

SFTrack: A Robust Scale and Motion Adaptive Algorithm for Tracking Small and Fast Moving Objects

Inpyo Song¹ and Jangwon Lee¹

Abstract—This paper addresses the problem of multi-object tracking in Unmanned Aerial Vehicle (UAV) footage. It plays a critical role in various UAV applications, including traffic monitoring systems and real-time suspect tracking by the police. However, this task is highly challenging due to the fast motion of UAVs, as well as the small size of target objects in the videos caused by the high-altitude and wide-angle views of drones. In this study, we thus introduce a simple yet more effective method compared to previous work to overcome these challenges. Our approach involves a new tracking strategy, which initiates the tracking of target objects from low-confidence detections commonly encountered in UAV application scenarios. Additionally, we propose revisiting traditional appearance-based matching algorithms to improve the association of low-confidence detections. To evaluate the effectiveness of our method, we conducted benchmark evaluations on two UAV-specific datasets (VisDrone2019, UAVDT) and one general object tracking dataset (MOT17). The results demonstrate that our approach surpasses current state-of-the-art methodologies, highlighting its robustness and adaptability in diverse tracking environments. Furthermore, we have improved the annotation of the UAVDT dataset by rectifying several errors and addressing omissions found in the original annotations. We will provide this refined version of the dataset to facilitate better benchmarking in the field.

I. INTRODUCTION

Multi-object tracking (MOT) plays a crucial role in various UAV applications, ranging from real-time suspect tracking conducted by the police to interactions between humans and drones [1]–[3]. The primary goal of MOT in such applications is to accurately estimate and trace the trajectories of multiple objects in real-time video streams captured by cameras mounted on drones. However, MOT in UAV applications poses unique challenges due to the fast motion of UAVs, as well as the small size of target objects caused by the high-altitude and wide-angle views of drones. These inherent challenges can significantly hinder the overall accuracy of object tracking performance.

To address these challenges, most contemporary MOT approaches employ a fusion of motion cues and deep learning-based appearance similarity features. Nonetheless, the dynamic and unconstrained mobility characteristic of UAVs introduces notable hurdles for conventional motion cue processing techniques like the Kalman Filter [4]. Additionally, rapid changes in viewpoints and the small scale of objects can present difficulties for approaches relying on appearance similarity, including the application of deep learning-based




Fig. 1. Illustration of challenges encountered during multi-object tracking in UAVs. These primarily include irregular patterns of UAV motion, along with low-confidence detections due to small scale objects and occlusions.

re-identification matching methods. This is primarily due to the frequent occurrence of low-confidence detections for small-scale objects, resulting in higher rates of both false positives and false negatives. Furthermore, modern deep learning-based re-identification modules often struggle to obtain high-quality visual features in scenarios where small-scale or partially obscured objects are involved, primarily due to the restricted image coverage of these objects. Therefore, this paper outlines two main objectives to tackle these challenges: (1) Develop a more robust technique to compensate for the irregular motion of UAVs. (2) Effectively manage the frequent occurrence of low-confidence detections in UAV applications, primarily stemming from blurred images, small object sizes, and occlusions.

We begin by introducing a simple yet effective technique called “UAV Motion Compensation” to address our first objective. This approach mitigates UAV motion by applying adjustments to the bounding boxes used for object tracking while preserving their aspect ratios. It can improve the traditional Intersection over Union (IoU) matching by leveraging the Kalman Filter to handle erratic UAV motion.

Concerning our second objective, the most straightforward approach would likely be to initiate tracking exclusively with high-confidence detections, a strategy commonly employed in most existing MOT methods. However, this method is prone to missing the tracking of many target objects, resulting in an increased number of false negatives. As a solution, this paper suggests initiating tracking for target objects based on low-confidence detections. However, this tracking strategy

¹Department of Immersive Media Engineering & Department of Computer Science Education, Sungkyunkwan University, Seoul, Korea e-mails: Inpyo Song(songinpyo@skku.edu) & Jangwon Lee(leejang@skku.edu)

is expected to produce more false positives. Therefore, there is a need for a method to reduce false positives while effectively tracking a larger number of target objects, including small-scale objects with low-confidence scores. To address this issue, we propose employing traditional appearance matching algorithms based on hand-crafted appearance features, such as color histograms and scaled image Mean Squared Error (MSE), for the effective association of these low-confidence detections. Our hypothesis is that these traditional features can enhance the data association step, particularly in scenarios where obtaining high-quality deep learning-based visual features is challenging, for instance, when dealing with small-scale or partially obscured objects. In contrast to deep learning-based features, traditional hand-crafted features can be more robust for image matching in situations with high occlusions and small-scale objects. In addition, the utilization of traditional techniques offers additional benefits, as they do not require a large volume of training data, and tend to be very general, and cost-efficient in contrast to deep learning-based features [5].

To assess the efficacy of our approach, we conducted experiments to measure object tracking performance on two UAV-specific datasets (VisDrone2019 and UAVDT) as well as a general dataset (MOT17), achieving state-of-the-art results. Moreover, throughout our experimentation we identified annotation errors in the UAVDT dataset, including inaccuracies and omissions. Consequently, we rectified numerous erroneous annotations and supplemented annotations for the many instances of missing data in the UAVDT dataset. We plan to make our enhanced version of the annotations accessible to the public, thereby facilitating further research in this field. To summarize, this paper offers the following contributions:

- We present an approach for UAV multi-object tracking that adeptly handles inherent challenges, including rapid and irregular UAV motion, small-scale objects, and occlusions.
- We introduce a novel strategy for initiating tracking from low-confidence detections, particularly beneficial in UAV contexts. To manage these low-confidence detections, we suggest revisiting traditional appearance matching algorithms based on hand-crafted features.
- We provide comprehensive evaluation of our approach on multiple datasets, including VisDrone2019, UAVDT, and MOT17, demonstrating its effectiveness and superior performance.
- We have identified and corrected missing and incorrect annotations in the UAVDT dataset, contributing to more accurate evaluation and further research.

II. RELATED WORK

A. Object Detection in Multi-Object Tracking (MOT)

Recent vision-based multi-object tracking systems mainly use a tracking-by-detection approach, a two-step process of detection and tracking [6]–[11]. This method involves first detecting objects in video frames and then associating these

detections over time to track their movement. Our research focuses on enhancing the association phase of this approach, specifically tailored to the unique challenges of UAV footage.

B. Data association for Multi-Object Tracking

The data association phase matches detection results with tracklets. It typically uses motion and appearance cues to compute similarities between objects and applying a matching strategy. **Motion cue** Motion-based strategies utilize object or camera movement information (motion cue) to effectively track objects. SORT [12] stands as a pioneering work in this field, employing Kalman Filters to predict object locations and compare these predictions with new detections. Tractor [8] suggested camera motion compensation by aligning frames through image registration [13]. OC-SORT [14] enhances the traditional Kalman Filter approach by focusing on object-centric adjustments, yielding improved tracking performance in scenarios with irregular motion. BoTSORT [7] proposed correcting Kalman Filter’s predicted bounding box locations using camera motion compensation based on feature point tracking [15], [16]. However, prior studies often fall short in handling extensive motion in UAV footage. This study tackles this issue by introducing a straightforward yet efficient UAV-specific motion compensation technique to enhance tracking accuracy and mitigate bounding box distortion.

Appearance cue Relying solely on motion cues can be challenging especially when bounding box predictions are imprecise. Consequently, strategies incorporating appearance cues have emerged to address these challenges. Earlier methods used simple techniques like color histograms [17], but modern approaches like DeepSORT [18] and JDE [9] use advanced neural networks to analyze object features. Similarly, MOTDT [19] used a Re-ID network [20] for feature extraction, while Tractor [8] utilized a siamese network. More recently, several attempts have been made to simultaneously train object detectors and Re-ID modules to enhance feature learning, as evidenced by JDE [9] and FairMOT [11]. Such approaches, however, have failed to effectively address the unique challenges presented by UAV-captured videos, including issues with small and fast objects as well as the common occurrence of occluded objects. Consequently, they often encounter difficulties in accurately tracking target objects. Hence, in this paper, we propose revisiting traditional hand-crafted appearance features like color histograms and scaled image MSE. Through experimentation, we demonstrate that these traditional appearance cues can exhibit enhanced robustness, particularly when dealing with small-scale or partially obscured objects.

Matching strategy In terms of the matching strategy, the Hungarian algorithm, originally introduced by Kuhn in 1955 [21], is considered the standard-de-facto approach for associating detections with tracklets. DeepSORT [18] took this a step further by introducing a cascaded matching algorithm that gives priority to recent tracklets. MOTDT [19] first applies the appearance cue, and then employs IoU for any tracklets that remain unmatched. ByteTrack

[6] introduces a secondary algorithm specifically designed to associate detections with lower confidence to existing tracklets. More recently, Transformer based architectures [22], [23] employ track queries to find object locations, enabling implicit matching through attention mechanisms without relying on the Hungarian algorithm. However, so far, very little attention has been paid to effectively handling and utilizing objects detected with low confidence scores, a common occurrence in UAV application scenarios. This is likely because using low-confidence detections without careful consideration of a new matching strategy may lead to more false positives. In contrast to prior studies, this paper proposes initializing new tracks from low-confidence detections by effectively managing them through traditional handcrafted feature-based appearance matching.

C. Multi-object Tracking on UAVs

For MOT in UAV videos, unique challenges arise from small object sizes and blurred images. Therefore, in recent years, competitions like VisDrone have emerged [24], driving the development of methods tailored for UAV-specific scenarios. V-IOU Tracker [25] enhances tracking using visual cues in the absence of detections. HMTT [26] employs a hierarchical approach with a re-identification subnet, while GIAOTracker [27] introduces the NSA Kalman Filter, which adjusts noise scale during state updates and several post-processing approaches. In contrast, UAVMOT [28] and FOLT [29] provide online tracking solutions. UAVMOT employs an adaptive motion filter for challenging scenarios, and FOLT utilizes optical flow to track small objects. In contrast to these methods primarily refining motion tracking for UAV scenarios, our framework integrates improvements in both motion and appearance cues, along with matching strategies, to create a more accurate and robust tracking system for the challenging environment of UAV footage.

III. APPROACH

In this section, we introduce our tracking strategy **SF-Track** (Small and Fast-moving Objects Tracking), designed to address the frequent occurrence of low-confidence detections in UAV applications. Unlike previous methods, such as [6], [7], [12], [18], [19], which take a conservative approach by starting tracking only on high-confidence detections, we actively leverage low-confidence detections as the starting point for tracking. Furthermore, we enhance tracking accuracy by revisiting traditional appearance matching algorithms for data association, such as color histogram similarity and scaled image MSE. Additionally, we employ a motion compensation technique designed for irregular and rapid movements in UAV footage.

The method details are outlined following the steps in Alg. 1. Our tracking algorithm takes a video sequence V as input, along with an object detector Det , with defined parameters: a detection confidence threshold τ , and a low-confidence detection tracking threshold ρ . The output comprises a set of tracks T each containing bounding box coordinates across frames while preserving the identity of the objects.

Algorithm 1: Pseudo-code of SFTrack.

Input : A video sequence V ; object detector Det ; detection score threshold τ ; tracking initiate threshold ρ

Output: Tracks T of the video

```

1 Initialization:  $T \leftarrow \emptyset$ 
2 for frame  $f_k$  in  $V$  do
3    $D_k \leftarrow Det(f_k)$ 
4    $D_{high} \leftarrow \emptyset$ 
5    $D_{low} \leftarrow \emptyset$ 
6   for  $d$  in  $D_k$  do
7     if  $d.det\_score > \tau$  then
8        $D_{high} \leftarrow D_{high} \cup \{d\}$ 
9     else
10       $D_{low} \leftarrow D_{low} \cup \{d\}$ 
11    end
12  end
13  UAV motion compensation Section 3.A #/
14  Figure 2 #/
15   $M_k \leftarrow$  UAV motion from  $f_{k-1}$  to  $f_k$ 
16  for  $t$  in  $T$  do
17     $t \leftarrow KalmanFilter(t)$ 
18     $t \leftarrow M_k(t)$ 
19  end
20  1st association: High detections #/
21  Associate  $T$  and  $D_{high}$ 
22   $D_{high\_remain} \leftarrow$  unmatched  $D_{high}$ 
23   $T_{remain} \leftarrow$  unmatched  $T$ 
24  2nd association: Low detections Section 3.B #/
25  Figure 3 #/
26  Associate  $T_{remain}$  and  $D_{low}$ 
27   $D_{low\_remain} \leftarrow$  unmatched  $D_{low}$ 
28   $T_{re\_remain} \leftarrow$  remaining tracks from  $T_{remain}$ 
29  Delete unmatched tracks #/
30   $T \leftarrow T \setminus T_{re\_remain}$ 
31  Initialize new tracks Section 3.C #/
32  Init T from unmatched high-score detections #/
33  for  $d$  in  $D_{high\_remain}$  do
34     $T \leftarrow T \cup \{d\}$ 
35  end
36  Init T from low-score detections #/
37  for  $d$  in  $D_{low\_remain}$  do
38    if  $d.appearance\ similarity\ with\ D_{high} > \rho$ 
39      then
40         $T \leftarrow T \cup \{d\}$ 
41      end
42  end
43 return  $T$ 

```

This process begins by employing the object detector to predict bounding boxes and confidence scores. Subsequently,



Fig. 2. Comparison of tracking results on a low-altitude video. (a) ByteTrack without motion-compensation (MC), fails to maintain tracks. (b) BoTSORT’s MC leads to distorted bounding boxes. (c) SFTrack with UAV MC, ensures consistent and accurate tracking.

similar to BYTE [6], we then categorize detections into high-confidence D_{high} and low-confidence D_{low} groups, guided by the score threshold τ (refer to lines 1-12 in Alg. 1).

A. UAV Motion Compensation

After separating low-score and high-score detection boxes, we apply an enhanced motion compensation algorithm by incorporating a simple yet effective idea, drawing inspiration from Aharon et al.’s method [7]. While the conventional motion compensation is effective, it often falters in low-altitude UAV scenes due to errors predominantly caused by bounding box ratio distortion. Our solution is a straightforward adjustment, that preserving the bounding box ratio. To implement this, we calculate the affine transformation matrix M_k capturing camera motion by extracting feature points [15] and utilizing sparse optical flow [16] to track these features. This matrix, encompassing scale, rotation, and translation components, guides the transformation of bounding boxes. Subsequently, we update the state vector of tracks T using Kalman Filter predictions, applying the M_k to this updated state vector. Furthermore, to prevent significant distortion in the bounding box ratio during this process, we enforce constraints on the scale factor in M_k , specifically selecting the larger scale factor between the x and y axes and uniformly applying it to both axes. This approach maintains a consistent bounding box ratio across frames, even during intense UAV camera motion at low-altitudes (Figure 2).

After compensating for UAV motions, we first associate all tracks T with high-confidence detections D_{high} using IoU and cosine similarity of Re-ID features [30]. The Hungarian algorithm [21] is employed for the association based on the product of these measures. Unmatched high-confidence detections then form set D_{high_remain} , and the remaining tracks constitute set T_{remain} (see lines 20-23 in Alg. 1).

B. Matching low-confidence detections

Following the initial association, a second association is performed between the remaining tracks T_{remain} and low-confidence detections D_{low} . In contrast to the first association, which relied on deep learning-based Re-ID features, in this phase, we propose to utilize traditional hand-crafted appearance features for these low-confidence detections (see lines 24-27 in Alg. 1). Specifically, we employ color histogram similarity and scaled image MSE as appearance cue. Therefore, associations are established through the Hungarian algorithm, relying on the product of three key values: IoU, color histogram similarity, and scaled image MSE.

Color histograms offer an efficient mechanism to capture the color distribution characteristics of an object within an image. We categorize the color intensity into 8 uniformly spaced levels (0-31, 32-63, ..., 224-255) for each color channel (R, G, B), and record the frequency of pixel color values. The color histogram similarity is then calculated by the Bhattacharyya distance [31]. Additionally, we employ scaled image MSE to measure image dissimilarity. To ensure a fair comparison regardless of object scale, images cropped by bounding boxes are first resized to the same size. Subsequently, the MSE is computed between these two images. To convert MSE to a similarity score within a range of 0-1, we subtract the normalized MSE value from 1.

The motivation behind this is rooted in the limitations of deep learning methods, including Re-ID, which frequently encounter challenges in low-resolution settings or when objects are occluded. This is due to their inclination to prioritize the foreground object, potentially overlooking the object of interest when it is located in the background. In contrast, traditional matching algorithms, like color histogram similarity and scaled image MSE, consider the complete image, regardless of occlusion. Thus, these methods show superior performance in occlusion and low-resolution scenarios due to their holistic use of image information. As illustrated in Figure 3, our approach to low-confidence detection association outperforms deep learning-based methods, particularly in low-resolution and occlusion scenarios.

After this phase of matching low-confidence detections, any tracks that remain unmatched following both the high- and low-confidence detection matching phases are removed from the tracking pool. (see Alg. 1, lines 28-29). In this context, we provide a grace period before removing these tracks, waiting for 30 frames where they don’t find a match in a row. This waiting period, influenced by the BYTE approach [6], prevents premature termination of a valid track due to temporary occlusions or brief detection failures.

C. Initiation of tracking from low-confidence detections

Finally, in contrast to previous approaches such as BYTE [6], we introduce a distinctive method by initiating new tracks not only from unmatched high-confidence detections (D_{high_remain}) but also from low-confidence detections (see Alg. 1, lines 30-39). However, initiating tracks from low-confidence detections poses a challenge as this new strategy

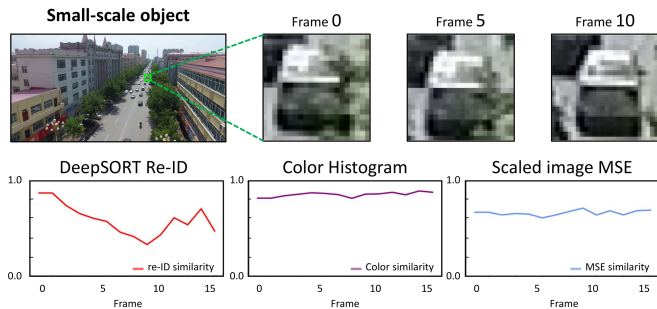


Fig. 3. This figure compares appearance similarity methodologies for small-scale objects. Re-ID faces challenges in low-resolution and occlusion scenarios, while color histogram and MSE matching offer consistent similarity throughout the frame.

is expected to generate more redundant, unnecessary, and potentially intrusive object candidates.

To resolve this challenge, we evaluate low-confidence detections by comparing their appearance similarity with high-confidence detections from the same object category using Re-ID features. This process functions as a filtering mechanism, allowing us to selectively initiate new tracks exclusively from unmatched low-confidence detections (D_{low_remain}) that surpass a predetermined appearance similarity threshold (ρ) with high-confidence detections. This double-checking procedure enhances the reliability of low-confidence detections, improving tracking of small-scale objects and enhancing the overall performance of our system.

IV. DATASETS AND EVALUATION METRICS

In this study, we conducted experiments on three distinct datasets: two UAV-specific datasets, VisDrone2019 and UAVDT, and a general dataset, MOT17. For evaluation, we employed CLEAR metrics [32], including Multiple Object Tracking Accuracy (MOTA), False Positives (FP), False Negatives (FN), and ID switches (IDs).

A. Datasets

VisDrone2019 [33] is employed for five tasks: Object Detection in Images, Object Detection in Videos, Single-Object Tracking, Multi-Object Tracking (MOT), and Crowd Counting in videos captured by UAVs. Among them, this paper specifically concentrates on the MOT task within the VisDrone2019 dataset, comprising 56 sequences in the training set, 7 sequences in the validation set, and 17 sequences in the test-dev set. To ensure a fair comparison, we limit our focus to five object categories: pedestrian, car, van, truck, and bus, aligning with the object categories used in [28].

UAVDT [34], is a UAV Detection and Tracking benchmark, offering 30 training sequences and 20 testing sequences exclusively focusing on car tracking for the MOT task. Unlike other datasets, UAVDT provides annotations with environmental attributes such as weather, altitude, and camera view, captured under diverse conditions.

MOT17 [35] serves as a prominent dataset in the field of multi-object tracking. It covers a wide array of urban scenarios involving pedestrians, vehicles, and other dynamic

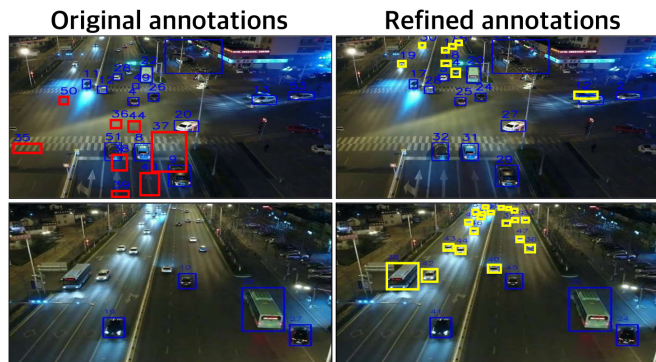


Fig. 4. This image provides a comparative visualization of the Original UAVDT and our Refined UAVDT annotations. Errors in the original annotations, which do not correspond to actual objects, are highlighted in ‘red’. The ‘yellow’ markers represent our additional annotations for visible objects within the Refined UAVDT dataset.

entities. We adhere to the protocol specified in [6], [10], splitting the training set equally for training and validation, and focusing on the pedestrian category for the MOT task. This dataset is the only non-UAV dataset used in our setup.

B. UAVDT refinement

During our study, we identified significant annotation errors in the UAVDT dataset, as illustrated in Figure 4. Significantly, certain objects lacked annotations, while numerous annotations either commenced before the object appeared or persisted beyond its exit in the image frame. To rectify these issues and improve the accuracy of our study, we refined annotations within 4721 frames. This effort, reflected in the ‘Refined UAVDT,’ introduced 43,981 additional annotations, elevating the count from 340,906 to 384,887 objects, along with 55 extra tracks. We plan to publicly release this enhanced dataset for future research.

C. Implementation details

Our experiments are conducted on a desktop server equipped with an Intel Core i9-10900X CPU @ 3.70GHz processor and single NVIDIA GeForce RTX 3090 GPU. We employed the YOLOX [36] object detector, pre-trained on the COCO dataset [37]. For VisDrone2019 and UAVDT experiments, the training was conducted on their respective training sets, with an input video resolution of 1920×1080 , adhering to the protocol set by [28]. For the MOT17 experiment, we trained on both the CrowdHuman [38] dataset and the first half of the MOT17 training set, while evaluation was performed on the second half of the MOT17 validation set. The input video resolution for this experiment was 1440×800 , in line with [6].

D. Evaluation metrics

In this section, we describe how we compute MOTA and IDF1, kindly providing a reminder. For further details on CLEAR metrics, please refer to [32]. MOTA is a composite metric that considers FP, FN, and IDs and is calculated as follows:

$$MOTA = 1 - \frac{(FP + FN + IDs)}{GT} \times 100 \quad (1)$$

TABLE I

COMPARISON OF DATA ASSOCIATION METHODS ON **VisDrone2019** TEST-DEV SET. (*INDICATES METHODS USING THEIR OWN DETECTOR).

Method	MOTA↑	IDF1↑	IDs↓	FP↓	FN↓	MT↑	ML↓
SORT [12]	25.7	46.4	865	9592	162216	366	856
MOTDT [19]	34.9	44.8	1692	13280	136283	440	805
DeepSORT [18]	35.1	47.0	934	12773	137078	421	810
UAVMOT* [28]	36.1	51.0	2775	27983	115925	520	574
BoTSORT [7]	39.6	53.6	1066	21333	118003	581	659
OCSORT [14]	41.3	52.4	941	12985	122445	429	703
FOLT* [29]	42.1	56.9	800	24105	107630	-	-
ByteTrack [6]	42.3	53.4	989	15689	117539	586	675
SFTrack (Ours)	47.2	62.1	557	27159	94910	753	518

TABLE II

COMPARISON OF DATA ASSOCIATION METHODS ON **UAVDT** TEST SET. (*INDICATES METHODS USING THEIR OWN DETECTOR).

Method	MOTA↑	IDF1↑	IDs↓	FP↓	FN↓	MT↑	ML↓
SORT [12]	32.4	55.1	165	45633	179092	442	405
MOTDT [19]	32.5	56.0	422	55664	168754	487	363
DeepSORT [18]	33.1	56.4	311	54016	168357	480	365
UAVMOT* [28]	46.4	67.3	456	66352	115940	624	221
OCSORT [14]	47.0	67.4	148	56309	119892	559	156
BoTSORT [7]	47.3	68.8	417	75547	99365	667	174
ByteTrack [6]	47.5	69.4	159	76181	98415	688	167
FOLT* [29]	48.5	68.3	338	36429	155696	-	-
SFTrack (Ours)	49.5	70.3	209	72866	95935	709	152

Here, GT denotes numbers of ground-truth objects.

We also adopt the IDF1 score from [39], which primarily assesses the identity preservation capability, recognizing that in the MOTA score, the impact of larger quantities of FP and FN is often more significant than that of ID

$$IDF1 = \frac{2 \cdot IDTP}{GT + C} \quad (2)$$

Here, IDTP represents the number of True Positives with correct identities, and C refers to the total computed detections. Additionally, we applied the standard metrics, Mostly Tracked (MT) and Mostly Lost (ML) [40]. MT and ML represent the proportions of ground truth trajectories that are covered by computed trajectories with correct and incorrect identities for at least 80% of their lifespan, respectively.

V. EXPERIMENTS AND RESULTS

A. Comparisons with other state-of-art approaches

We compared our approach with other prominent association methods using three datasets: VisDrone2019, UAVDT, and MOT17. However, it should be acknowledged that UAVMOT [28] and FOLT [29] do not strictly fall under the category of data association methods, thereby limiting their evaluation in an identical environment. Nevertheless, we referenced evaluation protocols from their original papers, and strive to align our evaluation settings, considering factors such as categories and input image resolution.

UAV Datasets On the UAV datasets, our SFTrack outperforms others significantly across most evaluation metrics. For example, our method achieved a 4.9% improvement in MOTA over the top-performing ByteTrack and a 5.2%

TABLE III

RESULTS ON **REFINED UAVDT** TEST SET USING "SAME DETECTOR".

Method	MOTA↑	IDF1↑	IDs↓	FP↓	FN↓	MT↑	ML↓
SORT [12]	35.6	55.3	217	33073	207951	419	448
MOTDT [19]	36.4	56.4	590	41505	196014	464	402
DeepSORT [18]	36.8	56.7	418	40166	195926	454	405
OCSORT [14]	52.1	69.8	186	37908	141340	535	192
BoTSORT [7]	52.5	69.7	441	56067	121304	646	201
ByteTrack [6]	53.0	70.3	215	56021	119674	668	193
SFTrack (Ours)	55.3	71.3	284	49946	117224	682	184

TABLE IV

RESULTS ON **MOT17** VALIDATION SET USING "SAME DETECTOR".

Detector	Method	MOTA↑	IDF1↑	IDs↓	FP↓	FN↓
YOLOX-x (AP = 0.624)	SORT	74.6	77.3	287	2797	10592
	DeepSORT	75.4	77.2	239	3399	9641
	MOTDT	75.8	77.7	263	2689	10069
	OCSORT	76.0	79.5	159	3325	9406
	BoTSORT	76.3	79.5	162	3521	9088
	ByteTrack	76.4	79.3	164	3418	9126
	SFTrack (Ours)	76.9	80.1	153	3237	9044
YOLOX-Nano (AP = 0.565)	SORT	69.4	67.0	705	2701	13092
	MOTDT	70.4	72.1	300	1730	13905
	DeepSORT	71.0	73.1	311	3141	12175
	OCSORT	71.5	72.0	231	2603	12486
	BoTSORT	71.8	71.7	236	2848	12158
	ByteTrack	72.1	71.3	239	2708	12066
	SFTrack (Ours)	73.0	74.7	220	2436	11918




Fig. 5. Experimental comparison with ByteTrack on VisDrone2019 video sequences. Improvements in MOTA and IDF1 scores correlate with higher MRA, indicating the presence of small and fast objects.

increase in IDF1 score compared to FOLT, the leading IDF1 scorer on VisDrone2019. However, we acknowledge the trade-offs of our approach, specifically between a reduction in false negatives (FN) and an increase in false positives (FP). This trade-off stems from our methodology, which actively utilizes low-confidence detections and employs traditional appearance matching algorithms for association. However, it is worth noting that the proposed SFTrack demonstrates more balanced tracking performance compared to previous state-of-the-art methods. Moreover, in many tracking scenarios, it is more crucial not to miss target objects than to briefly track additional objects, including those that are not of interest.

Evaluation on Challenging UAV Subsets To explicitly demonstrate SFTrack's enhancement in tracking small and fast-moving objects, we further examined its performance under challenging conditions within the VisDrone2019 and




Fig. 6. Experiments on **Refined UAVDT** under challenging conditions.

Refined UAVDT datasets. On VisDrone2019, we specifically selected videos with a high prevalence of small, fast-moving objects, categorized by high Mean Relative Acceleration (MRA) values [29]. Our findings, illustrated in Figure 5, reveal that SFTrack’s performance substantially improves in scenarios with elevated MRA, underscoring its adeptness at handling rapid object motions and small sizes of objects. On Refined UAVDT we conducted experiments using videos from high-altitude UAV flights (over 70m from the ground) and scenarios involving view changes (Figure 6). In these challenging conditions, our approach outperformed the next best method by a wider margin, demonstrating the robustness and efficacy of our methodology. These experiments affirm SFTrack’s capability to significantly improve tracking accuracy for small and fast-moving objects. Figure 7 provides qualitative results supporting this conclusion.

MOT17 We additionally evaluated our proposed method on the widely used MOT17 dataset, as detailed in Table IV. The results highlight the versatility of our approach, showcasing its effectiveness in non-UAV video environments. Notably, when applied to two detectors, YOLOX and YOLOX Nano, our method exhibited a substantial performance gap, particularly with the lower-performing YOLOX Nano. This underscores the efficacy of our method, especially in handling increased low-confidence detections resulting from diminished detection performance, as illustrated in Figure 8.

B. Ablation study

Finally, we conducted an ablation study to evaluate the impact of each component in our new tracking algorithm. To be specific, both BoTSORT and our SFTrack utilize ByteTrack as the backbone algorithm as a starting point. In this study, we thus demonstrate how our three proposed simple solutions enhance existing tracking methods. For the experiment setting, we focused on the car class in the VisDrone2019, as it is one of the most commonly appearing objects in the dataset. Table V provides a detailed overview of the distinct impact of each component. We initially replaced the motion compensation technique (MC) in BoTSORT with our proposed UAV-specific method that preserve aspect ratio, resulting in notable improvements, specifically a 1.4% increase in MOTA and 3.7% in IDF1. Initiating tracking with low-confidence detections further increased MOTA and IDF1 by 3.2%. Moreover, our association strategy for low-confidence detections, utilizing traditional

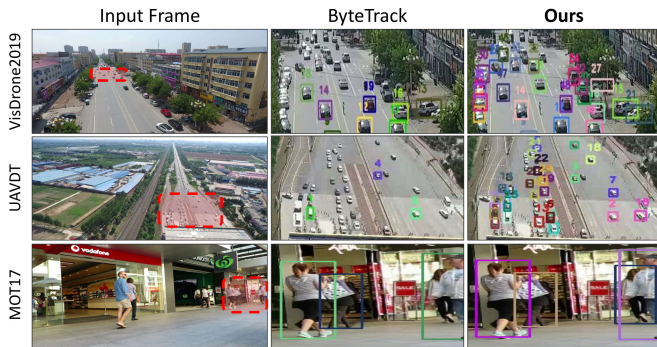


Fig. 7. Qualitative results on VisDrone2019 test-dev, UAVDT test, and MOT17 validation datasets. Our method significantly outperforms ByteTrack, especially for occluded or smaller-scale objects.




Fig. 8. Tracking comparison on a VisDrone2019 test-dev sequence, focusing on ‘pedestrian’ class with a far-reaching UAV view. Bounding boxes are categorized as High ($s \geq 0.7$) and Low ($s < 0.7$) confidence. The first row shows a tracking algorithm initiating tracks solely from high-confidence detections. In contrast, the second row illustrates our approach, initiating tracking from both high and low-confidence detections. This highlights our method’s effectiveness in tracking small-scale objects in UAV footage, particularly when detection confidence is low.

TABLE V
ABLATION STUDY ON CAR CLASS IN VISDRONE2019 TEST-DEV.

Model	1st Asso.	Motion Comp.		Tracking Init.	2nd Asso.	Evaluation Metrics	
	Re-ID	MC	UAV MC	w/ Low Det.	Traditional Algorithms	MOTA↑	IDF1↑
ByteTrack						58.6	67.7
	✓					58.9	68.0
BoTSORT	✓	✓				58.7	71.2
	✓		✓			60.1	74.9
	✓		✓	✓		63.3	78.1
SFTrack	✓		✓	✓	✓	65.0	78.3

appearance matching algorithms contributed to an additional increase of +1.7% in MOTA and +0.2% in IDF1. This result indicates that traditional methods effectively handle frequent low-confidence detections at minimal extra expense. Despite concerns about lighting and viewpoint variations, their performance remains robust in consecutive frames of high-FPS videos. In terms of tracking speed, integrating UAV MC has a similar impact on tracking speed as BoTSORT’s MC. Both our system (row 4) and BoTSORT achieve a tracking speed of 13 FPS. Initiating tracking from low-confidence detections slightly reduces this to 12 FPS, while applying traditional appearance matching algorithms further adjusts it to 10 FPS, aligning with UAVMOT’s tracking speed. While these improvements slightly impact tracking speed, our implementation in Python demonstrates the concept, with potential for speed enhancements by transitioning to C.

VI. CONCLUSION

In this study, we introduced a novel association approach with three simple yet effective ideas for object tracking, focusing on the unique challenges of UAV footage. Our strategies for low-confidence detections effectively address tracking small-scale objects and managing UAV motions. Our method outperforms established approaches across a variety of benchmarks, as demonstrated on three widely recognized datasets such as VisDrone2019, UAVDT, and MOT17. Furthermore, we identified and corrected errors in the existing annotations of the UAVDT dataset. The enhanced version will be publicly accessible, providing a valuable resource for more accurate benchmarking in the field.

ACKNOWLEDGMENT

This work was supported under the framework of international cooperation program managed by National Research Foundation of Korea (NRF-2021K1A3A1A49097955).

REFERENCES

- [1] J. Lee, H. Tan, D. Crandall, and S. Šabanović, "Forecasting hand gestures for human-drone interaction," in *ACM/IEEE International Conference on Human-Robot Interaction*, 2018.
- [2] S. MohaimenianPour and R. Vaughan, "Hands and faces, fast: mono-camera user detection robust enough to directly control a uav in flight," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2018.
- [3] F. Zhao, "The research on police uav investigation technology based on intelligent optimization algorithm," in *International Conference on Networking, Communications and Information Technology*, 2021.
- [4] R. E. Kalman, "A new approach to linear filtering and prediction problems," 1960.
- [5] N. O'Mahony, S. Campbell, A. Carvalho, S. Harapanahalli, G. V. Hernandez, L. Krpalkova, D. Riordan, and J. Walsh, "Deep learning vs. traditional computer vision," in *Advances in Computer Vision: Proceedings of the 2019 Computer Vision Conference*. Springer, 2020.
- [6] Y. Zhang, P. Sun, Y. Jiang, D. Yu, F. Weng, Z. Yuan, P. Luo, W. Liu, and X. Wang, "Bytetrack: Multi-object tracking by associating every detection box," in *European Conference on Computer Vision*, 2022.
- [7] N. Aharon, R. Orfaig, and B.-Z. Bobrovsky, "Bot-sort: Robust associations multi-pedestrian tracking," *arXiv preprint arXiv:2206.14651*, 2022.
- [8] P. Bergmann, T. Meinhardt, and L. Leal-Taixe, "Tracking without bells and whistles," in *IEEE/CVF International Conference on Computer Vision*, 2019.
- [9] Z. Wang, L. Zheng, Y. Liu, Y. Li, and S. Wang, "Towards real-time multi-object tracking," in *European conference on computer vision*, 2020.
- [10] X. Zhou, V. Koltun, and P. Krähenbühl, "Tracking objects as points," in *European Conference on Computer Vision*, 2020.
- [11] Y. Zhang, C. Wang, X. Wang, W. Zeng, and W. Liu, "Fairmot: On the fairness of detection and re-identification in multiple object tracking," *International Journal of Computer Vision*, vol. 129, 2021.
- [12] A. Bewley, Z. Ge, L. Ott, F. Ramos, and B. Upcroft, "Simple online and realtime tracking," in *IEEE international conference on image processing*, 2016.
- [13] G. D. Evangelidis and E. Z. Psarakis, "Parametric image alignment using enhanced correlation coefficient maximization," *IEEE transactions on pattern analysis and machine intelligence*, vol. 30, no. 10, 2008.
- [14] J. Cao, J. Pang, X. Weng, R. Khirodkar, and K. Kitani, "Observation-centric sort: Rethinking sort for robust multi-object tracking," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.
- [15] J. Shi *et al.*, "Good features to track," in *IEEE conference on computer vision and pattern recognition*, 1994.
- [16] J.-Y. Bouguet *et al.*, "Pyramidal implementation of the affine lucas kanade feature tracker description of the algorithm," *Intel corporation*, vol. 5, no. 1-10, 2001.
- [17] M.-Y. Ju, C.-S. Ouyang, and H.-S. Chang, "Mean shift tracking using fuzzy color histogram," in *International Conference on Machine Learning and Cybernetics*, 2010.
- [18] N. Wojke, A. Bewley, and D. Paulus, "Simple online and realtime tracking with a deep association metric," in *IEEE international conference on image processing*, 2017.
- [19] L. Chen, H. Ai, Z. Zhuang, and C. Shang, "Real-time multiple people tracking with deeply learned candidate selection and person re-identification," in *IEEE international conference on multimedia and expo*, 2018.
- [20] L. Zhao, X. Li, Y. Zhuang, and J. Wang, "Deeply-learned part-aligned representations for person re-identification," in *IEEE international conference on computer vision*, 2017.
- [21] H. W. Kuhn, "The hungarian method for the assignment problem," *Naval research logistics quarterly*, vol. 2, no. 1-2, 1955.
- [22] T. Meinhardt, A. Kirillov, L. Leal-Taixe, and C. Feichtenhofer, "Trackformer: Multi-object tracking with transformers," in *IEEE/CVF conference on computer vision and pattern recognition*, 2022.
- [23] F. Zeng, B. Dong, Y. Zhang, T. Wang, X. Zhang, and Y. Wei, "Motr: End-to-end multiple-object tracking with transformer," in *European conference on computer vision*, 2022.
- [24] H. Fan, D. Du, L. Wen, P. Zhu, Q. Hu, H. Ling, M. Shah, J. Pan, A. Schumann, B. Dong, *et al.*, "Visdrone-mot2020: The vision meets drone multiple object tracking challenge results," in *European conference on computer vision Workshops*, 2020.
- [25] E. Bochinski, T. Senst, and T. Sikora, "Extending iou based multi-object tracking by visual information," in *IEEE International Conference on Advanced Video and Signal Based Surveillance*, 2018.
- [26] S. Pan, Z. Tong, Y. Zhao, Z. Zhao, F. Su, and B. Zhuang, "Multi-object tracking hierarchically in visual data taken from drones," in *IEEE/CVF International Conference on Computer Vision Workshops*, 2019.
- [27] Y. Du, J. Wan, Y. Zhao, B. Zhang, Z. Tong, and J. Dong, "Giaotracker: A comprehensive framework for mcot with global information and optimizing strategies in visdrone 2021," in *IEEE/CVF International Conference on Computer Vision*, 2021.
- [28] S. Liu, X. Li, H. Lu, and Y. He, "Multi-object tracking meets moving uav," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.
- [29] M. Yao, J. Wang, J. Peng, M. Chi, and C. Liu, "Folt: Fast multiple object tracking from uav-captured videos based on optical flow," *arXiv preprint arXiv:2308.07207*, 2023.
- [30] H. Luo, Y. Gu, X. Liao, S. Lai, and W. Jiang, "Bag of tricks and a strong baseline for deep person re-identification," in *IEEE/CVF conference on computer vision and pattern recognition workshops*, 2019.
- [31] T. Kailath, "The divergence and bhattacharyya distance measures in signal selection," *IEEE transactions on communication technology*, vol. 15, no. 1, 1967.
- [32] K. Bernardin and R. Stiefelhagen, "Evaluating multiple object tracking performance: the clear mot metrics," *EURASIP Journal on Image and Video Processing*, 2008.
- [33] P. Zhu, L. Wen, D. Du, X. Bian, H. Fan, Q. Hu, and H. Ling, "Detection and tracking meet drones challenge," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 11, 2021.
- [34] D. Du, Y. Qi, H. Yu, Y. Yang, K. Duan, G. Li, W. Zhang, Q. Huang, and Q. Tian, "The unmanned aerial vehicle benchmark: Object detection and tracking," in *European conference on computer vision*, 2018.
- [35] A. Milan, L. Leal-Taixé, I. Reid, S. Roth, and K. Schindler, "Mot16: A benchmark for multi-object tracking," *arXiv preprint arXiv:1603.00831*, 2016.
- [36] Z. Ge, S. Liu, F. Wang, Z. Li, and J. Sun, "Yolox: Exceeding yolo series in 2021," *arXiv preprint arXiv:2107.08430*, 2021.
- [37] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *European Conference on Computer Vision*, 2014.
- [38] S. Shao, Z. Zhao, B. Li, T. Xiao, G. Yu, X. Zhang, and J. Sun, "Crowdhuman: A benchmark for detecting human in a crowd," *arXiv preprint arXiv:1805.00123*, 2018.
- [39] E. Ristani, F. Solera, R. Zou, R. Cucchiara, and C. Tomasi, "Performance measures and a data set for multi-target, multi-camera tracking," in *European Conference on Computer Vision Workshops*, 2016.
- [40] Y. Li, C. Huang, and R. Nevatia, "Learning to associate: Hybrid-boosted multi-target tracker for crowded scene," in *IEEE conference on computer vision and pattern recognition*, 2009.