

# The State of Aerial Surveillance: A Survey

K. Nguyen, *Member, IEEE*, C. Fookes, *Senior Member, IEEE*, S. Sridharan, *Life Senior Member, IEEE*, Y. Tian, *Fellow, IEEE*, F. Liu, X. Liu, *Senior Member, IEEE*, A. Ross, *Senior Member, IEEE*

**Abstract**—The rapid emergence of airborne platforms and imaging sensors are enabling new forms of aerial surveillance due to their unprecedented advantages in scale, mobility, deployment and covert observation capabilities. This paper provides a comprehensive overview of human-centric aerial surveillance tasks from a computer vision and pattern recognition perspective. It aims to provide readers with an in-depth systematic review and technical analysis of the current state of aerial surveillance tasks using drones, UAVs and other airborne platforms. The main object of interest is humans, where single or multiple subjects are to be detected, identified, tracked, re-identified and have their behavior analyzed. More specifically, for each of these four tasks, we first discuss unique challenges in performing these tasks in an aerial setting compared to a ground-based setting. We then review and analyze the aerial datasets publicly available for each task, and delve deep into the approaches in the aerial literature and investigate how they presently address the aerial challenges. We conclude the paper with discussion on the missing gaps and open research questions to inform future research avenues.

**Index Terms**—Aerial surveillance, Mass surveillance, Persistent surveillance, Eyes in the sky, Wide-Area motion imagery (WAMI)

## 1 INTRODUCTION

THIS paper presents the first ever review of the state-of-the-art research of the rapidly evolving area of aerial surveillance, covering the unique challenges in performing the human centric tasks of detection, tracking, recognition, person re-detection and action recognition on aerial data. Large volumes of human centric aerial surveillance data are being collected in several major initiatives. For example in 2011, the Gorgon Stare project, led by the Pentagon in the United States, rolled out a MQ-9 drone equipped with an advanced Multi-Spectral Targeting System called ARGUS to Iraq and Afghanistan. With as many as 368 individual cameras, Gorgon Stare could capture 1.8 billion pixels per frame, enough imaging power to spot an object six inches wide from an altitude of 25,000 feet [93]. An entire city of size  $10 \times 10 \text{ km}^2$  could be continuously observed at a resolution sufficient to monitor any person or vehicle in the city. This all-seeing system enabled persistent and mass surveillance of a wide area 24/7 without awareness of the citizens in the city.

In 2019, a Stratollite - a giant stratospheric balloon, orbited continuously for 45 days in the stratosphere at an altitude of 65,000 feet to monitor North America [104]. Equipped with hi-tech radars and multi-spectral sensors, these solar-powered balloons can hover over a small area of interest, taking photos with a quality of five centimeters per pixel and beaming the footage down to the ground station. This resolution is many times superior than commercial

satellites and is sufficient to detect a mobile phone in a person's hand and track all individual vehicles day and night in an entire city for weeks, even months [104].

In the last four consecutive years, four "Vision meets drones" challenges in major academic conferences (ECCV 2018, ICCV 2019, ECCV 2020, ICCV 2021) were organized targeting aerial object detection, tracking and crowd counting from drone-based footage. The challenges were based on a large-scale aerial dataset, VisDrone, with more than 400 aerial videos formed by 265K frames and 2.6M bounding boxes or points of targets of frequent interests, such as pedestrians, cars, and bicycles [187]. Captured by the civilian DJI Mavic drones flying at an altitude of a few hundred feet in urban areas, the VisDrone footage shows the potential and feasibility of large data collection for analysis.

From the all-seeing eyes of the Gorgon Stare and the months-lasting Stratollites, to the off-the-shelf DJI drones, as illustrated in Fig. 1, the uptake of aerial surveillance has significantly broadened across military, industry and academic sectors. The academic research related to aerial surveillance has seen a huge boom. According to Scopus, there are more than 78K UAV/drone/aerial papers published <sup>1</sup>, 38K surveillance papers published <sup>2</sup> and 1.3K UAV/drone/aerial surveillance papers published <sup>3</sup> in the last six years as illustrated in Fig. 2.

Powered by recent breakthroughs in computer vision and deep learning, the implementation of the basic aerial surveillance tasks of detection, tracking, identification, and action recognition, on aerial data are emerging as an important and timely research area to be investigated. This paper dives deep into aerial surveillance tasks from a computer vision and pattern recognition perspective. Compared with

- K. Nguyen, C. Fookes and S. Sridharan are with the Research Program in Signal Processing, AI and Vision Technologies (SAIVT), Queensland University of Technology, Australia. E-mail: {k.nguyenthanh,c.fookes,s.sridharan}@qut.edu.au
- Yingli Tian is with The City College of New York, United States. E-mail: ytian@ccny.cuny.edu
- Xiaoming Liu, Feng Liu and Arun Ross are with Department of Computer Science and Engineering, Michigan State University, United States. E-mail: {liuxm,liufeng6,rossarun}@cse.msu.edu.

1. Query string: TITLE-ABS (uav OR drone OR aerial)
2. Query string: TITLE-ABS (surveillance) AND TITLE-ABS (detection OR tracking OR identification OR action OR event)
3. Query string: TITLE-ABS (uav OR drone OR aerial) AND TITLE-ABS (surveillance) AND TITLE-ABS (detection OR tracking OR identification OR action OR event)



Figure 1: The rise of aerial surveillance. (1) The first column shows the military all-seeing Gorgon Stare [93] deployed in Iraq, using a MQ-9 drone (costs \$17M USD) flying at an altitude of 25K feet, able to capture simultaneously an area of  $10 \times 10$  km, with a 1.8 GB sensor at a Ground Sampling Distance (GSD) of 15 cm. (2) The second column is Stratollite, a giant solar-powered balloon flying at an altitude of 65K feet [104], which can hover for months, able to detect a mobile phone in a person's hand at a GSD of 5 cm. (3) The third column is the off-the-shelf DJI Mavic 2 (costs \$3K USD) flying at an altitude of 1K feet, able to capture simultaneously an area of  $1 \times 1$  km, at a GSD of 20 cm [187].

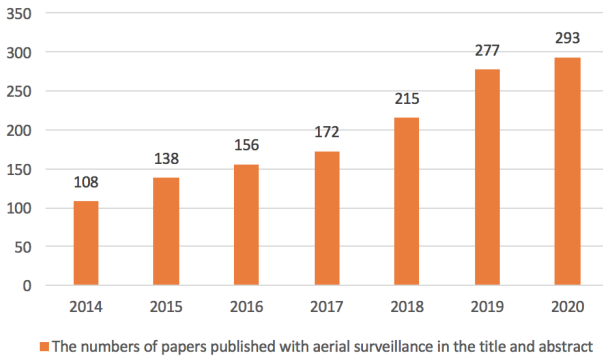


Figure 2: According to Scopus, there are more than 1,300 papers published with one aerial keyword, *i.e.* UAV or Drone or Aerial, and one surveillance keyword, *i.e.* detection or tracking or identification or action or event in the title and abstract in the last six years.

ground-based surveillance, aerial surveillance provides unprecedented scope in surveillance scale, mobility and covert observation. However, it also poses distinct challenges for the computer vision and deep learning community to address including extreme object viewing angles, low-resolution of objects and non-linear object distribution with diverse backgrounds from moving cameras. In surveillance, it is of great interest to detect and track humans, identify

and re-identify humans from a camera or across multiple cameras, understand their behaviors and interactions, as well as search for humans with specific characteristics. From a computer vision and machine learning perspective, these tasks are categorized as detection, identification (biometrics, soft-biometrics, and re-ID), action recognition and crowd analysis. While these tasks have been actively studied in generic ground-based images, the challenges with aerial images are daunting but targeted research to address these challenges are rapidly emerging in the aerial domain. This paper will investigate each task from 4 angles:

- (1) *aerial surveillance challenges*: discuss unique challenges for performing the task in the aerial domain, and the amount of performance drop when ground-based approaches shift to the aerial domain,
- (2) *aerial surveillance datasets*: analyze the aerial datasets publicly available for the task,
- (3) *approaches to solve aerial surveillance challenges*: delve deep in the state-of-the-art approaches in the aerial literature and how they address the challenges unique to aerial data,
- (4) *techniques to improve aerial surveillance tasks*: further investigate techniques to improve the performance of the aerial task.

The rest of the paper is structured as follows. Section §2 identifies the advantages and challenges of aerial surveillance, and the statistics of public datasets and academic papers published. Sections §3 - §6 review the state-of-the-

art research of the four central aerial surveillance tasks: detection, tracking, identification, and action recognition. Section §7 summarizes the state of aerial surveillance, open questions and future outlook. Section §8 concludes the paper.

## 2 AERIAL SURVEILLANCE

Aerial surveillance is the task of employing airborne platforms and imaging sensors mounted on them to detect, track, identify, and monitor behavior and activities of a person or a group of persons [80], [93], [187]. Aerial surveillance is also known as Wide Area Motion Imagery (WAMI) or mass surveillance since it is capable of monitoring a wide range area, such as an entire city with a large population persistently and unblinkingly 24/7. A wide range of airborne platforms are available with diverse characteristics in flying ranges and altitudes, endurance, speed, manoeuvrability, payload and vulnerabilities as summarized in Table 7 in Appendix. These airborne platforms are deployed with a wide range of sensors. In this paper, we consider mainly techniques for analyzing data acquired by RGB imaging sensors, which capture data with spatial details to be analyzed by computer vision and deep learning. Imaging sensors are categorized by their spectrum as summarized in Table 8 in Appendix.

To limit the scope of this paper, we only focus on the human-centered security-related surveillance tasks. Aerial surveillance empowers a wide range of human-centric applications, including border patrol [52], search and rescue [15], [121], maritime surveillance [128], protest monitoring, drug trafficking monitoring [145], military IED tracking [93] and crime fighting [160]. Due to these burgeoning applications of aerial surveillance, corporate aerial surveillance is rapidly growing. Along with this growth, there is a growing privacy threat. Aerial surveillance of individuals as well as mass surveillance can now be conducted at low cost using the drone technology. This poses a serious threat to privacy since there are not well-established privacy protections to prevent widespread and indiscriminate aerial surveillance [125]. In this review we will not be covering the privacy issues of aerial surveillance which is a very important topic. Since aerial surveillance has immense potential to provide safety and security to the public, it is important to ensure that the regulations that are imposed to protect privacy do not hamper the development of the area of aerial surveillance.

### 2.1 Advantages and challenges

Aerial surveillance can be employed independently or complementary with ground-based surveillance. Compared with standard ground-based surveillance, there are four distinct advantages to observation from the air:

- *Scale*: high resolution imaging sensors mounted on airborne platforms enable covering a wide area of observation from the air at multiple resolutions (depending on the altitude of the aerial platform) with less occlusion [80], [173].
- *Mobility*: airborne platforms such as UAVs and drones can move rapidly to a target. Once at the destination,

they can circle or hover and adjust positions over the target destination for optimum viewing [53], [128].

- *Deployment*: airborne platforms can be deployed at any time (day and night), for any terrain (land and sea), and launched from a long distance [52], [127], [128].
- *Covert observation*: airborne platforms can adjust the flying altitudes to either quietly hide themselves, stay out of visibility reach or explicitly reveal their presence [33], [93].

However aerial surveillance opens up a plethora of challenges that must be addressed. Aerial surveillance not only inherits all challenges of unconstrained and outdoor surveillance, but also exhibits additional unique challenges to be addressed. From a data perspective, there are seven key challenges to contend with:

- *Small resolutions*: objects may appear extremely small in aerial data due to the high flying altitude [27], [173].
- *Multiple scales*: multiple instances of the same class, e.g. person, can appear drastically different in sizes and scales [92], [148].
- *Extreme views*: objects can appear in overhead views, i.e. top views and angle views, which rarely exist in generic object detection [30], [111].
- *Moving cameras*: the view of objects may continuously change due to the moving of the camera mounted on the airborne platform [8], [80]. This also adds additional challenge of motion blur and camera stabilization [16], [146].
- *Non-uniform distribution*: in many cases, objects are distributed non-uniformly: clustered with high density, e.g. in a busy urban village [72], [163]; or cluttered with low density in a wide area, e.g. in a search and rescue task [158], [166].
- *Illumination*: non-linear local strong and/or low illumination and lighting due to a wide area coverage [187].
- *Noise*: the scene may be obstructed, e.g. cloud, fog, haze and rain [69], [158]; The wind may add more challenge in video stabilization for sharp footage.

### 2.2 State of aerial surveillance

We provide the first overview of the state of aerial surveillance tasks via analyzing public datasets available for four key tasks, i.e. detection, tracking, identification and action recognition, and by summarizing the public interest in this area as illustrated in Fig. 3. While this does not reflect efforts by military and defense in these areas, which are usually protected or classified, the state of the public dataset collections is expected to represent a driving factor for the academic research community towards new aerial surveillance tasks.

- *Publication distribution*: Fig. 3.a shows that a majority of aerial surveillance papers that are published in open literature focus on detection and tracking, while the number of papers on identification is limited. This would be due to the fact that compared to other tasks, aerial human identification requires high quality visual details of subjects from expensive and purchase-regulated cameras/sensors, which may not be available or too costly to acquire for public use.
- *Dataset distribution*: Fig. 3.b shows that a majority of public datasets are for research in detection and action

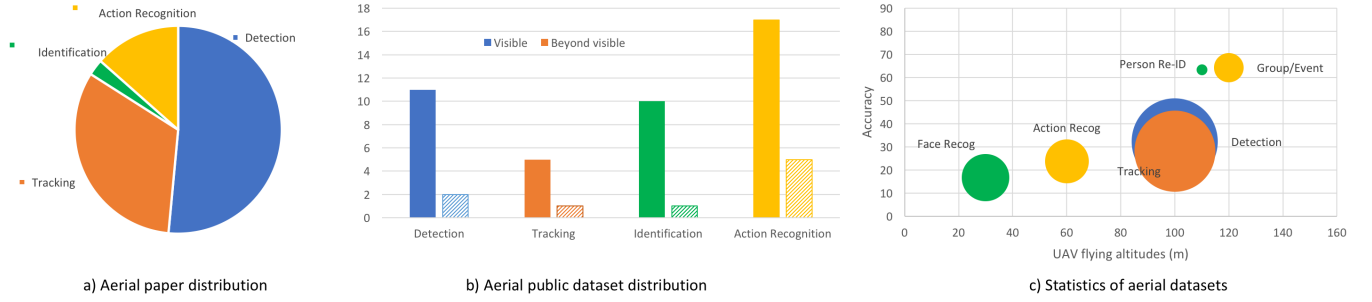


Figure 3: Overview of the state of aerial surveillance tasks. a) The distribution of academic papers across four aerial surveillance tasks: detection, tracking, identification and action recognition; b) The distribution of public aerial datasets across four aerial surveillance tasks: detection, tracking, identification and action recognition. Solid columns denote datasets in the visible spectrum, stripe columns denote datasets in other spectrum; c) The statistics of aerial datasets for aerial surveillance tasks. The size of circles denotes the scale of datasets (*i.e.* #images), the horizontal axis denotes UAV flying altitudes, and the vertical axis denotes the state-of-the-art accuracies on these datasets.

recognition, while the number of datasets for identification is limited. Recent datasets have gone beyond visible cameras, capturing footage from other spectral bands such as thermal, night vision infrared and depth.

- *Dataset characteristics:* Fig. 3.c shows that: (i) Flying altitudes: all of the datasets are captured while flying drones or UAVs at most 120 meters, which is the regulated maximum altitude for civilian drones; (ii) Scale: datasets for aerial human detection have large scales, while others usually have small scales; (iii) State-of-the-art accuracies: aerial group/event recognition tasks have already achieved high accuracies in available databases, while face recognition struggles to perform well in the aerial setting.

The following sections (Sections §3 - §6) dive deep into the four central aerial surveillance tasks.

### 3 AERIAL HUMAN DETECTION

Object detection aims at detecting and localizing instances of visual objects of a certain class, *i.e.* human in our surveillance setting, in digital images and videos. Akin to ground-based surveillance, aerial object detection also plays a fundamental role as a foundation for high-level tasks in aerial surveillance. However, due to the unique characteristics of aerial imaging settings, aerial object detection exhibits new challenges to be addressed.

In this section, we first review generic ground-based object detection landscape to form a foundation for the object detection discussion. Then, we focus on discussing the aerial human detection in terms of challenges, datasets, approaches, and future outlook.

#### 3.1 Challenges for aerial human detection

Despite the great success of the generic object detection methods trained on ground-to-ground images, a huge performance drop is observed when they are directly applied to images captured by UAVs [36], [171]. Examples of the performance drop of state-of-the-art detectors are illustrated in Fig. 5. The Cascade R-CNN [17] drops the performance by 50% in the aerial VisDrone [187] dataset compared with the ground-based COCO and Pascal VOC datasets [36]. The

Faster R-CNN [117] also drops the performance by 30% in the aerial TinyPersons [173] dataset compared with the ground-based COCO and Pascal VOC datasets [171].

The unsatisfactory performance is owing to the domain shift when compared to ground-based data caused by the high flying altitude and the camera characteristics. Aerial human detection shares all seven challenges discussed in Section §2.1. Examples of these challenges are illustrated in Fig. 4.

#### 3.2 Datasets for aerial human detection

The number of aerial datasets has quickly increased in the last few years, partially due to the affordable availability of off-the-shelf drones such as DJI. We summarize public datasets and their statistics in Table 1. Table 1 shows a wide range of public datasets where two notable large-scale ones are VisDrone [187] and TinyPersons 2020 [173].

- *VisDrone 2018-2021* [187]: The VisDrone team has compiled a dedicated large-scale drone benchmark and organized four consecutive challenges in object detection in ECCV/ICCV from 2018 to 2021. It consists of 400 video clips formed by 265K frames and 10K static images with 2.6M bounding boxes, captured by various drone-mounted cameras, covering a wide range of aspects including location (14 cities), environment (urban and country), objects (pedestrian, vehicles, bicycles, *etc.*), and density (sparse and crowded scenes).
- *TinyPersons 2020* [173]: this dataset is behind two “Tiny Object Detection” challenges in ICCV 2019 and ECCV 2020. The unique characteristic of this dataset is the tiny resolution of humans. A majority of human instances appear as small as [20, 32] pixels and as tiny as [2, 20] pixels. In total, there are 72K objects with bounding boxes that have been manually annotated.

There are a wide range of other datasets with specific characteristics targeting various applications.

- *Context:* there is a diverse range of contexts, from urban villages [15], [100], [187], rural [80], [187], university campuses [118], internet [173], forest [12], to agricultural areas [75].
- *Flying altitudes:* there is a diverse range of flying altitudes, from a very close distance of less than 10m [71],



Figure 4: Challenges for aerial object detection: (i) low resolution, (ii) a wide range of scales, (iii) arbitrary viewing angles, (iv) non-uniformly distributed, (v) illumination. Images from the TinyPersons [173] and VisDrone [187] datasets.

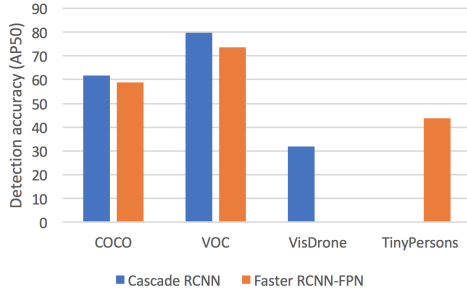


Figure 5: State-of-the-art generic detectors drop the performance when shifting to the aerial data. The figure compares the detection accuracy (AP50) of Cascade R-CNN [17] and Faster R-CNN [117] on two ground-based (COCO and VOC) and two aerial (VisDrone and TinyPersons) datasets.

to a close range 10-50m [100], [118], to a middle range 50-120m [75], [80], to above the recreational limit >120m [5], [173].

**Multimodal:** many datasets also provide multi-modal data from other on-board sensors such as GPS, time, altitude, IMU, velocity, weather conditions. The AU-AIR [14] labeled each frame with time, GPS, IMU, altitude, velocities of the UAV. [158] annotated and released altitude, viewing angle, and weather for the VisDrone dataset. These auxiliary details can be used to improve detection [158].

**Task-specific:** many aerial human detection datasets are aimed for specific tasks. The HERIDAL [15], [90] and SARD [120] datasets were collected to support the search and rescue mission with drones. The AgriDrone [75] dataset was collected in an agricultural context for rural applications.

**Beyond visible:** sensors from other spectrum have also been employed to complement visible cameras. Infrared sensors complement visible cameras in such adverse conditions as night time or low light. Both the BIRDAI [12] dataset and UAV-Human [80] dataset provides thermal-infrared images of humans for detection. Depth sensors complement visible cameras in 3D object details and can be found in the recent UAV-Human [80] dataset.

**Competitions:** the task of aerial human detection has attracted great attention from the computer vision commu-

Table 1: Public datasets for aerial human detection. ‘Spec.’, ‘#Con’, ‘#Vid’, ‘#ID’, and ‘#Box’ respectively represent for the imaging spectrum (V: Visible, T: Thermal-infrared, Z: Depth), the context where data was collected, the number of videos, identities and bounding boxes.

No.	Dataset	Spec.	Con.	#Vid.	#ID	#Box.
1	VisDrone [187]	V	Urban	400	50	2M
2	TinyPerson [173]	V	Internet	2	632	1,264
3	AU-AIR [14]	V	Real	8	192	1,011
4	MiniDrone [13]	V	Carpark	21	100	200
5	HERIDAL [15]	V	Urban	2	632	1,264
6	StanfordDrone [118]	V	Campus	2	450	900
7	SARD [120]	V	Urban	6	491	38,271
8	UAV123 [100]	V	Urban	123	632	1,264
9	UAVDT [169]	V	Urban	2	502	3,012
10	PDT-ATV [115]	V	Urban	2	1,467	14,000
11	AgriDrone [75]	V	Agri	2	412	8,240
12	BIRDAI [12]	VT	Forest	2	119	34K
13	UAV-Human [80]	VTZ	Various	64K	1,144	41K

nity. Multiple challenges have been organized in top-tier conferences. The challenges “Vision meets drones” [36] including detection, tracking and crowd counting have been organized in the last four consecutive years in ECCV 2018, ICCV 2019, ECCV 2020 and ICCV 2021 using the VisDrone dataset. The TinyPerson challenges [171], organized in ICCV 2019 and ECCV 2020, focus on the persons from a very long distance with a wide view using the TinyPerson dataset.

Details of these datasets can be found in Appendix §B.

### 3.3 Approaches for aerial human detection

Network architecture is the key consideration for all deep-learning-based approaches. The aerial human detection community has adopted two-stage network architecture [50], [66], [180], one-stage network architecture [50], [162],

[170], anchor-free architecture [20], [135] and network ensembles [36], [36].

#### *Two-stage networks*

Due to its performance, two-stage detectors such as Faster R-CNN [117] are popular in the literature [30], [50], [58], [77]. In aerial object detection, a strong multi-scale representation is crucial due to the variance of object sizes and resolutions, which reflects the popularity of the FPN [84] backbone in the literature [50], [66], [180]. The multi-scale challenge also makes multi-stage detectors well-suitable for aerial object detection. Many recent approaches [2], [148], [175] achieved good performance with Cascade R-CNN [17]. In fact, many entries, including the winning and runner-up entries, of the VisDrone detection challenges combined Cascade R-CNN [17] with other networks to achieve high detection performance [36]. Similarly, two of top-three detectors in the TinyPersons challenge employed Cascade R-CNN [171].

#### *One-stage networks*

One-stage networks are fast and less computationally intensive than their multi-stage counterparts, hence they have also been employed in many aerial object detection approaches. RetinaNet [85] is among the most popular one-stage networks for aerial object detection by [50], [162], [170]. EfficientDet [134] has also been employed in the aerial setting by [69], [92]. Real-time detectors such as YOLO and its variants [11] are also a popular one-stage network for aerial object detection by [30], [114]. Pelee [151], a real-time object detection system on mobile devices, has also been utilized in aerial object detection by [142].

#### *Anchor-free*

In aerial images where object instances vary drastically in resolution and size in a single image, finding good prior anchor sizes is not feasible. LaLonde *et al.* [72] also showed that it is easier to locate extremely small objects with points instead of anchors, hence many approaches have shifted to using anchor-free network designs. CenterNet [39] with a Hourglass backbone [101] has been a popular choice such as [51], [153]. [31] took advantages of FCOS [136], [137]. Tang *et al.* [135] employed a CornerNet design [73]. Zhang *et al.* [175] employed a FreeAnchor design [181]. Chen *et al.* [20] also designed a point based detector called RRNet and observed that point-based detectors usually outperform all anchor-based detectors in the VisDrone test set.

#### *Network ensemble*

Each detector may have different strengths and weaknesses. While the multi-stage detectors tend to produce more false negatives, which means that multi-stage detectors fail to detect some objects, single-stage detectors generally propose more bounding boxes with less quality [2], [188]. Hence, combining them may predict more bounding boxes than multi-stage detectors and the quality of the single-stage detector predictions may be enhanced by the multi-stage one. Alibaba *et al.* [2] showed that combining a multi-stage detector, *i.e.* Cascade R-CNN [17], and a single-stage detector, *i.e.* CenterNet [39], yields higher accuracy than individual detectors. The ensemble strategy has been widely employed in practice since it is effective to improve the accuracy of object detection. For example, the winner of the VisDrone object detection challenge 2020 [187], DPNNetV3 [36], ensembles a few powerful backbones such as HRNet-W40 [149], Res2Net

[45], Balanced Feature Pyramid Network [107] and Cascade R-CNN [17]. The second-ranked detector also uses different combinations of multiple models (*i.e.* Cascade R-CNN [17], HRNet [149], and ATSS [176]) to fuse the detection results.

### 3.4 Techniques to solve aerial detection challenges

The distinct aerial challenges discussed in Section §3.1 arise in a large number of domains, across which an effective aerial detection model has to stay robust.

#### 3.4.1 Low resolution and scale variance

The most challenging factor in aerial object detection is the low resolution or small size of objects due to the high flying altitude. A human may appear as tiny as a few pixels in an image. For example, the size of humans in the TinyPerson dataset [173] ranges [2, 32], in practice [2, 20] is considered as tiny and [20, 32] is considered as small. In the VisDrone dataset, human sizes only ranges from 0.00014% to 5.59% and a mean of 0.044% and the human objects occupy only three pixels in images within a frame resolution of  $1,916 \times 1,078$  pixels [187]. The scale of object resolutions also varies largely from  $10^1$  to  $10^3$  pixels not just within the dataset, but also within a single image [20]. There are three key approaches to deal with small object detection in aerial data: (i) improve feature maps for small objects [86], [138], (ii) incorporate context information of small objects [82], [97], and (iii) data augmentation [36].

#### 3.4.2 Multi-scale detection

Multi-scale detection is commonly used in aerial object detection due to the dominance of small objects and the co-presence of object instances with a wide range of scales. Similar to generic multi-scale object detection, many aerial object detectors [58], [83], [148], [180] employ FPN [84] and its variants. New techniques have also been proposed to improve the fusion between feature maps. Gong *et al.* [50] introduced a fusion factor to weigh the aggregation between adjacent layers. Gong *et al.* [51] also proposed an adaptive feature selection scheme to aggregate feature maps. Wang *et al.* [148] proposed to refine multi-scale features by increasing the receptive field size for high-level semantic features. Differently, Yang *et al.* proposed a sequential approach to employ local high-resolution features for small objects in their QueryDet [162]. Liu *et al.* proposed to combine both feature pyramid and image pyramid to improve the aggregated features.

#### 3.4.3 Viewing angle

Extreme views due to the high altitude of cameras, birds-eye views, and highly angled views can make objects appear very differently from popular ground-based data. Objects, *i.e.* human, in aerial images may only have a top view and could be arbitrary rotated [165]. If the training data contains these novel views and rotations, most networks can learn to cope with these large variations. However, this requires a very large scale dataset to cover all possible variations. Data augmentation [36], [171] and generative models such as GANs [22] can be used to generate extra data to deal with this challenge during training.

Table 2: Winning entries of the VisDrone-DET2020 aerial object detection challenge (AP - Average Precision) [36].

Method	AP(%)	AP50(%)	AP75(%)
DPNetV3	37.37	62.05	39.10
SMPNet	35.98	59.53	37.41
DBNet	35.73	59.63	36.92
ECascade RCNN	34.09	56.77	35.30
FPAFS CenterNet	32.34	56.46	32.39
DOHR RetinaNet	21.68	44.59	18.73

### 3.4.4 Non-uniform distribution

The aerial community has departed from the uniform cropping strategy [114], [142] by learning density-specific regions to perform detection in parallel with detection on the original image. Li *et al.* [77] designed a multi-column CNN [182] to predict a density map and utilized a sliding window on the density map to generate proposal regions for cropping. Zhang *et al.* [174] designed a Difficult Region Estimation Network to estimate the regions that contain difficult targets and utilized a sliding window on the difficulty map to generate proposal regions for cropping. Yang *et al.* [163] designed a Cluster Proposal Network (CPNet) and Tang *et al.* [135] designed a coarse network (CPEN) to predict cluster chips which have similar object scales.

## 3.5 Techniques to improve aerial human detection

*Context modeling:* While feature pyramid networks have the capacity to aggregate multi-scale features to deal with the small and tiny object instances in aerial data, context information can provide strong cues on the presence of humans [97]. DBNet [36], which ranked the third place in the VisDrone challenge, added a global context block to improve detection. Contextual information can also be captured by dilated convolution [159].

*Auxiliary meta-data:* Wu *et al.* [158] proposed to incorporate UAV-specific nuisance annotations which are freely available as meta-data in aerial data. For example, for UAVDT [169] and VisDrones [187], three nuisance annotations (altitude, weather, and view angle) are used simultaneously with object classes to train the representation network by a combining a object detection loss and a nuisance prediction loss. Kiefer *et al.* [69] proposed a multi-domain strategy to leverage those nuisance annotations. The multi-domain strategy improves representation, which subsequently boosts detection accuracy on both datasets.

## 3.6 Insights from the competitions

*VisDrone challenge:* results from the VisDrone object detection challenge organized and reported in ECCV 2020 [36], as presented in Table 2, provide two major insights:

- The winning entry, DPNetV3, of the challenge achieved an AP50 of 62.05% [36]. This huge increase in performance compared to the generic detectors such as Cascade R-CNN (16.09%) and approach the accuracy level in the ground-based dataset MSCOCO.
- High-performing entries mainly focus on either (i) directly dealing with aerial challenges, or (ii) combining powers via a network ensemble and data augmentation. For example, DPNetV3 combines Cascade R-CNN with HRNet, Res2Net, FPN and Cascade R-CNN [36].

## 3.7 Aerial human detection beyond visible

Despite the benefits of performing human detection in spectrum other than visible such as 24h all weather scenarios (*i.e.* fog), very few studies have been conducted regarding this topic. Bondi *et al.* showed that Faster-RCNN and YOLO only achieved an accuracy of 18.1% and 10.4% respectively on the BIRDSAI [12] dataset. Haider *et al.* [55] proposed an convolutional autoencoder architecture to map the human heat signature in the input thermal image to the spatial density maps. Portman *et al.* [115] and Ma *et al.* [88] showed handcrafted approaches would struggle to deal with aerial thermal human detection due to a large number of variations. Schedl *et al.* showed that aerial thermal person detection under occlusion conditions can be notably improved by combining multi-perspective images before classification [121]. Using a synthetic aperture imaging technique, they achieve this with a precision and recall of 96% and 93%, respectively on their own dataset.

# 4 AERIAL HUMAN TRACKING

Object tracking, aiming to estimate the location and scale of an object (*i.e.* pedestrian in our case) in a video with an initial bounding box given in the first frame, is a key step in video analysis. Object tracking is the process of estimating the trajectory of an object in a sequence and consists of four components, *i.e.* object initialization, appearance model, motion prediction and object positioning. To this objective, numerous approaches have been proposed by tackling the following questions: how to model the objects' motion, appearance and shape? which image feature is suitable for tracking? which prior information could be leveraged? The answers to these questions heavily depend on the scenario in which the tracking is performed. For instance, there are many differences between aerial object tracking technology and standard ground object tracking technology due to unique or additional challenges in aerial surveillance settings.

## 4.1 Challenges for aerial human tracking

Section §2.1 and Fig. 4 detailed the potential challenges in the aerial-based imagery comparing to the ground-based ones. However, there still exists some unique or additional challenges for aerial human tracking. Examples of the performance drop of state-of-the-art trackers are illustrated in Fig. 6.

- *Blurred imagery:* Targets for which the image blurred due to the shaking of the UAV fuselage, obstacle avoidance or long imaging distance.
- *Background clutter:* The background around the target has similar appearances as the target, *i.e.* the building shadows.
- *Severe occlusion:* The tracked objects may be occluded for a long time, and even disappear for several frames.
- *Fast-moving targets:* Targets may exhibit fast movement in the image plane.
- *Illumination variation:* The illumination of the target changes significantly across frames.

Table 3: Public datasets for aerial human tracking. ‘Spec.’, ‘#Con.’, ‘#Vid.’, ‘#ID’, and ‘#Box.’ respectively represent for the imaging spectrum (V: Visible, T: Thermal-infrared, Z: Depth), the context where data was collected, the number of videos, identities and bounding boxes.

No.	Dataset	Spec.	Con.	#Vid.	#ID	#Box.
1	VisDrone [187]	V	Urban	400	50	2M
2	UAV123 [100]	V	Internet	123	632	1,264
3	Campus [118]	V	Real	8	192	11K
4	DTB70 [13]	V	Youtube	70	100	200
5	UAVDark135 [76]	V	Urban	2	632	1,264
6	BIRDSAI [12]	VT	Forest	2	119	34K

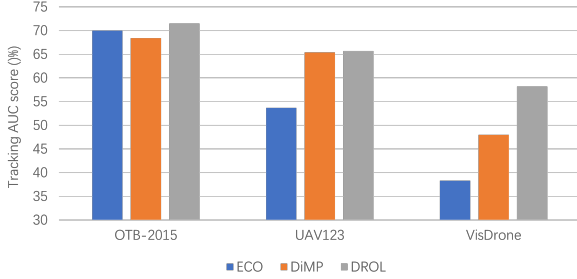


Figure 6: State-of-the-art generic tracking drop the performance when shifting to the aerial data. The figure compares the area-under-the-curve (AUC) score of ECO [25], DiMP [9] and DROL [184] on one ground-based (DTB-2015 [157]) and two aerial (UAV123 [100] and VisDrone [187]) datasets.

## 4.2 Datasets for aerial human tracking

To date, there is still a limited availability of annotated datasets specific to UAVs where human trackers can be rigorously evaluated or trained for precision and robustness in aerial scenarios. Some drone-captured datasets do not contain either pedestrians or annotations for pedestrians, such as UAVDT dataset [37]. We thus do not list them here.

We summarize public datasets and their statistics in Table 3. There are two notable large-scale datasets: VisDrone [187] and UAV123 [100].

- **VisDrone 2018-2021:** The VisDrone team has compiled a dedicated large-scale drone benchmark and organized challenges for tracking, *i.e.* VisDrone-SOT2018 [156] and VisDrone-SOT2019 [38]. The VisDrone-SOT2018 consists of 132 videos with 106K frames. Compared with VisDrone-SOT2018, VisDrone-SOT2019 introduces 35 new sequences. To further increase the diversity of videos and assess the performance of trackers in the wild, VisDrone-SOT2020 [42] conducts extensive evaluation of more tracking algorithms using the same dataset in VisDrone-SOT2019. VisDrone2021 further increases the dataset size to 400 videos with more diverse scenarios.

- **UAV123 2021:** this dataset [100] contains 123 fully annotated HD sequences over 110K frames taken from UAV platforms. Each video has 12 attribute categories. A video may have a variety of attributes by the shooting conditions. The captured targets include *pedestrian*, vehicles, boats, groups and *etc.* The video resolution is between 720p and 4K.

**Beyond visible:** datasets captured at spectrum other than visible have also been collected. BIRDSAI [12] was collected by a thermal camera mounted on drones. UAVDark 135 [76] targets aerial tracking at night at various scenes such as crossings, t-junctions, roads, and highways.

Details of these datasets can be found in Appendix §C.

## 4.3 Approaches for aerial human tracking

As mentioned earlier, aerial tracking presents its own challenges compared to generic tracking. In aerial imagery, the target appearance changes severely. However, most of prevailing aerial trackers are adopted from generic tracking techniques. We review the representative approaches based on the following two main streams: discriminative correlation filter (DCF)-based and deep-learning-based.

### Discriminative correlation filter (DCF)-based trackers

Due to the simplicity and efficiency of the models, DCF-based trackers are widely used for object tracking in aerial videos. Their main advantage is that they can generate plenty of cyclic shift candidates and learn filters in the frequency domain efficiently [65], [81]. Wang *et al.* [152] developed a stability measurement metric based on the peak-to-sidelobe ratio, which makes the DCF-based aerial tracker more robust to complicated appearance variations. Huang *et al.* [62] proposed an aberrance repressed correlation filter, which is capable of suppressing aberrances that is caused by other background noise introduced by conventional DCF and appearance changes of the tracked objects. He *et al.* [57] introduced a unified tri-attention framework to leverage multi-level visual information, including contextual, spatio-temporal and dimension attention to improve UAV tracking robustness and efficiency. Ye *et al.* [167] proposed a novel tracking framework based on a multi-regularized correlation filter, which leads to favorable adaption to object appearance variations and enhancement of discriminability. Zhang *et al.* [178] exploited a two-stage scheme that combines a detection-based network (IoU-Net) with DCF-based tracker for object tracking in aerial videos.

### Deep-learning-based trackers

Even though deep-learning-based trackers have gained tremendous success in ground-view videos, they find it hard to cope with the nuances of aerial videos. The main reason is lack of large-scale well-labeled *training* datasets and view-invariant appearance model for fast-moving targets [129]. To address these challenges, CRAC [129] introduced a GAN-based tracker to model contextual relation and transfer the relations from ground-view to the aerial-view videos while retaining the discriminative features. C2FT [179] presented a coarse-to-fine reinforcement learning architecture to address the aspect ratio variation of targets in aerial tracking. COMET [91] introduced a context-aware IoU-guided tracker that exploits a multitask two-stream network and an offline

reference proposal generation strategy to improve the performance of aerial-view trackers.

#### 4.4 Techniques to solve aerial tracking challenges

##### 4.4.1 Blurred imagery

Environmental factors and equipment jitter lead to target blur. Deblurring techniques could be used to mitigate this effect. For example, [116] constructs a nonlinear blurred core with multiple moving components. A blind deconvolution technique that used a piecewise linear model was introduced to estimate the unknown kernels, which provides an efficient solution for deblurring motion blurred images.

##### 4.4.2 Small targets

Tracking small targets (*i.e.* human in our case) involves major difficulties comprising the lack of sufficient target information to distinguish them from the background and a large number of possible locations where the subject may appear. Extracting complementary features by employing additional information such as the context leads to more robust feature representations for small objects. The work of [91] introduces a context-aware IoU-guided tracker that fully exploits target-related information by multi-scale feature learning and attention modules.

##### 4.4.3 Occluded or fast-moving targets

Handling occluded or fast-moving targets is of crucial importance in several applications of aerial tracking. The typical solution is to determine the changes in the aspect ratio since targets encounter with fast motion or occlusion often cause aspect ratio variation [78].

#### 4.5 Techniques to improve aerial human tracking

*Integration of tracking and detection:* In a long aerial video, the targets may frequently leave and re-enter the view (full occlusion). A promising scheme to address this challenges is to incorporate a detection process into tracking for re-detecting the tracked targets [186].

*Motion modeling:* Motion information is crucial for distinguishing tracked targets from clutter background. Motion (including camera and target motion) modeling [164] is thus necessary in the aerial tracking.

*Domain adaptation:* GANs have been exploited in the context of unsupervised domain adaptation [144]. To alleviate the lack of large-scale well-labeled human videos in aerial tracking, one may consider to transfer the data or feature with diverse variations from ground-view to aerial-view via domain adaptation techniques. For instance, [129] proposes a robust tracker with a dual GAN learning mechanism, which can model contextual relation and transfer the ground-view features to the aerial-view ones.

#### 4.6 Insights from the competitions

We review the state-of-the-art aerial tracking methods based on their superior performances on *VisDrone single object tracking 2020 challenge*. Based on the observation, there are several potential directions that warrant future research.

- *Data augmentation:* Data augmentation is an essential part in network training with limited training data. The

augmentation operators for aerial tracking methods include: resale, horizontal flop, rotation, shift, image contrast by Gamma correction and Laplacian operator.

- *Searching region:* Since the fast-moving and occlusion often occur in aerial setting, the tracking would benefit from a large search region.
- *Spatio-temporal context:* the spatio-temporal context information takes an important role for improving the robustness of the trackers. The trackers based on RNN or 3D-CNN, which leverages the spatio-temporal information, are more effective in coping with the target appearance variations across frames.

#### 4.7 Aerial human tracking beyond visible

As near-infrared (NIR) and thermal infrared (TIR) sensors become more affordable in aerial image acquisition, there are increasing interests in tracking human through 24h all weather scenarios (*i.e.* fog). There are obvious advantages in NIR and TIR imagery, such as low sensitivities to illumination variations and good capabilities for nighttime and bad weather. In addition, different sensory modalities can complement each other in developing a more robust multi-modality tracking framework (*e.g.* visible + infrared). However, the challenges of background clutter, small target size and occlusion still remain in infrared tracking. Besides, public datasets for infrared aerial tracking, not surprisingly, are very scarce. Thus, few studies [18], [89] have been conducted regarding this topic.

### 5 AERIAL HUMAN IDENTIFICATION

Beyond detection and tracking, recognizing identity of the object, *i.e.* human, is of paramount importance to aerial surveillance. To identify humans, biological biometric traits, *e.g.* face, periocular, iris, fingerprint, or behavioral biometric traits, *e.g.* gait, keystroke, voice, signature, cognitive, have been investigated. Aerial human identification is emerging quickly as an important area of research as evidenced by the recent call (in 2021) for Biometric Recognition and Identification at Altitude and Range (BRIAR) from The Intelligence Advanced Research Projects Activity (IARPA) [63]. The BRIAR program aims to identify or recognize individuals at long-range (*e.g.* 300+ meters), through atmospheric turbulence, or from elevated and/or aerial sensor platforms (*e.g.*  $> 20^\circ$  sensor view angle from watch towers or UAV). However, due to the unique characteristics of aerial footage, human identification from aerial footage is very challenging, even for humans [10], [43].

This section reviews face and gait recognition, and person re-identification in the aerial surveillance setting.

#### 5.1 Aerial Face Recognition

Aerial face recognition has sparked enormous interest in recent years, from both positive and negative perspectives. The prospect that police or army could use drones equipped with facial recognition technology to monitor and recognize the identity of each individual in a protest has raised radical concerns about privacy [53]. While facial recognition technology of military and polices is usually kept secret, a multitude of commercial, federal and academic efforts have

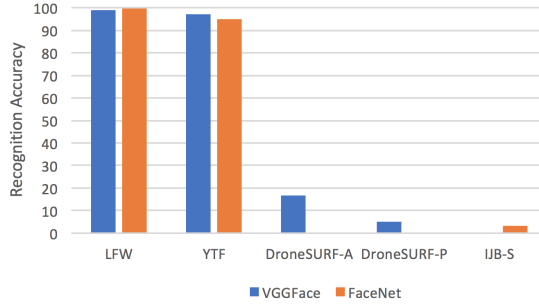


Figure 7: State-of-the-art generic face recognizers drop the performance significantly when shifting to the aerial data. The figure compares the accuracy of VGGFace [108] and FaceNet [122] on two ground-based (LFW and YTF) and two aerial (DroneSURF (Active and Passive) and IJB-S) datasets.

been made to progress in this area. For example, FA6 Drone claimed to identify a face from a distance of 800 meters at a flying altitude of 100 meters [41]. Another example is the IARPA initiated benchmark called IJB-S (IARPA Janus Surveillance Video Benchmark) with surveillance videos captured by cameras mounted on a small fixed-wing UAV [67]. The recent Face In Video Evaluation (FIVE) challenge organized by US National Institute of Science and Technology (NIST) has identification challenges from facial images captured from drones and other aerial vehicles.

#### 5.1.1 Challenges for aerial face recognition

Despite the great success of the generic face recognition methods trained and tested on ground images, an enormous performance drop is observed when they are directly applied to images captured by UAVs [40], [68]. Compared with other tasks, aerial face recognition may exhibit the most drop in performance when migrating from ground to aerial data. Examples of the performance drop of state-of-the-art face recognizers are illustrated in Fig. 7. The VGGFace [108] reduces the accuracy of face recognition from 99.13% in a ground dataset (LFW) to 16.78% in an aerial dataset (DroneSURF-Active) where subjects are cooperative and 4.95% in a dataset (DroneSURF-Passive) where subjects are covertly surveilled [68]. The FaceNet [122] drastically dropped its performance from 99.63% in a ground dataset (LFW) to 3.33% in a aerial dataset (IJB-S) [40].

The unsatisfactory performance is owing to the domain shift caused by the high flying altitude and the camera characteristics. The small size of a face itself [99], [172] tends to amplify all seven challenges discussed in Section §2.1, making aerial face recognition extremely challenging.

#### 5.1.2 Datasets for aerial face recognition

Despite the wide fear of aerial facial recognition with respect to privacy issues, the actual number of public datasets and research in this area is extremely limited. The lack of public datasets may be partially attributed to the privacy concern. There are only three public datasets for this task: DroneFace [60], IJB-S-UAV [67], and DroneSURF [68]. **DroneFace 2017** [60] is not real aerial footage where a commercial sport camera, *i.e.* GoPro Hero3+, was set up at a number of fixed altitudes (1.5, 3, 4.5m) and distances (2:0.5:17m) to simulate

Table 4: Public datasets for aerial face recognition. ‘Con.’, ‘#ID’, ‘#Frm’ and ‘#Vid.’ respectively represent the context where data was collected, the number of identities, frames and videos.

No.	Dataset	Height (m).	Con.	#ID.	#Frm	#Vid.
1	DroneFace [60]	1-5	Campus	11	2,057	-
2	IJB-S [67]	10	Marketplace	202	632	10
3	DroneSURF [68]	-	Various	58	411K	200



(a) DroneSURF



(b) IJB-S

Figure 8: Examples of aerial facial images from two datasets: (a) DroneSURF [68] and (b) IJB-S [67].

the context that a drone seeks lost people on the streets. **IJB-S-UAV 2018** [67] only contains a small number of videos, *i.e.* 10 videos, captured by a fixed-wing UAV flying over a marketplace.

**DroneSURF 2019** [68] is the most realistic dataset for aerial face recognition. It contains 200 videos of 58 subjects, captured across 411K frames, having over 786K face annotations. The dataset demonstrates variations across two surveillance use cases: (i) active and (ii) passive, two locations, and two acquisition times. DroneSURF encapsulates challenges due to the effect of motion, variations in pose, illumination, background, altitude, and resolution, especially due to the large and varying distance between the drone and the subjects. Examples of images from two datasets are shown in Fig. 8.

Details of these datasets can be found in Appendix §D.

#### 5.1.3 Approaches for aerial face recognition

Surprisingly the number of published work on aerial face recognition is scarce. There is only one paper on the DroneFace [26], one paper on the DroneSURF [3], and four papers on the IJB-S-UAV [40], [47], [48], [49]. While targeting aerial facial recognition, all of these had no mechanism to deal with the unique challenges of aerial footage. For example, both [26] and [3] simply employed a classification approach with a cross-entropy loss to classify a probe face into a closed gallery list. All work [47], [48], [49] on the IJB-S-UAV

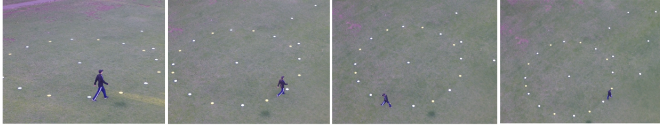


Figure 9: Challenges in aerial gait recognition: large tilt angles and low resolution. Images from the Aerial Gait Dataset [110].

simply investigated different strategies to aggregate frames in a video. However, these early attempts have shown the difficulty of aerial facial recognition. For the IJB-S-UAV dataset, the best model [47] only achieved a rank-1 accuracy of 7.63% in a closed-set setting and 3.13% in an open-set setting. For the DroneSURF dataset, the best model [3] only achieved a recognition accuracy of 24.25% on the active use case and 3.72% on the passive use case. When probe images are carefully cropped, the accuracies can reach 60.87% on the active use case and 45.84% on the passive use case.

## 5.2 Aerial Gait Recognition

Gait, the walking pattern of an individual, is one of important behavioral biometric modalities which can operate at a distance without users' cooperation. Gait recognition has advantages in the aerial scenario, since the dynamic information might be relatively more persistent and evident at low-resolution imagery than the static biometric characteristics such as the face.

### 5.2.1 Challenges for aerial gait recognition

Despite the great progress of generic gait recognition methods trained on ground-to-ground images, the performance of a ground-data-trained model often degrades on aerial data due to the domain shift issue caused by the large discrepancy of pose, resolution and clothing appearance (as shown in Fig. 9).

**Large pose discrepancy:** Most gallery gaits are captured with ground cameras while the probe gaits are from aerial views with large tilt angles, which makes gait recognition challenging.

**Low image resolution:** Given the altitude of the aerial sensor, the probe gait videos are often captured at lower image resolutions, which diminishes the gait information.

**Changing clothing:** This is the classic challenge for gait recognition where the subject in gallery and probe could wear completely different clothing, resulting in large appearance discrepancy.

### 5.2.2 Datasets for aerial gait recognition

There is only one aerial gait recognition dataset in the literature: **Aerial Gait Dataset** [110]. The gait sequences are captured by a drone with heights ranging from 10m to 45m while the subject was walking along a circle. All videos in the dataset are in HD format (1920 × 1080). However, this dataset is of a small scale with only 17 video sequences.

### 5.2.3 Approaches for aerial gait recognition

There are very few existing gait recognition methods specifically working in the aerial domain. [126] proposes a two-step framework which first detects human in a video based on a Single Shot Multi-box Detection (SSD), and then leverages a LSTM-based recurrent processing structure to perform gait recognition. [110] estimates the gait sequence and movement trajectory of a person from an aerial video. The proposed solution consists of perspective correction module, feature extraction module, pose estimation module, and trajectory estimation module.

## 5.3 Aerial Person Re-ID

Person Re-ID aims at retrieving a person of interest across multiple non-overlapping cameras. Person Re-ID seeks to answer questions such as "Where and when has this person been seen in the surveillance network?". The query person can be represented by visual cues, *i.e.* images or videos, or/and textual descriptions [168]. Similar to ground-based surveillance, aerial person re-ID is usually the most practical human recognition tool in unconstrained surveillance conditions where noisy cameras and uncooperative subjects hinder precise measurements of biometric traits such as the face. However, due to the unique characteristics of aerial footage, aerial person re-ID exhibits unique challenges to be addressed.

### 5.3.1 Challenges for aerial person re-ID

Compared with other tasks, person re-id may exhibit the least drop in performance when migrating from ground to aerial data. Examples of the performance drop of state-of-the-art person re-id algorithms are illustrated in Fig. 10. The state-of-the-art Bag of Tricks [87] reduces the accuracy from 94.2% and 89.1% in two ground datasets, Market1501 and DukeMTMC, to 63.4% in the aerial UAV-Human dataset [80]. Similarly, the OSNet [185] reduces the accuracy from 84.9% and 73.5% in two ground datasets, Market1501 and DukeMTMC, to 42.1% in the aerial PRAI-1581 dataset [177].

The performance gap is owing to the domain shift caused by the high flying altitude and the camera characteristics. Aerial Person Re-ID exhibits all seven challenges discussed in Section §2.1. Examples of these challenges are illustrated in Fig. 11.

### 5.3.2 Datasets for aerial person re-ID

One of the first datasets for aerial person re-ID is MRP which was collected by Layne *et al.* [74] in 2014. Since then, due to the affordable availability of off-the-shelf drones with higher resolution cameras, the number of aerial datasets has quickly emerged in the last three years. We summarize public datasets and their statistics in Table 5. Two notable large-scale datasets are PRAI-1581 [177] and UAV-Human [80].

- **PRAI-1581 2020** [177]: Zhang *et al.* recently collected a large dataset for aerial person re-ID. The images were shot by two DJI drones at an altitude ranging from 20 to 60 meters. The dataset consists of 39k images of 1581 unique subjects. The resolution of persons is low, ranging from 30 to 150 pixels. The high flying altitude makes the diversity of views, poses more extreme.

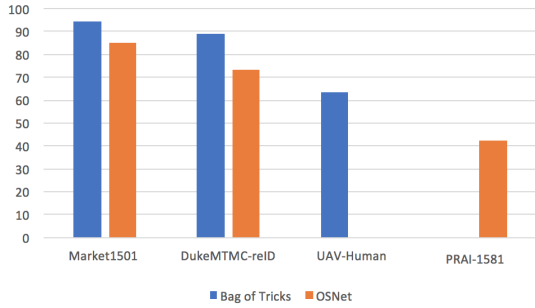


Figure 10: State-of-the-art generic person re-ID algorithms drop the performance when shifting to the aerial data. The figure compares the re-identification accuracy of Bag of Tricks [87] and OSNet [185] on two ground-based (Market1501 and DukeMTMC) and two aerial (UAV-Human and PRAI-1581) datasets.



Figure 11: Challenges for aerial person re-ID due to elevation angle, blur, lighting/illumination, and occlusion. Images from two aerial person re-ID datasets: PRAI-1581 [177] and UAV-Human [80].

- **UAV-Human 2021 [80]:** Li *et al.* just published a new dataset for aerial person re-identification in CVPR 2021. They flew UAVs in multiple urban and rural districts in both daytime and nighttime over three months, hence covering extensive diversities w.r.t. subjects, backgrounds, illuminations, weathers, occlusions, camera motions, and flying altitudes. It contains videos and annotations for multiple tasks, including person re-id and attribute recognition. There are 41,290 frames and 1,144 identities for person re-id and 22,263 frames for attribute recognition. The unique characteristic of the dataset is that multimodal aerial data was captured using a depth sensor (Azure DK), a fisheye camera, and a night-vision camera.

**Long-term Aerial Re-ID:** two datasets, MEVA [24] and p-DESTRE [71] provide footage captured over multiple days where the actors may change in clothing and accessories.

**Hetero Aerial Re-ID IR:** two datasets, UAV-Human [80] and MEVA [24] provide both RGB and thermal IR footage that can be used for aerial person re-ID across RGB and IR spectrum.

**Aerial Human Attribute Recognition:** UAV-Human [80]

Table 5: Public datasets for aerial person re-ID and search. ‘Height.’, ‘#Con.’, ‘#ID’ and ‘#Frm.’ respectively represent the altitude, the context where data was collected, the number of subjects, and frames.

No.	Dataset	Height.	Con.	#ID.	#Frm	#Vid.
1	MRP 2014 [74]	10	Campus	28	4,096	-
2	DroneHIT [67]	25	Campus	101	46K	-
3	p-DESTRE [71]	5-6	Campus	269	14.8M	-
4	PRAI-1581 [177]	20-60	Various	1,581	39K	-
5	UAV-Human [80]	10	Various	1,144	41K	-
6	MEVA [24]	10	Various	100	-	-

provides 7 human attributes (gender, hat, backpack, upper clothing color and style, as well as lower clothing color and style), and p-DESTRE [71] provides 16 human attributes (gender, age, height, body volume, ethnicity, hair colour, hairstyle, beard, moustache, glasses, head accessories’, ‘body accessories’, ‘action’ and ‘clothing information’ (x3)). These attributes can be used for aerial attribute recognition, person search or to support the aerial re-ID task.

Details of these datasets can be found in Appendix §E.

### 5.3.3 Approaches for aerial person re-ID

#### Closed-world Aerial Person Re-ID

Closed world person re-ID, refers to the search where the person in the probe image is definitely present as one of the candidates in the gallery images. The majority of prevailing approaches for closed world person re-ID focus on domain invariant settings, *i.e.* ground-ground and aerial-aerial person re-ID. Classical handcrafted features and classical distance metrics have been prevalent in the early stages of the development of closed world person re-ID techniques.

**Handcrafted feature representation:** Before the explosion of deep learning in 2012, handcrafted features had been employed for aerial person re-ID. Oreifej *et al.* [106] directly estimated the similarity between two human blobs by the Earth Mover Distance (EMD). Layne *et al.* [74] classified human detections into a pre-defined ID list using a number of SVM variants. Schumann *et al.* [123] employed a covariance descriptor [6] and a geodesic distance between two covariance descriptors to measure their similarity.

**Learnable feature representation:** All modern approaches have employed CNNs to learn representation directly from data. There are two categories of applying CNNs in person re-ID: discriminative and pairwise. Discriminative approaches classify each human detections into a pre-defined ID list. Schumann *et al.* [6], [124] designed their own CNN architecture with Inception [133] and Residual layers [56] to learn representation. One of the branch in the Grigorev *et al.*’s framework [54] employed Resnet-50 [56] as a base network. They trained the base network using a large margin Gaussian mixture (L-GM) loss [147]. Pairwise approaches seek to directly calculate the similarity between two human detections, eliminating the need for a pre-collected gallery

list. Most approaches in the literature combine both discriminative losses and pairwise losses to train the backbone network. Grigorev *et al.* [54] trained a backbone ResNet by both a triplet loss and a L-GM loss [147]. Zhang *et al.* also combine a triplet loss with an identification loss to train the deep network. They experimented with multiple backbone networks coupled with a subspace pooling layer to learn a compact representation [177].

*Data augmentation:* Data augmentation during training can be applied to increase the robustness against such error sources as viewpoint variations and occlusions. Moritz *et al.* proposed two augmentation approaches, random rotation (RR) and random cropped rotation (RCR), to specifically improve the robustness against diverse perspectives in UAV-based person re-id [96]. Testing with the OSNet baseline [185], these augmentation techniques can boost the accuracy from 50.0% to 53.0% in the PRAI-1581 [177] dataset and from 80.1% to 83.7% in the p-DESTRE [71] dataset.

*Transfer learning:* Many approaches [54], [123], [124] take advantages of a large number of existing datasets in a classic ground-based setting such as for transfer learning. A person re-ID model is first trained with these ground-based datasets, then either applied directly [123] or fine-tuned [54], [124] on the aerial dataset. While this strategy has shown to be working, it is still an open question to what extent the pre-training can help, since the characteristics of aerial images and ground-based images differ.

#### Open-world Aerial Person Re-ID

Open-world person re-ID is a more difficult problem of person re-ID where the person in the probe may or may not be in the gallery. Many approaches for open world person re-ID depart from the closed-world setting and explore aerial person re-ID across multiple modalities. This includes across domains (aerial-ground [123]), across spectrum (visible-infrared [80]), and across modalities (aerial-attribute [71]).

*Long-term aerial person re-ID:* Long-term aerial person re-ID relaxes the time constraints between the probe and gallery imaging moment. Existing re-ID methods rely on appearance features like clothes, shoes, hair, *etc.* Such features, however, can change drastically from one day to the next, leading to inability to identify people over extended time periods. The p-DESTRE dataset [71] also provides a baseline for long-term aerial person re-id based on facial and body features. The facial feature representation was obtained using the ArcFace model [28] and the body feature representation was obtained using the COSAM [131] model. The Euclidean norm was used as distance function between the concatenated representations. This approach achieved an mAP of 34.9%, but it is noticeable that the p-DESTRE footage was captured with very low flying altitudes, *i.e.* 5-6m.

*Aerial-Ground person re-ID:* Person re-ID between aerial images and ground-based images is of significant interest to surveillance in such tasks as large scale search. However, the differences in human appearances such as views, poses and resolutions make it very challenging. The representation has to be robust to these variations. There exists one work in aerial-ground person re-id by Schumann *et al.* [123] where the authors attempted to encode a robust representation

via covariance descriptors. However, the views of the aerial footage and the ground-based footage are similar. In addition, the dataset is very small with only 1217 probe images and 4244 gallery images [123]. The field is lacking large scale datasets of aerial and ground-based footage for deep networks to learn representation robust to the challenges in views, poses and resolutions. In addition, the field is also lacking approaches to explicitly deal with these challenges. Such generative models such as Generative Adversarial Networks (GANs) [154] can be used to synthesize new views, poses or resolutions for robust matching.

*Aerial-Attribute human recognition:* Searching for persons in aerial data based on attribute descriptions such as gender, height, clothing, *etc.* is also of practical significance in surveillance, which is imperative when the visual image of query person cannot be obtained. The re-id task now can be compiled as a human attribute recognition problem [161]. Kumar *et al.* [71] has addressed this issue by collecting the aerial dataset, p-DESTRE, and annotating persons with 16 attributes. They employed the COSAM algorithm [131] to search for human with facial and body attributes, which can deal with the long-term re-id task.

## 6 AERIAL HUMAN ACTION RECOGNITION

### 6.1 Challenges for aerial action recognition

Despite the great success of the generic human action recognition methods trained on ground-to-ground videos, a huge performance drop is observed when they are directly applied to images captured by UAVs [24], [80]. Examples of the performance drop of the state-of-the-art action recognition algorithm, I3D [19] are illustrated in Fig. 12. While I3D [19] reached 98.0% and 80.9% on two popular ground-based action datasets, UCF-101 [130] and HMDB-51 [70] respectively, it only achieved 23.86%, 28.7% and 16.8% on three aerial action recognition datasets: UAV-Human [80], TinyVIRAT [27], and UCF-Aerial [141].

The unsatisfactory performance is owing to the domain shift caused by the high flying altitude and the camera characteristics. Aerial action recognition exhibits all seven challenges discussed in Section §2.1. Especially, the novel elevation views require the action recognition task to explore more semantic attributes to be invariant to the change in views and poses of subjects. The non-uniform distribution requires effective algorithms to explore robust attention mechanism and attend to partial features to be effective. Examples of these challenges are illustrated in Fig. 13.

### 6.2 Datasets for aerial action recognition

A large number of datasets are captured and annotated for aerial action recognition [8], [80], [111]. Most popular actions of interest are running, walking, sitting, and standing [8], [80], [111]. Among them, many actions are performed in relation to the interaction with objects [8], [80] and other humans [61], [119], [127]. Popular human-object interactions are reading, carrying, pushing/pulling [8], fishing and cutting tree [80], entering and exiting a vehicle [24], [105]. Popular human-human interactions can be in forms of violence [111], [127], social interactions [80]. Two most notable aerial

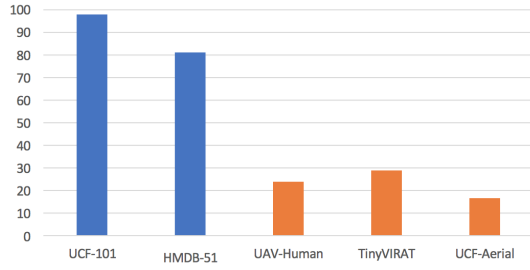


Figure 12: State-of-the-art generic action recognition algorithms drop the performance significantly when shifting to the aerial data. The figure compares the accuracy of I3D [19] on two ground-based (UCF-101 and HMDB-51) and three aerial (UAV-Human, TinyVIRAT and UCF-Aerial) datasets.

action and interaction datasets are UAV-Human [80] and MEVA [24].

- *UAV-Human* [80]: beside the data and annotations for person re-id and attribute recognition as discussed before, the UAV-Human team has also compiled 67K multi-modal video sequences and 119 subjects for action recognition. The unique characteristic in multi-modal aerial data including a depth sensor (Azure DK), a fisheye camera, and a night-vision camera makes this dataset very attractive to the research community.
- *MEVA* [24]: this collection observed approximately 100 actors performing scripted scenarios and spontaneous background activity over a three-week period at an access-controlled venue, collecting in multiple modalities with overlapping and non-overlapping indoor and outdoor viewpoints. The resulting data includes videos from 38 RGB and thermal IR cameras, 42 hours of UAV footage, as well as GPS locations for the actors.

**Task-specific:** many aerial human detection datasets are aimed for specific tasks such as sport [23], search and rescue [94], synthetic data from game engines [132], and multiview [113]. AVI [127] is for violent recognition from aerial videos.

**Beyond visible:** other sensors have also been employed in parallel with visible cameras. UAV-Human [80] employed depth and infrared sensors in parallel to the visible cameras. MEVA [24] employed infrared sensors simultaneously with the visible cameras.

**Fine-grained gesture recognition:** some datasets are collected to perform gesture recognition from aerial images/videos, which are usually employed for controlling drones/UAVs, such as UAV-Gesture [112] and DDIR [7]. Due to the fine-grained recognition, the images/videos are usually captured at short distances and low flying altitude.

**Large scale crowd behaviors:** Compared with ground-based surveillance configurations, the aerial configuration is much relevant to group activity understanding and crowd analysis considering the large view of aerial sensors [98]. Datasets such as DroneCrowd [155] and DLR-ACD [5] focus on crowd counting and density estimation while ERA [98] classifies aerial videos into action classes.

Details of these datasets can be found in Appendix §F.

Table 6: Public datasets for aerial human action recognition. ‘Spec.’, ‘#Con.’, ‘#Vid’ and ‘#Box.’ respectively represent for the action type (A: Action and Interaction, G: Gesture, C: Group event/activity crowd analysis), the context where data was collected, the number of videos, identities and bounding boxes.

No.	Dataset	Lev.	Con.	#Vid.	#ID	#Hei.	#Act.
1	UAV-Human [80]	A	Urban	67K	119	-	155
2	Drone-Action [111]	A	YT	240	10	8-12	13
3	Okutama-Action [8]	A	Urban	43	9	10-45	12
4	UCF-ARG [141]	A	Carpark	480	12	-	10
5	UCF-Aerial [140]	A	Carpark	480	12	-	10
6	MEVA [24]	A	Urban	4.6h	100	-	37
7	MOD20 [113]	A	YT	2K	-	-	20
8	VIRAT [105]	A	Carpark	17	-	-	23
9	NEC-Drone [23]	A	Gym	5K	19	-	16
10	Game-Action [132]	A	Game	200	-	-	7
11	AVI [127]	A	Urban	2K	25	2-8	5
12	DIASR [94]	A	Disaster	7K	-	10-40	-
13	Youtube-Aerial [132]	A	YT	400	-	-	6
14	UT-Interaction [119]	A	Forest	60	6	-	6
15	UAV-Gesture [112]	G	Forest	119	10	-	13
16	DDIR [7]	G	Various	-	26	2-10	9
17	ERA [98]	C	YT	2.8K	-	-	25
18	DLR-ACD [5]	C	YT	2.8K	-	-	25
19	VisDroneCC [187]	C	YT	2.8K	-	-	25

### 6.3 Approaches for aerial action recognition

This section of the report discusses current and state-of-the-art approaches to action recognition trained on the aerial or drone datasets discussed in the previous section.

*Single-frame classification:* A number of approaches took advantage of existing 2D classification networks to perform classification of single frames and subsequently fuse the classification outputs of these frames in a video [46], [95], [98]. Similar to other tasks, ResNet and InceptionNet are among the most popular networks in such work as [46], [94], [95], [98]. More advanced networks such as MobileNet and DenseNet have also been used [7], [44], [98]. Mou *et al.* experimented multiple networks on the ERA dataset [98] and showed that DenseNet achieved the best performance for the aerial action recognition task. Multiple fusion approaches have been employed subsequently such as majority votes [95] or LSTM-based [46].

*Two-stream CNNs:* A number of researchers have employed two-stream CNNs for aerial action recognition since they are a natural way to combine appearance and motion [8], [111], [113]. These two-stream CNNs can also be extended to improve the representation capacity. Perera *et al.* employed a subspace representation called KRP-FS to improve representation [113]. Perera *et al.* employed a pose network



Figure 13: Challenges for aerial action recognition due to the small size of subjects, elevation angle, extreme lighting/illumination, and occlusion. Images from two aerial datasets: UAV-Human [80] and Okutama-Action [8].

to detect body keypoints, cropping the input into multiple patches to focus explicitly on body parts for more accurate recognition [111], [112]. Body keypoints can also be used as inputs for a conventional classifiers such as SVM for aerial action classification [127].

**3D CNNs:** 3D CNNs are still the most popular networks for aerial action recognition. Among the modern networks, I3D [19] has been widely adopted for aerial action recognition [23], [27], [80], [98], [132]. C3D [139] has also been utilized for aerial action recognition [24], [98]. Others have also upgraded existing 2D networks such as Inception-ResNet [56], [133] with 3D convolutions to make the suitable for video processing [109]. Mou *et al.* experimented multiple 3D CNNs, *i.e.* C3D, I3D, P3D, TRN, on the ERA dataset [98] and showed that TRN [183] achieved the best action recognition [98]. Ding *et al.* [32] replaced the backbone network of TSN [150] with a lightweight network, *i.e.* MobileNet [59], coupled with a focal loss and the modern self-attention mechanism in Transformer [143].

## 6.4 Techniques to solve aerial action recognition challenges

The distinct aerial challenges discussed in Section §6.1 constitute a large number of fine-grained domains, across which an effective aerial action recognition model has to stay robust.

### 6.4.1 Low-resolution

The most challenging factor in aerial action recognition is the low resolution or small size of objects due to the high flying altitude. A human may appear as tiny as ten pixels in an aerial video. For example, the size of humans in the TinyVIRAT dataset [27] only accounts for less than 0.01% of a video spatial size, compared with 0.15% in the classic ground-based UCF-101 dataset [130]. The most explicit approach for the low resolution challenge is super-resolution [103]. Demir *et al.* [27] proposed to employ a super-resolution network first to increase the resolution of low-resolution aerial videos before action classification. They proposed a dedicated foreground branch in parallel to the video super-resolution network to better focus on foreground objects and humans, which are more critical for accurate action recognition. The proposed approach improved the F1 score

of the baseline I3D [19] from 28.73% to 34.49% on the aerial dataset TinyVIRAT [27].

### 6.4.2 Lack of data

Due to the high cost of capturing and labeling large scale aerial videos with diverse actions, the aerial action recognition community has addressed this challenge by domain adaptation [23] and knowledge distillation [32]. Choi *et al.* [23] proposed domain adaptation approaches to leverage existing annotated action datasets and unannotated aerial videos. For same source and target label datasets, they proposed unsupervised domain adaption with a cross-entropy loss for classification and a domain adversarial loss for cross-domain learning. For different source and target datasets, the authors employed a triplet loss instead of the cross-entropy loss. Domain adaptation from UCF-101 [130] would increase the accuracy of the I3D baseline up to 70% on the aerial dataset NEC-Drone [23]. Ding *et al.* [32] proposed to use I3D [19] pre-trained on the Kinetics dataset as a teacher network and learned a student network to fit both the softmax output of the teacher network and the aerial data labels. This knowledge distillation scheme improved the recognition accuracy of the I3D baseline from 78.85% to 87.81% on their own UAV action recognition dataset [32].

### 6.4.3 Viewing angle

The unique top and vertical views of aerial footage are challenging for action recognition. To deal with the lack of these views, Sultani *et al.* [132] proposed to use game engines such as GTA-5 and FIFA to synthesize action videos where multiple views of one action can be easily generated. The authors also proposed using conditional Wasserstein Generative Adversarial Networks (WGANs [4]) to generate additional aerial images from real ground features. Additional images with aerial views help to improve the recognition accuracy of the baseline from 49.7% to 68.2%.

### 6.4.4 Fish-eye cameras

Fisheye cameras are widely used in UAVs due to their ultra-wide view angles; however, their angles would cause severe distortion in videos, which is much more challenging for action recognition than in conventional RGB cameras. Inspired by Spatial Transformer Network [64], Li *et al.* [80] proposed a guided transformer module, GT-Module, which

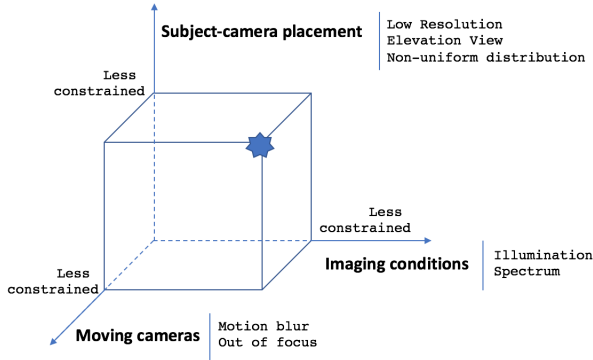


Figure 14: Three challenge dimensions of aerial surveillance: low resolution, vertical perspective, and moving cameras.

can be inserted immediately before maxpooling layers in the original I3D [19] to warp “pixel” in the featuremap extracted from fisheye distorted aerial videos by learning a series of unbounded transformation. This scheme improved the action recognition accuracy on fisheye aerial videos from 20.76% to 23.24% on the UAV-Human aerial dataset.

## 7 AN OUTLOOK: AERIAL SURVEILLANCE IN NEXT ERA

### 7.1 Challenges and Under-investigated Tasks

To fully exploit the potential of aerial surveillance, the challenges of unconstrained and uncooperative environment have to be addressed. The high flying altitudes of airborne platforms on which surveillance cameras and sensors are mounted results in unique settings of aerial surveillance compared to conventional ground-based surveillance in regards to subject-camera placement, imaging conditions, and moving cameras, as illustrated in Fig. 14. These settings cause the key challenges in small resolution, elevation view, non-uniform distribution, illumination, motion blur, and out of focus data. In the in-depth analysis in four surveillance tasks, we have shown that generic networks drop the performance drastically when shifting to the aerial data. Early approaches have been attempted to tackle these challenges to improve the performance of the aerial surveillance tasks.

Our review of existing work has shown that four aerial surveillance tasks, *i.e.* detection, tracking, re-id and action recognition, have received most attention since they do not require the fine details of human of interest in the scene to perform well. There are a wide range of resources, including public datasets, papers and even competitions available addressing these tasks. In contrast, the human identification task is the least investigated task arising from the small size of face biometric when captured from an aerial platform. The task of recognizing faces from long distance with cameras in a moving aerial platform is extremely challenging and requires further significant research efforts to achieve satisfactory operational performance. On the other hand we make a somewhat surprising observation that the relatively easier task of gait precognition from aerial data has received significantly less attention.

### 7.2 Open Issues

We discuss the open issues from two perspectives: data and model development and model deployment.

#### 7.2.1 Data

From a data perspective, three key aspects needing further research are: scale, diversity, and annotation of data.

**Large-scale Aerial Surveillance Data:** ImageNet has demonstrated the great importance of leveraging large, comprehensive, and challenging benchmarks to advance computer vision tasks. There is also a critical need for *large-scale public datasets* in the aerial surveillance tasks. More large-scale accurately annotated datasets akin to VisDrones and UAV-Human with more diversified imaging conditions will help to boost the attention and research.

**Heterogeneous Aerial Surveillance Data:** Current data collections from aerial platforms have been mainly using monocular RGB data even though there has been a few multi-modal collection of RGB and thermal. To address the main challenges of aerial surveillance including the reduced number of pixels per sq.cm, extreme view angles, poor visibility, and motion blur, multi-modal fusion techniques could be explored. Multiple heterogeneous modalities beyond visible RGB cameras such as radars, LiDARs, multispectral and hyperspectral sensors would generate rich information for surveillance tasks.

**Human Annotation Minimization:** Collecting large scale aerial data has been challenging, but labeling a large amount of data compounds this challenge as the process is very expensive and labor-intensive. Besides unsupervised learning techniques, active learning and learning from virtual data can also help to minimize annotation requirement. However it will require additional techniques, *e.g.* domain adaptation and knowledge distillation, to adapt to the new data.

#### 7.2.2 Model Development

From a model perspective, there are three key aspects that need further research: model architecture, effects of unconstrained data on model trustworthiness, prediction uncertainties and model explainability.

**Model Architecture:** Many new network architectures such as Transformers [34] and Res2Net [45], are emerging through computer vision tasks. It would be interesting to explore their performance in the aerial surveillance domain.

**Effects of Unconstrained Data on Model Trustworthiness:** The unconstrained nature of aerial surveillance settings results in low quality data with noise. In the current literature, there is no work to report their performance on adverse conditions of blur arising from atmospheric turbulence and motion, out of focus, low light, illumination and weather conditions such as rain and fog. The ramifications of these unconstrained data acquisition conditions are further accentuated by the large stand-off distances and the uncooperative nature of the human targets under surveillance. Conventional image and video enhancement techniques would be of limited value in aerial surveillance due to the high levels of data impairments and further focused research efforts in this direction is needed to mitigate the effects of unconstrained data on model trustworthiness. The

recent open call in 2021 by the US government [63] for the development of BRIAR is an example of the importance of the need for focused research to unlock the full capabilities and potential of aerial surveillance.

**Adversarial Attacks on Aerial Surveillance Tasks:** Recent research has proven that machine learning models are vulnerable to adversarial attacks, both digitally and physically. Understanding the robustness of aerial surveillance models against these attacks, e.g. [35], is critical to the security of these models.

**Prediction of Uncertainties in Aerial Surveillance:** While many models are emerging, there does not exist any approach to quantify how certain each model is of its predictions of the output. Two sources of uncertainties are from data and model. For example, out of distribution data usually leads to adversarial attacks [1]. For high stake application such as aerial surveillance, it is important for a model to notify when the output is not reliable.

**Explainability of Models in Aerial Surveillance:** Most current surveillance models in general and aerial surveillance models in particular are opaque and only output a scalar prediction value, there is no method providing users semantically understandable explanations for why the model predicts what it predicts. The blackbox-ness could lead to unexpected and detrimental outcomes. Explainable AI (XAI) techniques such as [21] could help to make these models more transparent and understandable to humans.

### 7.2.3 Deployment-aware Model Development

It is important to design efficient and adaptive models to address scalability issue for practical model deployment.

**Lightweight models:** Deep learning models usually contain tens to hundreds of millions parameters with hundreds of layers, e.g. ResNet-101 has 44.6M parameters and 347 layers [56]. Designing lightweight models is an important step in deploying resource-constrained computing platforms. Lightweight CNNs employ advanced model compression techniques such as pruning, quantization and knowledge distillation [29] to efficiently trade-off between resource and accuracy, minimizing their model size and computations in term of the number of floating point operations (FLOPs), while retaining high accuracies.

**Resource-Constrained Model Design:** Most of the current aerial surveillance systems are designed to work either off-line or data is live streamed and processed at the ground stations. There is an urgent need for them to be run on-board. Onboard computing would enable rapid response to surveillance events, reduce the cost and complexity of communication with ground stations, reduce power consumption and improve autonomy of the aerial surveillance systems. However, the limited resources of onboard computing hardware should be taken into account when designing the onboard models. To deal with this, differential constrained design approaches [102] should be explored.

## 8 CONCLUDING REMARKS

Human-centric aerial surveillance, with its advantages in scale, mobility, deployment and observation, presents new capabilities to enable surveillance in geographically-difficult

conditions. The emerging of recent public datasets and initiatives such as VisDrones has further boosted attention and research from the academic community. However, the research in this area is still very daunting. The physical imaging conditions of airborne sensors leads to new challenges in small resolutions, multiple scales, extreme views, motion blur, atmospheric turbulence, uniformity of distribution, illumination and noise, and calls for new innovations from the computer vision signal processing and machine learning community. This paper presents the very first comprehensive survey with in-depth analysis to understand the state of four key aerial surveillance tasks: detection, tracking, identification (including re-identification) and behavior analysis. While there are still many open issues which require significant more research to progress, the future outlook of aerial surveillance is very promising. We believe this survey will provide important guidance for future aerial surveillance research.

## 9 ACKNOWLEDGMENT

This research was supported by an Australian Research Council (ARC) Discovery grant DP200101942.

## REFERENCES

- [1] M. Abdar, F. Pourpanah, S. Hussain, D. Rezazadegan, L. Liu, M. Ghavamzadeh, P. Fieguth, X. Cao, A. Khosravi, U. R. Acharya, V. Makarenkov, and S. Nahavandi. A review of uncertainty quantification in deep learning: Techniques, applications and challenges. *Information Fusion*, 76:243–297, 2021.
- [2] B. M. Albaba and S. Ozer. Synet: An ensemble network for object detection in UAV images. In *ICPR*, 2020.
- [3] G. Amato, F. Falchi, C. Gennaro, F. V. Massoli, and C. Vairo. Multi-resolution face recognition with drones. In *International Conference on Sensors, Signal and Image Processing*, 2020.
- [4] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein generative adversarial networks. In *ICML*, 2017.
- [5] R. Bahmanyar, E. Vig, and P. Reinartz. MRCNet: Crowd Counting and Density Map Estimation in Aerial and Ground Imagery. In *BMVC Workshop*, 2019.
- [6] S. Bak and F. Brémond. Re-identification by covariance descriptors. In *Person Re-Identification*, 2014.
- [7] O. L. Barbed, P. Azagra, L. Teixeira, M. Chli, J. Civera, and A. C. Murillo. Fine grained pointing recognition for natural drone guidance. In *CVPR Workshop*, 2020.
- [8] M. Barekatin, M. Martí, H. Shih, S. Murray, K. Nakayama, Y. Matsuo, and H. Prendinger. Okutama-action: An aerial view video dataset for concurrent human action detection. In *CVPR Workshop*, 2017.
- [9] G. Bhat, M. Danelljan, L. V. Gool, and R. Timofte. Learning discriminative model prediction for tracking. In *ICCV*, 2019.
- [10] M. Bindemann, M. C. Fysh, S. S. K. Sage, K. Douglas, and H. M. Tummon. Person identification from aerial footage by a remote-controlled drone. *Nature Scientific reports*, 7(1):13629–10, 2017.
- [11] A. Bochkovskiy, C. Wang, and H. Liao. Yolov4: Optimal speed and accuracy of object detection. *arXiv*, 2004.10934, 2020.
- [12] E. Bondi, R. Jain, P. Aggrawal, et al. Birdsai: A dataset for detection and tracking in aerial thermal infrared videos. In *WACV*, 2020.
- [13] M. Bonetto, P. Korshunov, G. Ramponi, and T. Ebrahimi. Privacy in mini-drone based video surveillance. In *FG*, 2015.
- [14] I. Bozcan and E. Kayacan. Au-air: A multi-modal UAV dataset for low altitude traffic surveillance. In *ICRA*, pages 8504–8510, 2020.
- [15] D. Bozic-Stulic, Z. Marusic, and S. Gotovac. Deep learning approach in aerial imagery for supporting land search and rescue missions. *IJCV*, 127(9):1256–1278, 2019.
- [16] P. Burdzykowski. A novel method for the deblurring of photogrammetric images using conditional generative adversarial networks. *Remote Sensing*, 12(16), 2020.
- [17] Z. Cai and N. Vasconcelos. Cascade R-CNN: Delving Into High Quality Object Detection. In *CVPR*, 2018.

- [18] Y. Cao, G. Wang, D. Yan, and Z. Zhao. Two algorithms for the detection and tracking of moving vehicle targets in aerial infrared image sequences. *Remote Sensing*, 8(1):28, 2016.
- [19] J. Carreira and A. Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 2017.
- [20] C. Chen, Y. Zhang, Q. Lv, S. Wei, X. Wang, X. Sun, and J. Dong. Rrnet: A hybrid detector for object detection in drone-captured images. In *ICCV Workshop*, 2019.
- [21] L. Chen, Z. Liu, and M. Zhang. Explainable person re-identification with attribute-guided metric distillation. In *ICCV*, 2021.
- [22] Y. Chen, J. Li, Y. Niu, and J. He. Small object detection networks based on classification-oriented super-resolution gan for UAV aerial imagery. In *Chinese Control And Decision Conference (CCDC)*, 2019.
- [23] J. Choi, G. Sharma, M. Chandraker, and J.-B. Huang. Unsupervised and Semi-Supervised Domain Adaptation for Action Recognition from Drones. In *WACV*, 2020.
- [24] K. Corona, K. Osterdahl, R. Collins, and A. Hoogs. Meva: A large-scale multiview, multimodal video dataset for activity detection. In *WACV*, 2021.
- [25] M. Danelljan, G. Bhat, F. Shahbaz Khan, and M. Felsberg. Eco: Efficient convolution operators for tracking. In *CVPR*, 2017.
- [26] A. Deeb, K. Roy, and K. D. Etoh. Drone-based face recognition using deep learning. In *Advanced Machine Learning Technologies and Applications*, 2021.
- [27] U. Demir, Y. S. Rawat, and M. Shah. Tinyvirat: Low-resolution video action recognition. In *ICPR*, 2021.
- [28] J. Deng, J. Guo, N. Xue, and S. Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *CVPR*, 2019.
- [29] L. Deng, G. Li, S. Han, L. Shi, and Y. Xie. Model compression and hardware acceleration for neural networks: A comprehensive survey. *Proceedings of the IEEE*, 108(4):485–532, 2020.
- [30] S. Deng, S. Li, K. Xie, W. Song, X. Liao, A. Hao, and H. Qin. A global-local self-adaptive network for drone-view object detection. *IEEE Transactions on Image Processing*, 30:1556–1569, 2021.
- [31] K. Ding, G. He, H. Gu, Z. Zhong, S. Xiang, and C. Pan. Train in dense and test in sparse: A method for sparse object detection in aerial images. *IEEE Geoscience and Remote Sensing Letters*, pages 1–5, 2020.
- [32] M. Ding, N. Li, Z. Song, R. Zhang, X. Zhang, and H. Zhou. A lightweight action recognition method for unmanned-aerial-vehicle video. In *IEEE International Conference on Electronics and Communication Engineering*, pages 181–185, 2020.
- [33] A. W. Dorn. Aerial surveillance: Eyes in the sky. In *Air Power in UN Operations: Wings for Peace*, chapter 7, pages 119–134. Ashgate Publishings, UK, 2014.
- [34] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021.
- [35] A. Du, B. Chen, T. Chin, Y. W. Law, M. Sasdelli, R. Rajasegaran, and D. Campbell. Physical adversarial attacks on an aerial imagery object detector. *CoRR*, abs/2108.11765, 2021.
- [36] D. Du et al. VisDrone-DET2020: The Vision Meets Drone Object Detection in Image Challenge Results. In *ECCV Workshop*, 2020.
- [37] D. Du, Y. Qi, H. Yu, Y. Yang, K. Duan, G. Li, W. Zhang, Q. Huang, and Q. Tian. The UAV benchmark: Object detection and tracking. In *ECCV*, 2018.
- [38] D. Du, P. Zhu, L. Wen, X. Bian, H. Ling, Q. Hu, J. Zheng, T. Peng, X. Wang, Y. Zhang, et al. Visdrone-sot2019: The vision meets drone single object tracking challenge results. In *ICCV Workshop*, 2019.
- [39] K. Duan, S. Bai, L. Xie, H. Qi, Q. Huang, and Q. Tian. Centernet: Keypoint triplets for object detection. In *ICCV*, 2019.
- [40] J. A. Duncan, N. D. Kalka, B. Maze, and A. K. Jain. End-to-end protocols and performance metrics for unconstrained face recognition. In *ICB*, 2019.
- [41] Face-Six. Fa6 face recognition drone skysapience. Face-Six, 2018.
- [42] H. Fan, L. Wen, D. Du, P. Zhu, Q. Hu, H. Ling, M. Shah, B. Wang, B. Dong, D. Yuan, et al. Visdrone-sot2020: The vision meets drone single object tracking challenge results. In *ECCV*, 2020.
- [43] M. Fysh and M. Bindemann. Person identification from drones by humans: Insights from cognitive psychology. *Drones*, 2(4):32, 2018.
- [44] H. Gammulle, S. Denman, S. Sridharan, and C. Fookes. Predicting the future: A jointly learnt model for action anticipation. In *ICCV*, 2019.
- [45] S. H. Gao, M. M. Cheng, K. Zhao, X. Y. Zhang, M. H. Yang, and P. Torr. Res2net: A new multi-scale backbone architecture. *T-PAMI*, 43(2):652–662, 2021.
- [46] R. Galdes, A. Gonçalves, T. Lai, M. Villerabel, W. Deng, A. Salta, K. Nakayama, Y. Matsuo, and H. Prendinger. UAV-based situational awareness system using deep learning. *IEEE Access*, 7:122583–122594, 2019.
- [47] S. Gong, Y. Shi, and A. K. Jain. Low Quality Video Face Recognition: Multi-Mode Aggregation Recurrent Network (MARN). In *ICCV Workshop*, 2019.
- [48] S. Gong, Y. Shi, A. K. Jain, and N. D. Kalka. Recurrent embedding aggregation network for video face recognition. *arXiv*, 1904.12019, 2019.
- [49] S. Gong, Y. Shi, N. D. Kalka, and A. K. Jain. Video Face Recognition: Component-wise Feature Aggregation Network (CFAN). In *ICB*, 2019.
- [50] Y. Gong, X. Yu, Y. Ding, X. Peng, J. Zhao, and Z. Han. Effective fusion factor in fpn for tiny object detection. In *WACV*, 2021.
- [51] Z. Gong and D. Li. Towards Better Object Detection in Scale Variation with Adaptive Feature Selection. *arXiv*, 2012.03265, 2020.
- [52] F. Granelli, C. Sacchi, R. Bassoli, R. Cohen, and I. Ashkenazi. A dynamic and flexible architecture based on UAVs for border security and safety. In *Advanced Technologies for Security Applications*, 2020.
- [53] F. Greenwood. Can a police drone recognize your face? *Slate*, 2020.
- [54] A. Grigorev, Z. Tian, S. Rho, J. Xiong, S. Liu, and F. Jiang. Deep person re-identification in UAV images. *EURASIP Journal on Advances in Signal Processing*, 2019(1):54 – 64, 2019.
- [55] A. Haider, F. Shaukat, and J. Mir. Human detection in aerial thermal imaging using a fully convolutional regression network. *Infrared Physics and Technology*, 116, 2021.
- [56] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [57] Y. He, C. Fu, F. Lin, Y. Li, and P. Lu. Towards robust visual tracking for UAV with tri-attentional correlation filters. In *IROS*, 2020.
- [58] S. Hong, S. Kang, and D. Cho. Patch-level augmentation for object detection in aerial images. In *ICCV Workshop*, 2019.
- [59] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv*, 1704.04861, 2017.
- [60] H.-J. Hsu and K.-T. Chen. DroneFace: An Open Dataset for Drone Research. In *ACM on Multimedia Systems Conference*, 2017.
- [61] W. Hu, H. Liu, Y. Du, C. Yuan, B. Li, and S. J. Maybank. Interaction-aware spatio-temporal pyramid attention networks for action classification. *T-PAMI*, pages 1–1, 2021.
- [62] Z. Huang, C. Fu, Y. Li, F. Lin, and P. Lu. Learning aberrance repressed correlation filters for real-time UAV tracking. In *CVPR*, 2019.
- [63] I. A. R. P. A. (IARPA). Biometric Recognition and Identification at Altitude and Range (BRIAR), 2020.
- [64] M. Jaderberg, K. Simonyan, A. Zisserman, and k. kavukcuoglu. Spatial transformer networks. In *NeurIPS*, 2015.
- [65] M. Jain, A. Tyagi, A. V. Subramanyam, S. Denman, S. Sridharan, and C. Fookes. Channel graph regularized correlation filters for visual object tracking. *IEEE Transactions on Circuits and Systems for Video Technology*, pages 1–10, 2021.
- [66] R. Jin and D. Lin. Adaptive anchor for fast object detection in aerial image. *IEEE Geoscience and Remote Sensing Letters*, 17(5):839–843, 2020.
- [67] N. D. Kalka, B. Maze, J. A. Duncan, K. O'Connor, S. Elliott, K. Hebert, J. Bryan, and A. K. Jain. Ijb-s: Iarpa janus surveillance video benchmark. In *BTAS*, 2018.
- [68] I. Kalra, M. Singh, S. Nagpal, R. Singh, M. Vatsa, and P. B. Sujit. Dronesurf: Benchmark dataset for drone-based face recognition. In *FG*, pages 1–7, 2019.
- [69] B. Kiefer, M. Messmer, and A. Zell. Leveraging domain labels for object detection from UAVs. *arXiv*, 2101.12677, 2021.
- [70] H. Kuehne, H. Jhuang, E. Garrote, T. Poggio, and T. Serre. Hmdb: A large video database for human motion recognition. In *ICCV*, 2011.
- [71] S. V. A. Kumar, E. Yaghoubi, A. Das, B. S. Harish, and H. Proença. The p-destre: A fully annotated dataset for pedestrian detection, tracking, and short/long-term re-identification from aerial devices. *IEEE Transactions on Information Forensics and Security*, 16:1696–1708, 2021.
- [72] R. LaLonde, D. Zhang, and M. Shah. Clusternet: Detecting small objects in large scenes by exploiting spatio-temporal information. In *CVPR*, 2018.
- [73] H. Law and J. Deng. Cornernet: detecting objects as paired

- keypoints. In *ECCV*, 2018.
- [74] R. Layne, T. Hospedales, and S. Gong. Investigating open-world person re-identification using a drone. In *ECCV Workshop*, 2015.
- [75] A. Leipnitz, T. Strutz, and O. Jokisch. Spatial resolution-independent CNN-based person detection in agricultural image data. In *Interactive Collaborative Robotics*, pages 189–199, 2020.
- [76] B. Li, C. Fu, F. Ding, J. Ye, and F. Lin. All-day object tracking for UAVs. *arXiv preprint arXiv:2101.08446*, 2021.
- [77] C. Li, T. Yang, S. Zhu, C. Chen, and S. Guan. Density map guided object detection in aerial images. In *CVPR Workshop*, 2020.
- [78] F. Li, Y. Yao, P. Li, D. Zhang, W. Zuo, and M.-H. Yang. Integrating boundary and center correlation filters for visual tracking with aspect ratio variation. In *ICCV Workshop*, 2017.
- [79] S. Li and D.-Y. Yeung. Visual object tracking for UAVs: A benchmark and new motion models. In *AAAI*, 2017.
- [80] T. Li, J. Liu, W. Zhang, Y. Ni, W. Wang, and Z. Li. UAV-Human: A Large Benchmark for Human Behavior Understanding with UAVs. In *CVPR*, 2021.
- [81] Y. Li, C. Fu, F. Ding, Z. Huang, and G. Lu. Autotrack: Towards high-performance visual tracking for UAV with automatic spatio-temporal regularization. In *CVPR*, 2020.
- [82] X. Liang, J. Zhang, L. Zhuo, Y. Li, and Q. Tian. Small object detection in UAV images using feature fusion and scaling-based single shot detector with spatial context analysis. *IEEE Transactions on Circuits and Systems for Video Technology*, 30(6):1758–1770, 2020.
- [83] H. Lin, J. Zhou, Y. Gan, C.-M. Vong, and Q. Liu. Novel up-scale feature aggregation for object detection in aerial images. *Neurocomputing*, 411:364–374, 2020.
- [84] T. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017.
- [85] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár. Focal loss for dense object detection. In *ICCV*, 2017.
- [86] Y. Liu, P. Sun, N. Wergeles, and Y. Shang. A survey and performance evaluation of deep learning methods for small object detection. *Expert Systems with Applications*, 172:114602, 2021.
- [87] H. Luo, Y. Gu, X. Liao, S. Lai, and W. Jiang. Bag of Tricks and a Strong Baseline for Deep Person Re-Identification. In *CVPRW*, 2019.
- [88] Y. Ma, X. Wu, G. Yu, Y. Xu, and Y. Wang. Pedestrian detection and tracking from low-resolution UAV thermal imagery. *Sensors*, 16(4):446, 2016.
- [89] Y. Ma, X. Wu, G. Yu, Y. Xu, and Y. Wang. Pedestrian detection and tracking from low-resolution UAV thermal imagery. *Sensors*, 16(4):446, 2016.
- [90] Z. Marusic, D. Bozic-Stulic, S. Gotovac, and T. Marusic. Region proposal approach for human detection on aerial imagery. In *International Conference on Smart and Sustainable Technologies*, 2018.
- [91] S. M. Marvasti-Zadeh, J. Khaghani, H. Ghanei-Yakhdan, S. Kasaei, and L. Cheng. COMET: context-aware IoU-guided network for small object tracking. In *ACCV*, 2020.
- [92] M. Messmer, B. Kiefer, and A. Zell. Gaining scale invariance in UAV bird's eye view object detection by adaptive resizing. *arXiv*, 2101.12694, 2021.
- [93] A. H. Michel. *Eyes in the Sky*. Mariner Books, 2019.
- [94] B. Mishra, D. Garg, P. Narang, and V. Mishra. Drone-surveillance for search and rescue in natural disaster. *Computer Communications*, 156:1–10, 2020.
- [95] H. Mliki, F. Bouhlef, and M. Hammami. Human activity recognition from UAV-captured video sequences. *Pattern Recognition*, 2020.
- [96] L. Moritz, A. Specker, and A. Schumann. A study of person re-identification design characteristics for aerial data. In *Pattern Recognition and Tracking XXXII*, 2021.
- [97] R. Mottaghi, X. Chen, X. Liu, N. Cho, S. Lee, S. Fidler, R. Urtasun, and A. Yuille. The role of context for object detection and semantic segmentation in the wild. In *CVPR*, 2014.
- [98] L. Mou, Y. Hua, P. Jin, and X. X. Zhu. Event and activity recognition in aerial videos using deep neural networks and a new dataset. In *IEEE International Geoscience and Remote Sensing Symposium*, 2020.
- [99] S. P. Mudunuri and S. Biswas. Low resolution face recognition across variations in pose and illumination. *T-PAMI*, 38(5):1034–1040, 2016.
- [100] M. Mueller, N. Smith, and B. Ghanem. A benchmark and simulator for UAV tracking. In *ECCV*, pages 445–461, 2016.
- [101] A. Newell, K. Yang, and J. Deng. Stacked hourglass networks for human pose estimation. In *ECCV*, 2016.
- [102] K. Nguyen, C. Fookes, and S. Sridharan. Constrained design of deep iris networks. *IEEE Transactions on Image Processing*, 29:7166–7175, 2020.
- [103] K. Nguyen, C. Fookes, S. Sridharan, M. Tistarelli, and M. Nixon. Super-resolution for biometrics: A comprehensive survey. *Pattern Recognition*, 78:23–42, 2018.
- [104] D. Oberhaus. Giant surveillance balloons are lurking at the edge of space. *WIRED*, 2019.
- [105] S. Oh, A. Hoogs, A. Perera, et al. A large-scale benchmark dataset for event recognition in surveillance video. In *CVPR*, 2011.
- [106] O. Oreifej, R. Mehran, and M. Shah. Human identity recognition in aerial images. In *CVPR*, 2010.
- [107] J. Pang, K. Chen, J. Shi, H. Feng, W. Ouyang, and D. Lin. Libra R-CNN: Towards Balanced Learning for Object Detection. In *CVPR*, 2019.
- [108] O. M. Parkhi, A. Vedaldi, and A. Zisserman. Deep face recognition. In *BMVC*, 2015.
- [109] H. Peng and A. Razi. Fully autonomous UAV-based action recognition system using aerial imagery. In *Advances in Visual Computing*, pages 276–290, 2020.
- [110] A. G. Perera, Y. W. Law, and J. Chahl. Human pose and path estimation from aerial video using dynamic classifier selection. *Cognitive Computation*, 10(6):1019–1041, 2018.
- [111] Y. W. Perera, Asanka G. and Law and J. Chahl. Drone-action: An outdoor recorded drone video. *Drones*, 2019.
- [112] Y. W. Perera, Asanka G. and Law, J. Chahl, and T. T. Ogunwa. UAV-GESTURE: A Dataset for UAV Control and Gesture Recognition. In *ECCV Workshop*, 2018.
- [113] Y. W. Perera, Asanka G. and Law, J. Chahl, and T. T. Ogunwa. A multiviewpoint outdoor dataset for human action recognition. *IEEE transaction on Human-Machine Systems*, 2020.
- [114] G. Plastiras, C. Kyrkou, and T. Theodoridis. Efficient convnet-based object detection for UAVs by selective tile processing. *arXiv*, 1911.06073, 2019.
- [115] J. Portmann, S. Lynen, M. Chli, and R. Siegwart. People detection and tracking from aerial thermal views. In *ICRA*, 2014.
- [116] N. N. Raj and A. S. Vijay. Adaptive blind deconvolution and denoising of motion blurred images. In *RTEICT*, 2016.
- [117] S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *NeurIPS*, 2015.
- [118] A. Robicquet, A. Sadeghian, A. Alahi, and S. Savarese. Learning social etiquette: Human trajectory understanding in crowded scenes. In *ECCV*, 2016.
- [119] M. S. Ryoo and J. K. Aggarwal. Spatio-Temporal Relationship Match: Video Structure Comparison for Recognition of Complex Human Activities. In *ICCV*, 2009.
- [120] S. Sambolek and M. Ivacic-Kos. Automatic person detection in search and rescue operations using deep CNN detectors. *IEEE Access*, 9:37905–37922, 2021.
- [121] D. C. Schedl, I. Kurmi, and O. Bimber. Search and rescue with airborne optical sectioning. *Nature Machine Intelligence*, 2020.
- [122] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In *CVPR*, 2015.
- [123] A. Schumann and J. Metzler. Person re-identification across aerial and ground-based cameras by deep feature fusion. In *Automatic Target Recognition XXVII*, 2017.
- [124] A. Schumann and J. Metzler. Adapted deep feature fusion for person re-identification in aerial images. In *Autonomous Systems: Sensors, Vehicles, Security, and the Internet of Everything*, 2018.
- [125] J. Scott and J. Wiener. Surveillance oversight. *epic.org*, 2020.
- [126] J. P. T. Sien, K. H. Lim, and P.-I. Au. Deep learning in gait recognition for drone surveillance system. In *Materials Science and Engineering*, 2019.
- [127] A. Singh, D. Patil, and S. Omkar. Eye in the sky: Real-time drone surveillance system (dss) for violent individuals identification using scatternet hybrid deep learning network. In *CVPR Workshop*, 2018.
- [128] G. Soldi, D. Gaglione, N. Forti, et al. Space-based global maritime surveillance. part i: Satellite technologies. *IEEE Aerospace and Electronic Systems Magazine*, 36:6 – 28, 2020.
- [129] W. Song, S. Li, T. Chang, A. Hao, Q. Zhao, and H. Qin. Cross-view contextual relation transferred network for unsupervised vehicle tracking in drone videos. In *CVPR*, 2020.
- [130] K. Soomro, A. R. Zamir, and M. Shah. UCF101: A dataset of 101 human actions classes from videos in the wild. *arXiv*, 1212.0402, 2012.
- [131] A. Subramaniam, A. Nambiar, and A. Mittal. Co-segmentation inspired attention networks for video-based person re-identification. In *ICCV*, 2019.
- [132] W. Sultani and M. Shah. Human action recognition in drone videos using a few aerial training examples. *CVIU*, 206:103186, 2021.
- [133] C. Szegedy, Wei Liu, Yangqing Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going

- deeper with convolutions. In *CVPR*, 2015.
- [134] M. Tan, R. Pang, and Q. Le. EfficientDet: Scalable and efficient object detection. In *CVPR*, 2020.
- [135] Z. Tang, X. Liu, and B. Yang. Penet: Object detection using points estimation in high definition aerial images. In *ICMLA*, 2020.
- [136] Z. Tian, C. Shen, H. Chen, and T. He. FCOS: Fully convolutional one-stage object detection. In *ICCV*, 2019.
- [137] Z. Tian, C. Shen, H. Chen, and T. He. FCOS: A simple and strong anchor-free object detector. *T-PAMI*, 2021.
- [138] K. Tong, Y. Wu, and F. Zhou. Recent advances in small object detection based on deep learning: A review. *Image and Vision Computing*, 97:103910, 2020.
- [139] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, 2015.
- [140] UCF. UCF-Aerial Action Data Set.
- [141] UCF. UCF-ARG Data Set.
- [142] F. O. Unel, B. O. Ozkalayci, and C. Cigla. The power of tiling for small object detection. In *CVPR Workshop*, 2019.
- [143] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. Attention is all you need. In *NeurIPS*, 2017.
- [144] R. Volpi, P. Morerio, S. Savarese, and V. Murino. Adversarial feature augmentation for unsupervised domain adaptation. In *CVPR*, 2018.
- [145] A. Von Moll, E. Garcia, D. Casbeer, M. Suresh, and S. Chandra Swar. Multiple-pursuer, single-evader border defense differential game. *Journal of Aerospace Information Systems*, 17(8):407–416, 2020.
- [146] A. Walha, A. Wali, and A. M. Alimi. Video stabilization with moving object detecting and tracking for aerial video surveillance. *Multimedia Tools and Applications*, pages 6745–6767, 2016.
- [147] W. Wan, Y. Zhong, T. Li, and J. Chen. Rethinking feature distribution for loss functions in image classification. In *CVPR*, 2018.
- [148] H. Wang, Z. Wang, M. Jia, A. Li, T. Feng, W. Zhang, and L. Jiao. Spatial attention for multi-scale feature refinement for object detection. In *ICCV Workshop*, 2019.
- [149] J. Wang, K. Sun, T. Cheng, B. Jiang, C. Deng, Y. Zhao, D. Liu, Y. Mu, M. Tan, X. Wang, W. Liu, and B. Xiao. Deep high-resolution representation learning for visual recognition. *T-PAMI*, 2019.
- [150] L. Wang, Y. Xiong, Z. Wang, Y. Qiao, D. Lin, X. Tang, and L. Van Gool. Temporal segment networks for action recognition in videos. *T-PAMI*, 2019.
- [151] R. J. Wang, X. Li, and C. X. Ling. Pelee: A real-time object detection system on mobile devices. In *NeurIPS*, 2018.
- [152] Y. Wang, L. Ding, and R. Laganier. Real-time UAV tracking based on PSR stability. In *ICCV Workshop*, 2019.
- [153] Y. Wang, Y. Yang, and X. Zhao. Object detection using clustering algorithm adaptive searching regions in aerial images. In *ECCV Workshop*, 2020.
- [154] Z. Wang, Q. She, and T. E. Ward. Generative adversarial networks in computer vision: A survey and taxonomy. *ACM Computing Surveys*, 54(2):1–38, 2021.
- [155] L. Wen, D. Du, P. Zhu, Q. Hu, Q. Wang, L. Bo, and S. Lyu. Detection, tracking, and counting meets drones in crowds: A benchmark. In *CVPR*, 2021.
- [156] L. Wen, P. Zhu, D. Du, X. Bian, H. Ling, Q. Hu, C. Liu, H. Cheng, X. Liu, W. Ma, et al. Visdrone-sot2018: The vision meets drone single-object tracking challenge results. In *ECCVW*, 2018.
- [157] Y. Wu, J. Lim, and M.-H. Yang. Object tracking benchmark. *T-PAMI*, 37(9):1834–1848, 2015.
- [158] Z. Wu, K. Suresh, P. Narayanan, H. Xu, H. Kwon, and Z. Wang. Delving into robust object detection from UAVs: A deep nuisance disentanglement approach. In *ICCV*, 2019.
- [159] J. Xu, W. Wang, H. Wang, and J. Guo. Multi-model ensemble with rich spatial information for object detection. *Pattern Recognition*, 99:107098, 2020.
- [160] J.-P. Yaacoub, H. Noura, O. Salman, and A. Chehab. Security analysis of drones systems: Attacks, limitations, and recommendations. *Internet of Things*, 11:100218–, 2020.
- [161] E. Yaghoubi, F. Khezeli, D. Borza, S. A. Kumar, J. Neves, and H. Proença. Human attribute recognition - a comprehensive survey. *Applied Sciences*, 10(16), 2020.
- [162] C. Yang, Z. Huang, and N. Wang. Querydet: Cascaded sparse query for accelerating high-resolution small object detection. *arXiv*, 2103.09136, 2021.
- [163] F. Yang, H. Fan, P. Chu, E. Blasch, and H. Ling. Clustered object detection in aerial images. In *ICCV*, 2019.
- [164] G. Yang and D. Ramanan. Volumetric correspondence networks for optical flow. *NeurIPS*, 2019.
- [165] X. Yang and J. Yan. Arbitrary-oriented object detection with circular smooth label. *ECCV*, 2020.
- [166] X. Yang, J. Yang, J. Yan, Y. Zhang, T. Zhang, Z. Guo, X. Sun, and K. Fu. Srdet: Towards more robust detection for small, cluttered and rotated objects. In *ICCV*, 2019.
- [167] J. Ye, C. Fu, F. Lin, F. Ding, S. An, and G. Lu. Multi-regularized correlation filter for UAV tracking and self-localization. *IEEE Transactions on Industrial Electronics*, 2021.
- [168] M. Ye, J. Shen, G. Lin, T. Xiang, L. Shao, and S. C. H. Hoi. Deep learning for person re-identification: A survey and outlook. *T-PAMI*, 2021.
- [169] H. Yu, G. Li, W. Zhang, Q. Huang, D. Du, Q. Tian, and N. Sebe. The UAV benchmark: Object detection, tracking and baseline. *IJCV*, 128(5):1141–1159, 2020.
- [170] W. Yu, T. Yang, and C. Chen. Towards resolving the challenge of long-tail distribution in UAV images for object detection. In *WACV*, 2021.
- [171] X. Yu et al. The 1st tiny object detection challenge: Methods and results. In *ECCV Workshop*, 2020.
- [172] X. Yu, B. Fernando, R. Hartley, and F. Porikli. Semantic face hallucination: Super-resolving very low-resolution face images with supplementary attributes. *T-PAMI*, 42(11):2926–2943, 2020.
- [173] X. Yu, Y. Gong, N. Jiang, Q. Ye, and Z. Han. Scale match for tiny person detection. In *WACV*, 2020.
- [174] J. Zhang, J. Huang, X. Chen, and D. Zhang. How to fully exploit the abilities of aerial image detectors. In *ICCV Workshop*, 2019.
- [175] R. Zhang, Z. Shao, X. Huang, J. Wang, and D. Li. Object detection in UAV images via global density fused convolutional network. *Remote Sensing*, 12(19), 2020.
- [176] S. Zhang, C. Chi, Y. Yao, Z. Lei, and S. Z. Li. Bridging the gap between anchor-based and anchor-free detection via adaptive training sample selection. In *CVPR*, 2020.
- [177] S. Zhang, Q. Zhang, Y. Yang, X. Wei, P. Wang, B. Jiao, and Y. Zhang. Person re-identification in aerial imagery. *IEEE Transactions on Multimedia*, 2021.
- [178] S. Zhang, L. Zhuo, H. Zhang, and J. Li. Object tracking in UAV videos via multifeature discrimination and instance-aware attention network. *Remote Sensing*, 12(16):2646, 2020.
- [179] W. Zhang, K. Song, X. Rong, and Y. Li. Coarse-to-fine UAV target tracking with deep reinforcement learning. *IEEE Transactions on Automation Science and Engineering*, 16(4):1522–1530, 2018.
- [180] X. Zhang, E. Izquierdo, and K. Chandramouli. Dense and small object detection in UAV vision based on cascade network. In *ICCV Workshop*, 2019.
- [181] X. Zhang, F. Wan, C. Liu, R. Ji, and Q. Ye. Freeanchor: Learning to match anchors for visual object detection. In *NeurIPS*, 2019.
- [182] Y. Zhang, D. Zhou, S. Chen, S. Gao, and Y. Ma. Single-image crowd counting via multi-column convolutional neural network. In *CVPR*, 2016.
- [183] B. Zhou, A. Andonian, A. Oliva, and A. Torralba. Temporal relational reasoning in videos. In *ECCV*, 2018.
- [184] J. Zhou, P. Wang, and H. Sun. Discriminative and robust online learning for siamese visual tracking. In *AAAI*, 2020.
- [185] K. Zhou, Y. Yang, A. Cavallaro, and T. Xiang. Omni-scale feature learning for person re-identification. In *ICCV*, 2019.
- [186] X. Zhou, V. Koltun, and P. Krähenbühl. Tracking objects as points. In *ECCV*, 2020.
- [187] P. Zhu, L. Wen, D. Du, X. Bian, Q. Hu, and H. Ling. Vision meets drones: Past, present and future. *arXiv*, 2001.06303, 2021.
- [188] Z. Zou, Z. Shi, Y. Guo, and J. Ye. Object detection in 20 years: A survey. *arXiv*, 1905.05055, 2019.

## APPENDIX

### APPENDIX A

#### AIRBORNE PLATFORMS AND IMAGING SENSORS FOR AERIAL SURVEILLANCE

A wide range of airborne platforms are available with diverse characteristics in flying ranges and altitudes, endurance, speed, manoeuvrability, payload and vulnerabilities as summarized in Table 7. These airborne platforms are deployed with a wide range of sensors. In this paper, we only consider techniques for analyzing data acquired by RGB imaging sensors, which capture data with spatial details to be analyzed by computer vision and deep learning. Imaging sensors are categorized by their spectrum as summarized in Table 8.

### APPENDIX B

#### DETAILS ON AERIAL HUMAN DETECTION DATASETS

##### VisDrone 2020

The VisDrone dataset [187] is collected by the AISKYEYE team at Tianjin University, China. This dataset is behind three “Vision meets drones” challenges in ECCV 2018, CVPR 2019 and ECCV 2020. The benchmark dataset consists of 400 video clips formed by 265,228 frames and 10,209 static images, captured by various drone-mounted cameras, covering a wide range of aspects including location (14 different cities separated by thousands of kilometers), environment (urban and country), objects (pedestrian, vehicles, bicycles, *etc.*), and density (sparse and crowded scenes). The dataset was collected using various drone platforms (*i.e.* drones with different models), in different scenarios, and under various weather and lighting conditions. These frames are manually annotated with more than 2.6 million bounding boxes or points of targets of frequent interests, such as pedestrians, cars, bicycles, and tricycles. Some important attributes including scene visibility, object class and occlusion, are also provided for better data utilization.

##### TinyPersons 2020

The TinyPerson dataset [173] is collected from Internet. This dataset is behind two “Tiny Object Detection” challenges in ICCV 2019 and ECCV 2020. The unique characteristic of this dataset is the tiny resolution of humans. A majority of human instances appear as tiny as [2,20] pixels, and as small as [20,32] pixels. In total, there are 72,651 objects with bounding boxes have been manually annotated. The dataset is split into a training set, which has 794 labeled images with 42,197 annotations, and a testing set, which has 816 labeled images with 30,454 annotations.

##### AU-AIR 2020

The AU-AIR dataset [14] is the first multi-modal UAV dataset for object detection collected by Aarhus University, Denmark. It meets vision and robotics for UAVs having the multi-modal data from different on-board sensors (*i.e.* visual, time, location, altitude, IMU, velocity). The AU-AIR dataset has more than 2 hours raw videos, with 32,823 labeled frames and 132,034 object annotations from 8 object categories related to traffic surveillance. Each frame is also labeled with time, GPS, IMU, altitude, linear velocities of the UAV.

##### BIRDSAI 2020

The BIRDSAI dataset [12] is a long-wave thermal infrared dataset containing nighttime images of animals and humans in Southern Africa. The dataset allows for benchmarking of algorithms for automatic detection and tracking of humans and animals with both real and synthetic videos. There are 48 real aerial TIR videos and 124 synthetic aerial TIR videos (generated with AirSim), for a total of 62k and 100k images, respectively. The dataset breaks these into labels of animals or humans, and also provides species information when possible, including for elephants, lions, and giraffes. Information about noise and occlusion for each bounding box is also included.

##### StanfordDrones 2016

The StanfordDrones dataset [118] is collected by Stanford University, originally aimed to predict human trajectory in crowded scenes but the labels can be employed for aerial human and vehicle detection. The dataset consists of eight unique scenes, recorded in a university campus. The dataset comprises more than 19K targets consisting of 11.2K pedestrians, 6.4K bicyclists, 1.3k cars, 0.3K skateboarders, 0.2K golf carts, and 0.1K buses.

##### UAV123 2016

The UAV123 dataset [100] is collected by King Abdullah University of Science and Technology, Saudi Arabia, originally aims for low altitude UAV target tracking. There are 123 videos with 112,578 fully annotated frames. The UAV123 dataset contains 3 subsets: (i) Set1 contains 103 sequences captured using an off-the-shelf professional-grade UAV (DJI S1000) following different objects at altitudes varying between 5–25 m. (ii) Set2 contains 12 sequences captured from a boardcam (with no image stabilization) mounted to a small low-cost UAV following other UAVs. (iii) Set3 contains 8 synthetic sequences captured by the proposed UAV simulator.

##### Mini-drone 2015

The Mini-drone dataset [13] is collected by the Ecole Polytechnique Fédérale de Lausanne (EPFL), originally aimed for drone-based surveillance, helping in managing parking spaces, controlling crowds and reporting useful information such as suspicious behaviors, mis-parked cars, number of free parking spots, *etc.* The created dataset consists of 38 different contents captured in full HD resolution, with a duration of 16 to 24 seconds each, shot with the mini-drone Phantom 2 Vision+ in a parking lot. The dataset contents can be clustered in three categories: normal, suspicious, and illicit behaviors. The low altitude and high resolution of videos are sufficient for body silhouette, face detection, vehicle license plate and accessories such as bags, backpack, sunglasses, hat, wallet or bottle.

##### HERIDAL 2018

The HERIDAL database [15], [90] is collected to support search and rescue missions using drones. It contains over 68,750 image patches of wilderness acquired from an aerial perspective, 29,050 positive samples containing person as well as 39,700 negative samples. Of these, approximately 3,000 image patches are synthetically generated; the others are cropped from real images. Additionally, the HERIDAL database contains approximately 500 labeled, full-size  $4,000 \times 3,000$  pixel real-world images. Full size images

Table 7: Airborne platforms for aerial surveillance.

	Range	Endurance	Speed	Altitude (km)	Manoeuvrability	Payload	Vulnerabilities
Drones	Low	Low	Slow	Low 0.5 km	High	Very low	Easily detected
UAVs	Low to High (1-1,000km)	Low to Medium (15 minutes to 20 hours)	Medium 10-300km/h	Low 5 km	High	Low (1-150kg)	Weather- dependent, esp. wind
Balloons (Free or tethered)	Low (up to 100km per day)	Medium (10 or more days)	Stationary or very slow	Low 5 km	Very low (wind- dependent)	Low to Medium (500kg)	Easily targeted
Rotary aircraft (manned helicopter)	Medium (300km)	Low (3 hours)	Medium (up to 350km/h)	Medium 10 km	Very High (easy turn and stationary capacity)	Medium (10,000kg)	Possible fatalities
Fixed-wing aircraft (manned)	High (10,000km)	Low (15 hours)	High	High 20 km	High (but cannot fly as slowly)	High (250,000kg)	Possible fatalities
Stratollite	High	High (months)	Slow	High 20 km	Low	Low (50kg)	
Satellite	Very High (but has fixed trajectory)	Very High (years but revisit time can be days)	Very High (25,000km/h)	Very High 100-1,000 km	Low (only certain types)	Medium (5,000kg)	Limited availability at specific time and place

Table 8: Imaging sensors for aerial surveillance.

Sensor Type	Categories	Wavelength	Resolution	Distance	Visibility	Weather Affected	Passive	Output	Cost
Electro-Optical	Visible	0.39-0.75um	High	Medium	Medium	High		2Dx3	Low
	Infrared	5-14um	Low	Long	Medium	Low		2Dx1	Medium
	Hyperspectral	0.4-2.5um			Medium	Medium		2DxM	Very high
Radio-Wave	Radar	1mm-300cm	Low	Very Long~ ~100km	Low	Low	x	2Dx1*	Medium
	LiDAR	0.9-1.5um	Medium		High	Medium	x	3Dx1	Very high

could be used for training (Fast or Faster R-CNN) as well as for testing purposes. At the moment 101 images have been selected for testing purposes.

#### SARD 2021

The SARD database [120] is also collected to support search and rescue missions using a DJI Phantom 4A drone. The videos were recorded at a resolution of  $1,920 \times 1,080$  at 50Hz. The drone flew at different altitudes ranging from 5 to 50 meters. The dataset comprises 1,981 manually labeled images from 9 actors.

#### AgriDrone 2021

AgriDrone is a self-captured data set with focus on person detection in agricultural applications. All 4,586 images have been captured by two different drones: DJI Mavic2 Enterprise and DJI Mavic Pro between Spring and Winter. They share the same resolution of  $3840 \times 2160$  pixels. The data set is split into 70% training, 10% validation and 20% test data. Due to the rural application area, the average number of humans per image is about two only, which results in a smaller dataset than in the case of the VisDrone data, while the humans are usually bigger.

## APPENDIX C

### DETAILS ON AERIAL PERSON TRACKING DATASETS

**UAV123:** UAV123 dataset [100] contains 123 fully annotated HD sequences over 110K frames taken from UAV platforms. Each video has 12 attribute categories. A video may have a variety of attributes by the shooting conditions. The captured targets include *pedestrian*, vehicles, boats, groups and *etc.* The video resolution is between 720p and 4K.

**Campus:** The large-scale campus dataset [118] has images and videos of various classes of targets that move and interact in a real-world university campus. The dataset comprises over 19K targets consisting of 11.2K *pedestrian*, 6.4K bicyclists, 1.3k cars, 0.3K skateboarders, 0.2K golf carts, and 0.1K buses. The resolution of these UAV videos is  $1920 \times 1080$ .

**DTB70:** Drone Tracking Benchmark (DTB) [79] collects 70 videos of UAV collected data where the bounding boxes are manually annotated. Some of the videos are captured in a campus, which are for the purpose of tracking *pedestrian*, animals and vehicles. In order to enhance the diversity of the scene, some videos are collected from YouTube. The original resolution of each video frame is  $1280 \times 720$ .

**VisDrone:** The VisDrone team has compiled a dedicated large-scale drone benchmark and organized challenges for tracking, *i.e.* VisDrone-SOT2018 [156] and VisDrone-SOT2019 [38]. The VisDrone-SOT2018 consists of 132 videos with 106K frames. Compared with VisDrone-SOT2018, VisDrone-SOT2019 introduces 35 new sequences. To further increase the diversity of videos and assess the performance of trackers in the wild, VisDrone-SOT2020 [42] conducts extensive evaluation of more tracking algorithms using the same dataset in VisDrone-SOT2019. VisDrone2021 further increases the dataset size to 400 videos with more diverse scenarios. The annotated targets of VisDrone datasets include *pedestrians*, dogs, bicycles, vehicles, *etc.*

**BIRDSAI:** BIRDSAI [12] is the first aerial dataset captured by a thermal infrared camera. This dataset includes 48 real videos with a variety of attributes, such as motion

blur, large camera motions, background clutter and high altitude. The dataset comprises over 154K targets consisting of 120K bounding boxes of wild animals (*i.e.* giraffes, lions, elephants, *etc.*) and about 34K *human* bounding boxes.

**UAVDark135:** UAVDark135 dataset [76] contains fully annotated 135 videos captured by a standard UAV at night. The benchmark includes various tracking scenes, *e.g.* crossings, t-junctions, road, highway, and consists of different kinds of tracked objects like *pedestrian*, boat, bus, car, truck, athletes, house, *etc.* To extend the covered scenes, the benchmark also contains some YouTube videos, which are shot on the sea. The total frames, mean frames, maximum frames, and minimum frames of the benchmark are 125, 466, 929, 4571, and 216 respectively, making it suitable for large-scale evaluation. The videos are captured with the resolution of  $1920 \times 1080$ . at a frame-rate of 30 frames/s (FPS),

## APPENDIX D

### DETAILS ON AERIAL FACE RECOGNITION DATASETS

**DroneFace 2017 [60]:** The DroneFace is an open dataset to simulate the context that a drone seeks lost people on the streets, and tries to recognize the specific target from the air based on the face recognition model established from a few portrait photos. The aerial images are captured by a commercial sports camera (GoPro Hero3+) mounted on a UAV flying at different altitude (1.5, 3, 4, and 5 meters) and captured still images of subjects from 17 meters away from the subjects to 2 meters with 0.5 meters ahead in each step. There are 11 subjects and 2,057 pictures including 620 aerial images, 1,364 frontal face images, and 73 portrait images. The resolution of face regions ranges from  $23 \times 31$  to  $384 \times 384$ . Examples of images from this dataset are shown in Fig. 8.

**IJB-S-UAV 2018 [67]:** The IJB-S dataset is an open-source surveillance video benchmark from the Intelligence Advanced Research Projects Activity (IARPA) to investigate performance of surveillance face recognition. IJB-S stands for IARPA Janus Surveillance Video Benchmark. Aerial facial videos are a part of the dataset. There are 10 UAV videos, captured by a small fixed-wing UAV flying over the collection area, specifically the marketplace, opportunistically capturing surveillance video. There are 5 recognition protocols, one of them is UAV Surveillance-to-Booking, which perform face recognition from UAV video probe to a curated gallery of multiple high resolution mug-shot style photos. Examples of images from this dataset are shown in Fig. 8.

**DroneSURF 2019 [68]:** The DroneSURF dataset is a benchmark dataset to investigate performance of face recognition on aerial footage. DroneSURF stands for Drone Surveillance of Faces. The dataset contains 200 videos of 58 subjects, captured across 411K frames, having over 786K face annotations. The proposed dataset demonstrates variations across two surveillance use cases: (i) active and (ii) passive, two locations, and two acquisition times. DroneSURF encapsulates challenges due to the effect of motion, variations in pose, illumination, background, altitude, and resolution, especially due to the large and varying distance between

the drone and the subjects. Examples of images from this dataset are shown in Fig. 8.

## APPENDIX E

### DETAILS ON AERIAL PERSON RE-ID DATASETS

**MRP 2014 [74]** Layne *et al.* collected one of the very first aerial datasets for person re-ID. They used a standard remote-operated quadcopter to capture videos at a resolution of  $640 \times 360$  at 5Hz. Two sets of data were collected. Set 1 contains three flights across an outdoor and indoor environment. These consists of 436, 652, 848 video frames, from which 233, 471, 797 person detections were obtained from 6, 7, 10 distinct people. Set 2 is significantly larger with 6 flights in three unconstrained and heavily crowded outdoor environments. Across each flight, there are between 10k and 30k frames and an average of 8.6k person detections from an unknown number of distinct people. Of this data, they selected 28 unique identities and 4,096 images.

**DroneHIT 2019 [54]** Grigorev *et al.* also employed a standard remote-operated quadcopter to collect aerial data around a university campus. The drone was flying at a altitude of 25m, capturing videos at a resolution of  $1920 \times 1080$  at 30fps. A total of 101 unique identities are extracted, where each person has about 459 images.

**P-DESTRE 2020 [71]** Kumar *et al.* collected a comprehensive dataset for re-identification across multiple days with the change in clothing of persons. The authors also provided 16 human attributes of person instances for the person search task. The annotations for the attributes include demographic information: gender, ethnicity and age, appearance information: height, body volume, hair color, hairstyle, beard, mustache; accessories information: glasses, head accessories, body accessories; clothing information and action information. The dataset was collected in two universities using DJI Phantom 4 drones at two altitudes: 5.5 and 6.7 meters. The videos are captured at a resolution of  $3,820 \times 2,160$  at 30fps. There are a total of 269 unique subjects with 14.8M person detections. The re-ID task can be performed based on visual input, *i.e.* image or video, or heterogeneously textual input of attributes.

**PRAI-1581 2020 [177]** Zhang *et al.* recently collected a large dataset for aerial person re-ID. The images were shot by two DJI drones at an altitude ranging from 20 to 60 meters. The dataset consists of 39k images of 1581 unique subjects. The resolution of persons is low, ranging from 30 to 150 pixels. The high flying altitude makes the diversity of views, poses more extreme.

**UAV-Human 2021 [80]** Li *et al.* just published a new dataset for aerial person re-identification in CVPR 2021. The dataset was collected by a flying UAV in multiple urban and rural districts in both daytime and nighttime over three months, hence covering extensive diversities w.r.t. subjects, backgrounds, illuminations, weathers, occlusions, camera motions, and flying altitudes. The dataset contains videos and annotations for multiple tasks, including action recognition, pose estimation and person re-identification. There are 41,290 frames and 1,144 identities for person re-identification and 22,263 frames for attribute recognition. The unique characteristic of the dataset is multimodal aerial

data was captured and provided using a depth sensor (Azure DK), a fisheye camera, and a night-vision camera.

**MEVA 2021 [24]** MEVA is a new and very-large-scale dataset for human activity recognition. The collection observed approximately 100 actors performing scripted scenarios and spontaneous background activity over a three-week period at an access-controlled venue, collecting in multiple modalities with overlapping and non-overlapping indoor and outdoor viewpoints. The resulting data includes video from 38 RGB and thermal IR cameras, 42 hours of UAV footage, as well as GPS locations for the actors.

## APPENDIX F

### DETAILS ON AERIAL HUMAN BEHAVIOR DATASETS

**UAV-Human 2021 [80]** Li *et al.* recently published a new dataset for aerial human behavior understanding in CVPR 2021. The dataset was collected by a flying UAV in multiple urban and rural districts in both daytime and nighttime over three months, hence covering extensive diversities w.r.t. subjects, backgrounds, illuminations, weathers, occlusions, camera motions, and flying altitudes. The dataset contains videos and annotations for multiple tasks, including action recognition, pose estimation and person re-identification. There are 67,428 multi-modal video sequences and 119 subjects for action recognition. The unique characteristic of the dataset is multimodal aerial data was captured and provided using a depth sensor (Azure DK), a fisheye camera, and a night-vision camera.

#### Drone-Action 2019

The Drone-Action dataset [111] was collated by a team of academics out of the University of South Australia, this dataset was developed in order to fill the gap of limited outdoor footage. Most datasets were filled with footage captured in doors. Published in 2019, Drone-Action consists of 10 actors performing 13 different actions; Walking, walking side, jogging, jogging side, running, running side, hitting with bottle, hitting with stick, stabbing, punching, kicking, clapping and waving hands. In total 240 videos were collated, consisting of 66,191 frames annotated with body joint estimations and bounding boxes.

#### Okutama-Action 2019

The Okutama-action dataset [8] is a comprehensive dataset published in 2019 out of the University of Tokyo. The dataset consists of fully annotated sequences similar to Drone-action this dataset consists of 12 action classes; hand-shaking, hugging, reading, drinking, pushing/pulling, carrying, calling, running, walking, lying down, sitting and standing. Which are all categorized as either human-human, human-object, or none-interaction. The dataset consists of 43 videos and 77365 frames fully annotated with bounding boxes and body joint estimations. Fig. 15 below depicts a sample from the Okutama dataset.

**UCF-ARG:** action features videos collected at an altitude of 400-450ft (122-137m). Collated in 2008, the videos were captured using a camcorder attached to a blimp [140]. UCF Aerial action is part of the UCF-ARG dataset which consists of footage collected from cameras stationed in different positions; Aerial, Rooftop and ground. It consists of 48 aerial videos of 12 actors performing 10 actions; boxing, carrying,

clapping, digging, jogging, open/closing trunk, running, throwing, walking and waving [141].

**Game-Action 2020:** A unique approach to collecting aerial videos is the use of video game characters instead of real people. UCF along with a team out of ITU Pakistan introduced the first of its kind dataset, games action dataset, which uses human action footage from video games [132]. This footage is of characters in Grand Theft Auto 5 (GTA5) performing 7 different human actions; cycling, fighting, soccer kicking, running, walking, shooting and skydiving. The team used the game FIFA to collect the footage of soccer kicking. The dataset consists of 200 videos, 100 of which are aerial footage. Sample of the aerial footage in the dataset is depicted in Fig. 16 below.

**DIASR 2020:** Search and rescue implementations of drone-based technology has garnered little interest from researchers so far, however, due to the increasing capabilities of drones and requirements to monitor remote areas drone surveillance has become popular. To fill the gap of a lack of search and rescue implementations [94] have proposed an search and rescue dataset; Drone image action dataset for search and rescue (DIASR) which consists of 30,000 frames of HD video and images of multiple actors performing 6 different actions. The actions performed are intended to mimic a person signaling / asking for help and include; waving, standing, sitting, laying and handshaking. As the name suggests the primary reason this dataset was curated was to cultivate the implementation of search and rescue implementations of drones using computer vision. A sample of the dataset is depicted in Fig. 17 below.

**AVI 2020:** The AVI dataset [127] was developed to identify violent individuals in public areas. The complete datasets consist of 2,000 images with 10,863 humans where 5,124 (48%) engaged in one or more of the five violent activities of (1) Punching, (2) Stabbing, (3) Shooting, (4) Kicking, and (5) Strangling. These activities are performed by 25 subjects between the ages of 18–25 years. These images are recorded from the parrot drone at four heights of 2m, 4m, 6m and 8m.

#### Youtube Aerial 2020:

##### UT-interaction 2009

Published in 2009, UT-interaction [119] is a dataset containing 6 different actions similar to Drone-Action; shake-hands, point, hug, push, kick and punch. This dataset is not as comprehensive as Drone-Action, UT-interaction is comprised of 20 videos, shot at 720x480 resolution.

**MOD 2020:** is another dataset produced by the team at the University of South Australia, it is their most recent dataset. The dataset is curated to combat the challenges of the angle / viewpoint of aerial footage and for this reason the dataset fills a gap in the current literature regarding a scarcity of multi viewpoint datasets. The dataset follows in the footsteps of the others in that it is specifically focused on outdoor scenes of action recognition. The dataset is a collection of 2324 video clips, 503,086 frames, sourced from YouTube or recorded on an unspecified drone by the team. [113].

##### UAV-Gesture 2018

The same team out of the University of South Australia produced another dataset, UAV-Gesture [112]. The aim of



Figure 15: Sample frames from the Okutama-action dataset [8].



Figure 16: Sample from the games-action dataset [132].

this dataset was that action recognition datasets where aimed specifically at actions similar to those discussed in the previous dataset. UAV-Gesture was designed to fill the gaps in command signaling/gesturing datasets, the team found there was a lack of data collected in an outdoor environment. Published the year prior to Drone-Action in 2018, UAV-Gesture comprises of 10 actors performing 13 different gestures, all of which are aimed at being able to control/command the drone to maneuver in a programmed way upon recognition of the action. The dataset contains 119 videos with 37151 frames annotated with bounding boxes and body joints estimations.

**DDIR 2020:** Systems that implement human action recognition are typically focused on identifying multiple different actions, such as all the datasets highlighted in this section so far. In contrast to this [7] have proposed an approach that builds a natural interaction system to guide autonomous drones. The approach focuses on fine grained variations of the same human action; pointing and in order to validate their framework, and foster future innovation in the same area they have curated a dataset with variability in user appearance, viewpoint camera distance and scenery. The dataset presented is called Direction Dataset for Interaction with Robots (DDIR) and is divided into 5 subsets denoted DDIR1 to DDIR5. The number of actors range from 5 to 7 with up to 22,628 frames of VGA and FWVGA video. Recorded on an unspecified drone, the dataset consists of 26 different actions, representing a different direction.



Figure 17: Sample from the DIASR dataset [94].

**DLR's Aerial Crowd Dataset (DLR-ACD):** The DLR-ACD dataset [5] is a collection of aerial images for crowd counting and density estimation, as well as for person localization at mass events. It contains 3 large RGB aerial images with average size of  $3619 \times 5226$  pixels acquired through 16 different flight campaigns at various mass events and over urban scenes involving crowds, such as sport events, city centers, open-air fairs and festivals.

The images were recorded using a camera system composed of three standard DSLR cameras (a nadir-looking and two side-looking cameras) mounted on an airborne platform installed on a helicopter flying at an altitude between 500 m to 1600 m with spatial resolution (or ground sampling distance – GSD) ranges from 4.5 to 15 cm/pixel. The dataset was labeled manually with point-annotations on individual people and contains 226,291 person annotations in total, ranging from 285 to 24,368 annotations per image.