level scores as the sorting basis, which means longer trajectories tend to have higher priority.

Experimental results (Table 4) show that our TrackNMS works well.

4. Experiment

4.1. Dataset and Metrics

The VisDrone MOT dataset consists of 96 sequences (39,988 frames in total). Each frame is manually labeled in high quality. Participants are asked to submit tracking results for 5 selected object classes in this challenge, including pedestrian, car, van, bus and truck. Unless otherwise stated, we use both training and validation dataset for training and use test-dev for validation.

The VisDrone2021 MOT Challenge uses the protocol in [4] to evaluate the tracking performance. Given tracking results which consist of a list of detections with confidence scores and corresponding identities, trajectories are sorted according to the average confidence of their detections. A trajectory is considered correct if the IoU overlap with the ground truth trajectory is larger than a threshold. The final ranking metric is calculated by averaging the mean average precision (mAP) per object class over different thresholds.

4.2. Implementation details

Detection. We fine-tune the ResNet50-based [30] DetectorS [53] detector on the VisDrone MOT train+val dataset, which is pre-trained on MS COCO [43] dataset. To avoid

Method	add	mAP	mAP-ped.	mAP-car	mAP-van	mAP-truck	mAP-bus
baseline [71]	-	21.53	14.21	30.37	20.34	22.22	20.49
+ORB	✓	29.69	14.32	51.68	26.97	25.37	30.14
+EMA Bank	✓	30.06	14.77	51.14	27.45	26.78	30.14
+NSA Kalman	✓	32.30	18.59	54.84	27.55	30.37	30.14
+OSNet	✓	32.66	19.61	55.74	27.68	30.12	30.14
+R2F-hard	×	32.29	19.39	55.83	30.18	29.61	26.46
+R2F-soft	✓	33.87	19.60	55.93	30.83	29.29	33.67
+UKF	✓	36.15	19.60	57.61	30.80	33.00	39.76

Table 1. Ablation studies of the online tracking stage. “✓” in the “add” line means the corresponding component is added to GIAOTracker. “R2F-hard” means “Rough2Fine” using “hard-vote” strategy and “R2F-soft” uses “soft-vote” (best in bold).

Method	mAP-P	Rank1-P	mAP-V	Rank1-V
baseline [71]	52.1	75.5	65.4	90.1
+train	76.4	89.7	80.7	94.0
+sat	80.3	91.8	83.8	95.9
+part	81.3	92.2	85.1	95.8

Table 2. Effects of different strategies for GIModel, which use ResNet50-TP [25] pretrained on ImageNet as baseline. “+train” means training it on the VisDrone2021 dataset. “+sat” represents implementing our self-attention based temporal modeling block. “+part” means adding part-level features supervised by triplet loss [32] (best in bold).

overfitting, 1 frame in every 5 frames is sampled while training. Training input scale is set to [1333, 800]. When testing, SoftNMS [11] and multi-scale testing [(1333, 800), (2000, 1200)] are used. It achieves 56.9 AP50 on the test-dev dataset, named DetV1. As the improvement, when training, we cut an image into 4 non-overlapping regions as input and use a larger input scale [1600, 1050]. When testing, we still use the entire image as input and change the input scale to [(3000, 1969), (3200, 2100), (3400, 2231), (3600, 2362), (3800, 2494), (4000, 2625)]. This improved method achieves 59.4 AP50. Then we fuse it with DetV1 using SoftNMS, whose AP50 is 63.2 on test-dev, named DetV2.

OSNet. For training and evaluating the OSNet, we create a ReID dataset based on VisDrone MOT dataset. The frame sampling rate is 5 and the minimum height/width is 32. Any object with an occlusion ratio or truncation ratio greater than 50% will be removed. The ratio of gallery to probe is set to 7:1. We use the ImageNet [18] pre-trained model and fine-tune it on our dataset. The input scale is set to [128, 256] for vehicles (i.e., car, van, truck and bus) and [256, 128] for pedestrian.

GIModel. Similar to the ReID dataset mentioned above, we also create a VideoReID-like [41, 72, 80] dataset for GIModel. The main differences are the sampling rate is set to 1 and the ratio of gallery to probe is set to 3:1. The in-

put size is set to [224, 224, 4] for vehicle tracklet clips and [224, 112, 4] for pedestrian. ImageNet [18] pre-trained model is also used. GIModel is evaluated by mAP and Rank1, which are generally used in ReID and VideoReID tasks [25, 61, 81].

Others. In EMA Bank, $L_b = 100$, $\alpha = 0.9$ as in DeepSORT [71] and JDE [69]. In the global link stage, $Th_a = 0.4$, $Th_t = 200$, $Th_s = 150$ and $\lambda_a = 40$, $\lambda_t = 1$, $\lambda_s = 1$.

4.3. Ablation

In this section, we describe the ablation results on VisDrone MOT test-dev dataset.

Online Tracking. In order to evaluate the effect of different components of our online tracking stage, we compare them with baseline [71] as shown in table 1. We have the following four observations: 1) Images matching based on ORB significantly improves the tracking performance which compensates for the camera movements. 2) Both stronger feature extractor OSNet and more robust feature storage strategy EMA Bank benefit tracking results. 3) NSA Kalman and UKF are far superior to linear Kalman filter algorithm. 4) As for Rough2Fine strategy, “soft-vote” exceeds “hard-vote” by a large margin.

GIModel. Table 2 presents the effects of different strategies evaluated on our VideoReID dataset for GIModel, which takes ResNet50-TP [25] as baseline (“P” for people, “V” for vehicle). Results show that both our self-attention based temporal modeling block and part-level features supervised by triplet loss [32] make GIModel perform better than baseline significantly.

Post-processing. We explore the influence of different post-processing procedures in table 3. Denoising and interpolation brings 0.14 mAP and 0.23 mAP gains respectively. We argue that interpolation is complementary to denoising, i.e., denoising removes redundant trajectories caused by replicate detections and interpolation restores missing detections. Furthermore, our rescoring method increases tracking accuracy by 1.64 mAP, which indicates the quality of trajectories is strongly correlated to their length.

Method	mAP	mAP-ped.	mAP-car	mAP-van	mAP-truck	mAP-bus
GIAOTracker-Global	38.71	23.63	57.95	35.19	36.35	40.43
+denoising	38.85	23.90	57.82	35.51	36.59	40.43
+interpolation	39.08	25.49	57.57	35.48	37.21	39.63
+rescoring	40.72	27.10	57.46	36.65	36.72	45.69

Table 3. The influence of different post-processing methods (best in bold).

Method	mAP	mAP-1	mAP-4	mAP-5	mAP-6	mAP-9
baseline [71]	21.53	14.21	30.37	20.34	22.22	20.49
GIAOTracker-Online	36.15	19.60	57.61	30.80	33.00	39.76
GIAOTracker-Global	38.71	23.63	57.95	35.19	36.35	40.43
GIAOTracker-Post	40.72	27.10	57.46	36.65	36.72	45.69
GIAOTracker-DetV2	43.12	37.44	58.95	39.93	35.31	43.95
GIAOTracker-Fusion	44.46	38.29	59.86	40.45	38.31	45.37
GIAOTracker*	82.97	77.69	79.72	71.15	73.33	91.67

Table 4. Overview of the results on test-dev dataset. "*" means taking annotation detections as input (best in bold).

Method	mAP
SOMOT	58.61
GIAOTracker (ours)	54.18
MMDS	52.68
Deep IoU Tracker	48.54
Yolo-Deepsort-VisDrone	46.70
CenterPointCF	44.03
MIYoT	39.35
DeepTAMA+Homography+Voting	39.25
HNet	24.71

Table 5. Top 9 results of the VisDrone2021 MOT Challenge [5].

Overview. Table 4 presents the overview results on VisDrone MOT test-dev dataset, which includes 16 sequences. Taking DetV1+DeepSORT as baseline, the second to fourth rows respectively show the tracking performance after adding three stages, i.e., online tracking, global link and post-processing (excluding fusion). The fifth row is the results after replace DetV1 with DetV2, which has better detection performance (from 56.9 to 63.2 AP50). After fusing the tracking results on the fourth and fifth row by our TrackNMS, we achieve 44.46 mAP (row 6). Moreover, the last row shows the tracking performance by taking annotation detections as input of GIAOTracker, which improves by 92% compared to our best single model results on test-dev dataset. It proves the effectiveness and potential of our GIAOTracker, whose tracking quality is greatly limited by detection accuracy.

4.4. Comparison to State-of-the-art

Finally, we compare our GIAOTracker with state-of-the-art results of the VisDrone2021 MOT Challenge. Table 5 lists the best 9 performing results and our GIAOTracker ranks 2nd, which proves the effectiveness of our framework. Note that our detector is simply trained on VisDrone train+val dataset, whose performance is limited by the amount of data and detection methods (the usage of extra data is allowed in the VisDrone2021 Challenge). As mentioned in Sec 4.3, we argue that more accurate detection results would improve the performance significantly.

5. Conclusion

In this paper, we provide a comprehensive framework for multi-class multi-object tracking in drone videos. Inspired by the hierarchical data association strategy, it consists of three stages, i.e., online tracking, global link and post-processing. The online tracking generates reliable tracklets with information including camera movements, object appearance and object motion, then they are associated into trajectories according to tracklet features and spatio-temporal distances in the global link stage. The final post-processing stage refines the tracking results through four methods. Our GIAOTracker achieves the state-of-the-art results on the VisDrone MOT dataset and ranks 2nd in the VisDrone2021 MOT Challenge.

6. Acknowledgements

This work is supported by Chinese National Natural Science Foundation under Grants (62076033, U1931202).

References

- [1] <http://aiskyeye.com/>. 1
- [2] https://www.youtube.com/watch?v=ironNC_6cHLs&t=1s. 1, 3, 4
- [3] <https://github.com/ultralytics/yolov5.2>
- [4] <http://image-net.org/challenges/LSVRC/2017>. 6
- [5] <http://www.aiskyeye.com/leaderboard/>. 8
- [6] Philipp Bergmann, Tim Meinhardt, and Laura Leal-Taixé. Tracking without bells and whistles. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 941–951, 2019. 2
- [7] 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. IEEE, 2016. 1, 2
- [8] Erik Bochinski, Volker Eiselein, and Thomas Sikora. High-speed tracking-by-detection without using image information. In *2017 14th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*, pages 1–6. IEEE, 2017. 2
- [9] Erik Bochinski, Tobias Senst, and Thomas Sikora. Extending iou based multi-object tracking by visual information. In *2018 15th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*, pages 1–6. IEEE, 2018. 1, 2
- [10] Alexey Bochkovskiy, Chien-Yao Wang, and Hong-Yuan Mark Liao. Yolov4: Optimal speed and accuracy of object detection. *arXiv preprint arXiv:2004.10934*, 2020. 2
- [11] Navaneeth Bodla, Bharat Singh, Rama Chellappa, and Larry S Davis. Soft-nms—improving object detection with one line of code. In *Proceedings of the IEEE international conference on computer vision*, pages 5561–5569, 2017. 1, 5, 6, 7
- [12] Guillem Brasó and Laura Leal-Taixé. Learning a neural solver for multiple object tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6247–6257, 2020. 2
- [13] Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: Delving into high quality object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6154–6162, 2018. 2
- [14] Jiahui Chen, Hao Sheng, Yang Zhang, and Zhang Xiong. Enhancing detection model for multiple hypothesis tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 18–27, 2017. 2
- [15] 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. 2
- [16] Peng Chu, Jiang Wang, Quanzeng You, Haibin Ling, and Zicheng Liu. Spatial-temporal graph transformer for multiple object tracking. *arXiv e-prints*, pages arXiv–2104, 2021. 1, 5
- [17] Afshin Dehghan, Yicong Tian, Philip HS Torr, and Mubarak Shah. Target identity-aware network flow for online multiple target tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1146–1154, 2015. 2
- [18] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 7
- [19] Thomas G Dietterich. Ensemble methods in machine learning. In *International workshop on multiple classifier systems*, pages 1–15. Springer, 2000. 6
- [20] Loïc Fagot-Bouquet, Romaric Audigier, Yoann Dhome, and Frédéric Lerasle. Improving multi-frame data association with sparse representations for robust near-online multi-object tracking. In *European Conference on Computer Vision*, pages 774–790. Springer, 2016. 2
- [21] Heng Fan, Dawei Du, Longyin Wen, Pengfei Zhu, Qinghua Hu, Haibin Ling, Mubarak Shah, Junwen Pan, Arne Schumann, Bin Dong, et al. Visdrone-mot2020: The vision meets drone multiple object tracking challenge results. In *European Conference on Computer Vision*, pages 713–727. Springer, 2020. 2
- [22] Kuan Fang, Yu Xiang, Xiaocheng Li, and Silvio Savarese. Recurrent autoregressive networks for online multi-object tracking. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 466–475. IEEE, 2018. 2
- [23] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Detect to track and track to detect. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3038–3046, 2017. 2
- [24] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981. 2, 3
- [25] Jiyang Gao and Ram Nevatia. Revisiting temporal modeling for video-based person reid. *arXiv preprint arXiv:1805.02104*, 2018. 5, 6, 7
- [26] Andreu Girbau, Xavier Giró-i Nieto, Ignasi Rius, and Ferran Marqués. Multiple object tracking with mixture density networks for trajectory estimation. *arXiv preprint arXiv:2106.10950*, 2021. 2
- [27] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015. 2
- [28] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 580–587, 2014. 2
- [29] Shoudong Han, Piao Huang, Hongwei Wang, En Yu, Donghaisheng Liu, Xiaofeng Pan, and Jun Zhao. Mat: Motion-aware multi-object tracking. *arXiv preprint arXiv:2009.04794*, 2020. 2
- [30] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 6
- [31] João F Henriques, Rui Caseiro, Pedro Martins, and Jorge Batista. High-speed tracking with kernelized correlation filters. *IEEE transactions on pattern analysis and machine intelligence*, 37(3):583–596, 2014. 2
- [32] Alexander Hermans, Lucas Bayer, and Bastian Leibe. In de-

- fense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737*, 2017. 5, 7
- [33] Andrea Hornakova, Roberto Henschel, Bodo Rosenhahn, and Paul Swoboda. Lifted disjoint paths with application in multiple object tracking. In *International Conference on Machine Learning*, pages 4364–4375. PMLR, 2020. 2
- [34] Chang Huang, Bo Wu, and Ramakant Nevatia. Robust object tracking by hierarchical association of detection responses. In *European Conference on Computer Vision*, pages 788–801. Springer, 2008. 2
- [35] Zdenek Kalal, Krystian Mikolajczyk, and Jiri Matas. Forward-backward error: Automatic detection of tracking failures. In *2010 20th international conference on pattern recognition*, pages 2756–2759. IEEE, 2010. 2
- [36] Rudolph Emil Kalman. A new approach to linear filtering and prediction problems. 1960. 3, 4
- [37] Margret Keuper, Siyu Tang, Bjoern Andres, Thomas Brox, and Bernt Schiele. Motion segmentation & multiple object tracking by correlation co-clustering. *IEEE transactions on pattern analysis and machine intelligence*, 42(1):140–153, 2018. 2
- [38] Chanh Kim, Fuxin Li, Arridhana Ciptadi, and James M Rehg. Multiple hypothesis tracking revisited. In *Proceedings of the IEEE international conference on computer vision*, pages 4696–4704, 2015. 2
- [39] Han-Ul Kim and Chang-Su Kim. Cdt: Cooperative detection and tracking for tracing multiple objects in video sequences. In *European Conference on Computer Vision*, pages 851–867. Springer, 2016. 2
- [40] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955. 5
- [41] Jianing Li, Jingdong Wang, Qi Tian, Wen Gao, and Shiliang Zhang. Global-local temporal representations for video person re-identification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3958–3967, 2019. 7
- [42] 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. 2
- [43] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 6
- [44] Anton Milan, S Hamid Rezatofighi, Anthony Dick, Ian Reid, and Konrad Schindler. Online multi-target tracking using recurrent neural networks. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017. 2
- [45] Anton Milan, Konrad Schindler, and Stefan Roth. Multi-target tracking by discrete-continuous energy minimization. *IEEE transactions on pattern analysis and machine intelligence*, 38(10):2054–2068, 2015. 2
- [46] Peter Ondruska and Ingmar Posner. Deep tracking: Seeing beyond seeing using recurrent neural networks. In *Thirtieth AAAI conference on artificial intelligence*, 2016. 2
- [47] Siyang Pan, Zhihang Tong, Yanyun Zhao, Zhicheng Zhao, Fei Su, and Bojin Zhuang. Multi-object tracking hierarchically in visual data taken from drones. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0, 2019. 1, 2
- [48] Bo Pang, Yizhuo Li, Yifan Zhang, Muchen Li, and Cewu Lu. Tubetk: Adopting tubes to track multi-object in a one-step training model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6308–6318, 2020. 2
- [49] 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
- [50] Jinlong Peng, Changan Wang, Fangbin Wan, Yang Wu, Yabiao Wang, Ying Tai, Chengjie Wang, Jilin Li, Feiyue Huang, and Yanwei Fu. Chained-tracker: Chaining paired attentive regression results for end-to-end joint multiple-object detection and tracking. In *European Conference on Computer Vision*, pages 145–161. Springer, 2020. 2
- [51] Jinlong Peng, Tao Wang, Weiyao Lin, Jian Wang, John See, Shilei Wen, and Erui Ding. Tpm: Multiple object tracking with tracklet-plane matching. *Pattern Recognition*, 107:107480, 2020. 1, 5
- [52] Hamed Pirsiavash, Deva Ramanan, and Charless C Fowlkes. Globally-optimal greedy algorithms for tracking a variable number of objects. In *CVPR 2011*, pages 1201–1208. IEEE, 2011. 2
- [53] Siyuan Qiao, Liang-Chieh Chen, and Alan Yuille. Detectors: Detecting objects with recursive feature pyramid and switchable atrous convolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10213–10224, 2021. 2, 3, 6
- [54] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016. 2
- [55] Joseph Redmon and Ali Farhadi. Yolo9000: better, faster, stronger. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7263–7271, 2017.
- [56] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018. 2
- [57] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28:91–99, 2015. 2
- [58] Ethan Rublee, Vincent Rabaud, Kurt Konolige, and Gary Bradski. Orb: An efficient alternative to sift or surf. In *2011 International conference on computer vision*, pages 2564–2571. Ieee, 2011. 2, 3
- [59] Roman Solovyev, Weimin Wang, and Tatiana Gabruseva. Weighted boxes fusion: Ensembling boxes from different object detection models. *Image and Vision Computing*, 107:104117, 2021. 6
- [60] ShiJie Sun, Naveed Akhtar, XiangYu Song, HuanSheng Song, Ajmal Mian, and Mubarak Shah. Simultaneous detection and tracking with motion modelling for multiple object tracking. In *European Conference on Computer Vision*, pages 626–643. Springer, 2020. 2
- [61] Yifan Sun, Liang Zheng, Yi Yang, Qi Tian, and Shengjin

- Wang. Beyond part models: Person retrieval with refined part pooling (and a strong convolutional baseline). In *Proceedings of the European conference on computer vision (ECCV)*, pages 480–496, 2018. 7
- [62] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. 5
- [63] Paul Voigtlaender, Michael Krause, Aljosa Osep, Jonathon Luiten, Berin Balachandar Gnana Sekar, Andreas Geiger, and Bastian Leibe. Mots: Multi-object tracking and segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7942–7951, 2019. 2
- [64] Eric A Wan and Rudolph Van Der Merwe. The unscented kalman filter for nonlinear estimation. In *Proceedings of the IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium (Cat. No. 00EX373)*, pages 153–158. Ieee, 2000. 3
- [65] Chien-Yao Wang, Alexey Bochkovskiy, and Hong-Yuan Mark Liao. Scaled-yolov4: Scaling cross stage partial network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13029–13038, 2021. 2
- [66] Gaoang Wang, Yizhou Wang, Renshu Gu, Weijie Hu, and Jenq-Neng Hwang. Split and connect: A universal tracklet booster for multi-object tracking. *arXiv preprint arXiv:2105.02426*, 2021. 2
- [67] Gaoang Wang, Yizhou Wang, Haotian Zhang, Renshu Gu, and Jenq-Neng Hwang. Exploit the connectivity: Multi-object tracking with trackletnet. In *Proceedings of the 27th ACM International Conference on Multimedia*, pages 482–490, 2019. 2
- [68] Xinchao Wang, Engin Türetken, Francois Fleuret, and Pascal Fua. Tracking interacting objects using intertwined flows. *IEEE transactions on pattern analysis and machine intelligence*, 38(11):2312–2326, 2015. 2
- [69] 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. 1, 2, 3, 7
- [70] Longyin Wen, Pengfei Zhu, Dawei Du, Xiao Bian, Haibin Ling, Qinghua Hu, Jiayu Zheng, Tao Peng, Xinyao Wang, Yue Zhang, et al. Visdrone-mot2019: The vision meets drone multiple object tracking challenge results. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0, 2019. 2
- [71] Nicolai Wojke, Alex Bewley, and Dietrich Paulus. Simple online and realtime tracking with a deep association metric. In *2017 IEEE international conference on image processing (ICIP)*, pages 3645–3649. IEEE, 2017. 1, 2, 3, 7, 8
- [72] Yu Wu, Yutian Lin, Xuanyi Dong, Yan Yan, Wanli Ouyang, and Yi Yang. Exploit the unknown gradually: One-shot video-based person re-identification by stepwise learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5177–5186, 2018. 7
- [73] Yu Xiang, Alexandre Alahi, and Silvio Savarese. Learning to track: Online multi-object tracking by decision making. In *Proceedings of the IEEE international conference on computer vision*, pages 4705–4713, 2015. 1, 2
- [74] Yihong Xu, Aljosa Osep, Yutong Ban, Radu Horaud, Laura Leal-Taixé, and Xavier Alameda-Pineda. How to train your deep multi-object tracker. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6787–6796, 2020. 2
- [75] Bo Yang and Ram Nevatia. Multi-target tracking by online learning of non-linear motion patterns and robust appearance models. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1918–1925. IEEE, 2012. 2
- [76] Amir Roshan Zamir, Afshin Dehghan, and Mubarak Shah. Gmcp-tracker: Global multi-object tracking using generalized minimum clique graphs. In *European conference on computer vision*, pages 343–356. Springer, 2012. 2
- [77] Jimuyang Zhang, Sanping Zhou, Xin Chang, Fangbin Wan, Jinjun Wang, Yang Wu, and Dong Huang. Multiple object tracking by flowing and fusing. *arXiv preprint arXiv:2001.11180*, 2020. 2
- [78] Li Zhang, Yuan Li, and Ramakant Nevatia. Global data association for multi-object tracking using network flows. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2008. 2
- [79] Yifu Zhang, Chunyu Wang, Xinggang Wang, Wenjun Zeng, and Wenyu Liu. Fairmot: On the fairness of detection and re-identification in multiple object tracking. *arXiv preprint arXiv:2004.01888*, 2020. 1, 2, 3
- [80] Liang Zheng, Zhi Bie, Yifan Sun, Jingdong Wang, Chi Su, Shengjin Wang, and Qi Tian. Mars: A video benchmark for large-scale person re-identification. In *European Conference on Computer Vision*, pages 868–884. Springer, 2016. 7
- [81] Kaiyang Zhou, Yongxin Yang, Andrea Cavallaro, and Tao Xiang. Omni-scale feature learning for person re-identification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3702–3712, 2019. 3, 7
- [82] Xingyi Zhou, Vladlen Koltun, and Philipp Krähenbühl. Tracking objects as points. In *European Conference on Computer Vision*, pages 474–490. Springer, 2020. 2
- [83] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *arXiv preprint arXiv:1904.07850*, 2019. 2
- [84] Pengfei Zhu, Longyin Wen, Xiao Bian, Haibin Ling, and Qinghua Hu. Vision meets drones: A challenge. *arXiv preprint arXiv:1804.07437*, 2018. 2
- [85] Pengfei Zhu, Longyin Wen, Dawei Du, Xiao Bian, Haibin Ling, Qinghua Hu, Haotian Wu, Qinqin Nie, Hao Cheng, Chenfeng Liu, et al. Visdrone-vdt2018: The vision meets drone video detection and tracking challenge results. In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, pages 0–0, 2018. 2