

# MITFAS: Mutual Information based Temporal Feature Alignment and Sampling for Aerial Video Action Recognition

Ruiqi Xian\*, Xijun Wang\*, Dinesh Manocha  
University of Maryland - College Park  
{rxian, xijun, dmanocha}@umd.edu

## Abstract

We present a novel approach for action recognition in UAV videos. Our formulation is designed to handle occlusion and viewpoint changes caused by the movement of a UAV. We use the concept of mutual information to compute and align the regions corresponding to human action or motion in the temporal domain. This enables our recognition model to learn from the key features associated with the motion. We also propose a novel frame sampling method that uses joint mutual information to acquire the most informative frame sequence in UAV videos. We have integrated our approach with X3D and evaluated the performance on multiple datasets. In practice, we achieve 18.9% improvement in Top-1 accuracy over current state-of-the-art methods on UAV-Human [30], 7.3% improvement on Drone-Action [41], and 7.16% improvement on NEC Drones [7]. The code is available at <https://github.com/Ricky-Xian/MITFAS>

## 1. Introduction

Unmanned aerial vehicles (UAVs) are increasingly used for different applications, including search and rescue, agriculture, security, construction and aerial surveillance. This results in many challenging perception problems related to detection, tracking, re-identification, and recognition. In particular, action recognition using UAV videos is an important problem. While deep learning based methods [5, 10] have achieved good performance for video action recognition on ground camera videos [5, 38], there are many challenges with respect to using them on aerial videos.

Compared to ground camera videos, the human actors in UAV videos appear rather small due to high camera altitude (see Figure 1). A wider area of the background occupies most of the pixels in the video frame, and only a small fraction (e.g., less than 10%) corresponds to a human action. Since these videos are captured from a moving (or

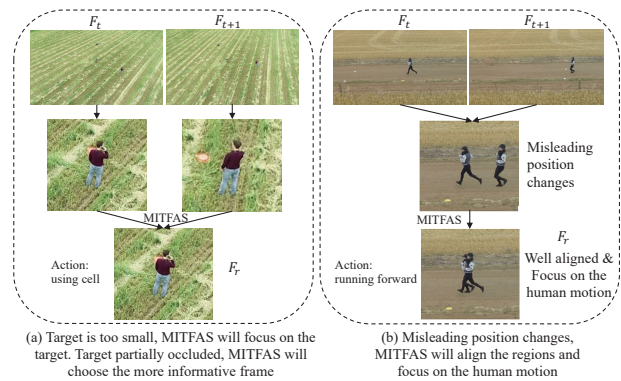


Figure 1.  $F_t$  and  $F_{t+1}$  are two frames at time  $t$  and  $t + 1$ , respectively, from the same UAV video. The human actor in the two frames occupies less than 10% of the pixels due to the high camera altitude (top images). (a) MITFAS will focus on the regions corresponding to salient motions and use mutual information to find the more informative frame. (b) Because of the UAV’s motion, the position of the human actor in  $F_{t+1}$  appears to be relatively behind compared to  $F_t$ . Our algorithm (MITFAS) computes and aligns these regions so that the recognition model will infer more from the human motions. As shown in the right image, the main body of the human actor in two frames overlaps after feature alignment.

dynamic) UAV, the position and orientation of the human actor may change considerably between the frames. This can result in making the model infer more from the background changes, as opposed to action information, during training. The motion of the UAV camera can also result in blurry frames and some techniques have been proposed to handle them [26, 30, 52].

It is harder to collect and annotate UAV videos. Overall, there are fewer and smaller UAV video datasets, as compared to ground video datasets. Additionally, because of continuous changes in the altitude and the camera angle, videos captured using UAVs tend to be more diversified and have unique viewpoints. Some parts of the human actor may be occluded, and not all parts of the human body that contribute to the action can be seen from the camera. Hence, some of the frames in the video are less informative, and this

\*These authors contributed equally to this work

reduces the overall accuracy [15, 44, 50, 53].

**Main Contribution:** We present a novel approach for video action recognition in UAV videos with dynamic backgrounds and moving cameras. We take advantage of the mutual information to obtain and align the useful features corresponding to the human actor in the temporal domain. Our alignment method is used to identify the region of the human action, and find the most similar features in the video sequence. As a result, our learning-based recognition model is able to focus more on the human action, rather than the background regions. Due to the varying viewpoints generated by the movement of a UAV camera, not all the human body parts that contribute to the action are visible. We present a novel frame sampling method based on joint mutual information for dynamic UAV videos, which can compute the most informative and distinctive frame sequence for training aerial action recognition models. We have integrated our temporal feature alignment and frame sampling methods with X3D [10] and use them for aerial action recognition (as shown in Figure 2). The novel components of our work include:

1. We use mutual information as a criterion to obtain and align the features at the same time. Our method takes the movement of the UAV into account and estimates the overlapping features by maximizing the mutual information. Given a reference image frame, our approach finds the most similar features in the subsequent frames.
2. We present a new frame sampling method for UAV videos. Our approach is designed to compute the most informative frame sequence in the video such that all the frames are mostly different from each other. We combine mutual information and joint mutual information to extract the frame. Our method is flexible and can deal with different variations in the video sequences. Extensive experiments show our sampling method overperforms peers.

We test our method on 3 public UAV video datasets. We achieve 20.2% improvement over the baseline method and 18.9% improvement over current state-of-the-art method on UAV-Human [30]. Our method improves the top-1 accuracy on Drone Action [41] by 16.6% over the baseline method and 7.3% over the current state-of-the-art methods. On NEC Drones [7], our method get 78.62% top-1 accuracy, which is 7.18% higher than the current state-of-the-art and 12.47% over baseline model by using 1/2 input frame size.

## 2. Related Work

### 2.1. Temporal Feature Alignment

Temporal alignment-based methods have been extensively studied in various video tasks. For instance, Cao et al. [4] proposed an ordered temporal alignment algorithm specifically designed for few-shot video classification. Similarly, Lu et al. [33] introduced an index-guided framework that utilizes indices to guide pooling and up-sampling operations, while Huang et al. [19] presented a method for learning transformation offsets of pixels to align up-sampled feature maps. Additionally, Huang et al. [20] proposed an aligned feature aggregation algorithm for aligning features of multiple resolutions, and Liu et al. [32] explored feature alignment in the context of multi-frame human pose estimation.

However, it is important to note that most of these methods rely on skeleton information and employ learning-based modules to align important joint points. Despite their success in ground-based videos, these methods may not be as effective in aerial videos due to two primary reasons. Firstly, there is a scarcity of labeled datasets for aerial videos compared to ground videos, making it more challenging to train these learning modules and difficult to transfer them to unseen domains. Secondly, the high flying altitude and moving camera in aerial videos make it challenging to accurately identify important body joint points. To address these limitations, our proposed method takes a different approach by avoiding complex training procedures and not relying on skeleton information.

### 2.2. Similarity Measurement

Various similarity measures have been proposed for comparing image patches. However, in the case of UAV videos with small actor resolution and moving cameras, conventional metrics like Euclidean distance [53] are affected by background changes and shaking frames. Cosine similarity [17], although used for high-dimensional data, neglects pixel value magnitudes. Peak Signal-to-Noise Ratio (PSNR) [18] focuses on pixel-level comparisons but is sensitive to dominant background changes. Structural Similarity Index Measure (SSIM) [45] evaluates luminance, contrast, and structure but is susceptible to structural variations like rotations and shifts commonly found in aerial videos.

Mutual information is used as a similarity measure between images by [34, 47]. As a similarity measure, mutual information has been widely used in the medical imaging domain [25, 43]. Liu et al. [32] have explored the possibility of using mutual information for person pose estimation tasks. Ji et al. [23] have proposed an unsupervised image clustering and segmentation method by maximizing the mutual information between spatial region pairs. Inspired by the success of mutual information for image processing, we

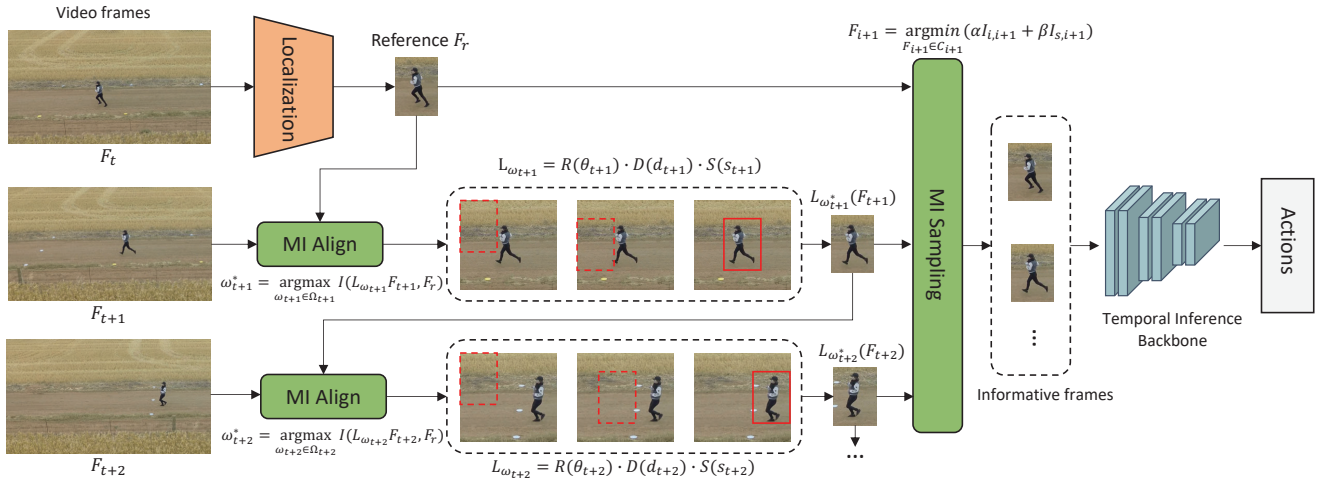


Figure 2. Given a starting frame  $F_t$  in a UAV video, we use a localization network to localize the human action and crop the region containing the human motion as the reference image  $F_r$ . At time  $t + 1$ , we use our feature alignment algorithm to estimate the optimal operation parameter  $\omega_{t+1}^*$  and find a region in  $L_{\omega_{t+1}^*}(F_{t+1}) \subset F_{t+1}$  that the mutual information between  $L_{\omega_{t+1}^*}(F_{t+1})$  and the reference image  $F_r$  is maximized. Next, we use  $L_{\omega_{t+1}^*}(F_{t+1})$  as the new reference image to find the optimal parameter  $\omega_{t+2}^*$  at time  $t + 2$  and repeat for subsequent frames. Then, we use the criterion illustrated in Section 3.2 Eq. 11 to find a sequence of the most distinctive and informative frames. We use a temporal inference backbone network (e.g., X3D [10]) to generate the predicted action label from the spatial-temporal features associated to the sampled frame sequence.

use this concept for temporal feature alignment and frame sampling. Compared to other similarity measures, mutual information measures the statistical dependence or information redundancy between two images using pixel value distributions, which makes it more robust.

### 2.3. Video Recognition for Aerial Videos

Aerial video action recognition is a challenging task, especially when the camera is moving. The performance of action recognition on ground-camera video datasets has increased as a result of recent advancements in deep learning techniques. However, we don't get a similar level of accuracy on videos captured using UAV cameras [40]. For aerial video, [14], [37], [36], [39], [1], [12], [39] apply 2D CNNs (e.g., ResNet, MobileNet) as the backbones to perform single-frame classification and combine the outputs of all frames in the video for recognition. [2], [41], [42] leverage two-stream CNNs to utilize attributes from the human motion and the appearance. [7], [8], [30], [39], [46] use I3D network [5] to learn from spatial-temporal features from human actors and surroundings. To better focus on the target actor in the video, [26,27] have proposed an attention mechanism with Fourier transform for better feature extraction. AZTR [48] proposes a general framework leveraging CNNs and attention mechanisms for aerial action recognition on both edge devices and decent GPUs. Our feature alignment and sampling method could also be combined with these action recognition methods to improve their accuracy.

Given a video captured from a UAV, classic feature rep-

resentation algorithms for aerial video action recognition are limited by the small size of the human actors in aerial videos. Sometimes, these approaches improperly identify the camera's motion as a feature [35, 49]. [21] have proposed 2D affine motion models to approximate the camera motion between the adjacent frames. [24] have proposed a method where the motion patterns of dense trajectories are clustered to characterize foreground-foreground or foreground-background relationships. Inspired by prior works, our method aligns the human-centered views that are transformed from UAV videos to learn from key features corresponding to the parts of the human body that contribute most to the actions.

## 3. Video Recognition using Mutual Information

We present a mutual information-based method for action recognition on UAV videos with moving cameras and dynamic backgrounds. Our method takes the characteristics of the UAV videos into consideration and uses mutual information as the criterion to compute and align the regions that existing salient motions in the video. We use joint mutual information to sample the frame sequences that convey most information about human action. Table 1 highlights the notation and symbols used in this section. We provide the foundational understanding of mutual information in Appendix A, urging readers to review it beforehand for an enhanced grasp of the paper.

Notation	Term
I()	Mutual information, joint mutual information
H()	Entropy, joint entropy
p	Probability mass function
h	Joint histogram
F	Frame sequence in the video
L	Operations to get aligned region
R	Rotation matrix
D	Translation matrix
S	Scaling operation
$\omega_t$	Operations parameters
M	Mapping function from frames to features
C	Candidate pool for frame sampling

Table 1. Notation and symbols used in the paper.

### 3.1. Temporal Feature Alignment

In this section, we describe our approach that uses mutual information (illustrated in Appendix. B) to obtain and align the features that correspond to salient motions in the temporal domain. In UAV videos, human actors appear significantly small in aerial data, and most pixels in the frame belong to the background. Therefore, we have redundant information about the background in the video that may decrease the performance of our learning model. Moreover, the position of the human actor may change considerably between adjacent frames, which makes the recognition model infer more from the pixels corresponding to redundant background information than the human body movements. Thus, our objective is to find the region that contains dominant information about the action for each frame in the video and the pixels related to the human actors are well matched.

Let's assume that all the images have the same 2D image coordinate with the origin positioned in the top left corner, with the  $x$  axis along the rows and  $y$  axis along the columns. Given a video  $V$ , which corresponds to a sequence of raw frames at different times,  $V = \{\cup F_t, t \in N\}$ . We generate the reference image  $F_r$  that is transformed from a region in the raw frame  $F_t$ . The reference image  $F_r$  is a human centred image that mainly contains salient actions of the human actor.

To compute  $F_r$ , suppose  $\Omega_t$  contains all feasible operation parameters  $\omega_t$ , such that for  $\omega_t \in \Omega_t$ , we can generate a region from  $F_t$  using an operation  $L_{\omega_t}$ . We can consider  $L_{\omega_t}$  as a transformation from 2D raw frame coordinates of  $F_t$  to the 2D reference frame coordinates corresponding to  $F_r$ , followed by scaling to the same size of  $F_r$ . Thus,  $L_{\omega_t}$  consists of rotation operation  $R(\theta_t)$ , translation operation  $D(d_t)$  and scaling operation  $S(s_t)$ , where  $\omega_t = (\theta_t, d_t, s_t) \in \Omega_t$ :

$$L_{\omega_t} = R(\theta_t) \cdot D(d_t) \cdot S(s_t) \quad (1)$$

Our objective is to find  $\omega_t^* \in \Omega_t$  for every  $t$  such that:

$$\omega_t^* = \arg \max_{\omega_t \in \Omega_t} I(L_{\omega_t}(F_t); F_r), \quad (2)$$

where

$$I(L_{\omega_t}(F_t); F_r) = H(L_{\omega_t}(F_t)) + H(F_r) - H(L_{\omega_t}(F_t), F_r). \quad (3)$$

We use this equation to compute the optimal parameter  $\omega_t^*$ , so as to compute the target region in  $F_t$  that is aligned with  $F_r$ .

We need to calculate the mutual information between two images  $L_{\omega_t}(F_t)$  and  $F_r$ . There is no exact mathematical model known to precisely calculate the actual probability distributions related to each image. In general, marginal and joint histograms are used [47] to approximate the respective distributions. Let  $v_{\omega_t}(p)$  denote the value of the pixel at position  $p$  in  $L_{\omega_t}(F_t)$  and  $z_{\omega_t}(p)$  the intensity of the corresponding pixel in  $F_r$ . The joint histogram  $h_{\omega_t}(v, z)$  can be computed by binning the values of the pixel pairs  $(v_{\omega_t}(p), z_{\omega_t}(p))$  for all possible  $p$ . We conduct ablation experiments on the impact of bin numbers that are used to generate histogram in Section. C. Then, the marginal probability distribution  $p_{V\omega_t}(v), p_{Z\omega_t}(z)$  and joint probability distribution  $p_{VZ\omega_t}(v, z)$  of  $v$  and  $z$  can be obtained by normalizing the joint histogram  $h_{\omega_t}(v, z)$ :

$$\begin{aligned} p_{VZ\omega_t}(v, z) &= \frac{h_{\omega_t}(v, z)}{\sum_{v, z} h_{\omega_t}(v, z)}, \\ p_{V\omega_t}(v) &= \sum_z p_{VZ\omega_t}(v, z), \\ p_{Z\omega_t}(z) &= \sum_v p_{VZ\omega_t}(v, z). \end{aligned} \quad (4)$$

The mutual information can be calculated as:

$$I(L_{\omega_t}(F_t); F_r) = \sum_{v, z} p_{VZ\omega_t}(v, z) \log \frac{p_{VZ\omega_t}(v, z)}{p_{V\omega_t}(v)p_{Z\omega_t}(z)} \quad (5)$$

Mutual information is computed using histograms of low-level pixel values on both target and reference patches, which is similar to the mean shift tracking. However, our method uses histograms to approximate the joint probability distribution and measure the inherent statistical dependence between target and reference patch. Also, it can be applied at the feature level. We use a feature extractor to get the features for both  $F_t$  and  $F_r$ . Suppose the mapping function between the RGB images to the features is  $M$ , the features extracted from  $F_t$  and  $F_r$  are  $M(F_t)$  and  $M(F_r)$ . Our objective reduces to finding a subset  $M_s(F_t) \subset M(F_t)$  such that

$$M_s(F_t)^* = \arg \max_{M_s(F_t) \subset M(F_t)} I(M_s(F_t); M(F_r)), \quad (6)$$




Figure 3. We sample the  $i + 1$ th frame  $F_{i+1}$  from the candidate pool by choosing the frame that is not only least similar to the previous frame but also all the previously sampled frames.

where

$$I(M_s(F_t); M(F_r)) = H(M_s(F_t)) + H(M(F_t)) - H(M_s(F_t), M(F_t)). \quad (7)$$

### 3.2. Mutual Information Sampling

Because of the high camera altitude, many parts of the human body are not visible. Some parts of the human body that result in the action may be occluded by some other parts that do not contribute to the action. Also, there are lots of "duplicated" frames because of the high frame rate, which essentially contains redundant information. Therefore, not all the video frames are useful for the training, and using some of them may even decrease the overall accuracy. To solve this issue, we present a novel frame sampling method using a combination of mutual information and joint mutual information to find the frame sequences that contain more information about action changes in the UAV videos.

The main idea behind our method is to find out more informative frame sequences in the video given a start frame. Consider a video as a sequence of frames across time. Suppose we have already sampled  $i$  frames and our goal is to find the  $i + 1$ th frame  $F_{i+1}$  in the candidate pool  $C_{i+1}$  where  $C_{i+1}$  consists of all the possible frames that we could choose for  $F_{i+1}$ . Let  $F_s = \{\cup F_0, F_1, F_2 \dots F_i\}$  denote the set that contains all the sampled frames. Our approach is to choose  $F_{i+1}$  that is the most distinctive as compared with  $F_i$  as well as the set of all previously sampled frames so that it provides more unseen features for the recognition model training, see Figure 3:

$$F_{i+1} = \arg \min_{F_{i+1} \in C_{i+1}} \alpha I(F_i; F_{i+1}) + \beta I(F_s; F_{i+1}). \quad (8)$$

The first term is used to minimize the mutual information between the current frame and the previous frame. It tends

to sample adjacent frames that are least similar so that the newly sampled frame will contain more information for training. The second term is used to minimize the joint mutual information with all the sampled frames, which could decrease the information redundancy over the whole sampling sequence. We can decompose it using the chain rule of joint mutual information:

$$\begin{aligned} I(F_s; F_{i+1}) &= I(F_0, F_1, F_2 \dots F_i; F_{i+1}) \\ &= \sum_{j=0}^i I(F_j; F_{i+1} | F_{j-1} F_{j-2} \dots F_0) \end{aligned} \quad (9)$$

In practice, the conditional mutual information is hard to compute as the conditional probability distribution is hard to calculate. However, to make the problem more tractable, we use the low-dimensional approximation to estimate the joint mutual information between  $F_{i+1}$  and  $F_s$  [3, 13].

$$I(F_s; F_{i+1}) \approx \frac{1}{i+1} \sum_{j=0}^i I(F_j; F_{i+1}) \quad (10)$$

So the overall expression becomes:

$$F_{i+1} = \arg \min_{F_{i+1} \in C_{i+1}} \alpha I(F_i; F_{i+1}) + \frac{\beta}{i+1} \sum_{j=0}^i I(F_j; F_{i+1}) \quad (11)$$

Here, we add weights  $\alpha, \beta$  to the two terms in Eq. 11 to adjust to different scenarios. We analyze the behavior of  $\alpha$  and  $\beta$  in the Appendix 4.4.

### 3.3. MITFAS: Aerial Video Recognition

In this section, we present our overall method for aerial video recognition (see Fig. 2). We use temporal feature alignment and frame sampling and combine them with a temporal inference backbone network (e.g, X3D [10]) to disentangle the human actor from superfluous backgrounds and learn from key features associated with the human motions.

In our benchmarks, most of the videos available are captured on a UAV camera with anti-shake technology which could stabilize the camera and reduce the camera vibration, we assume no rotation is needed, i.e.,  $R(\theta_t) = Identity$ . For general videos,  $R(\theta_t)$  is the rotation matrix represented and computed as a 2D transformation:

$$R(\theta_t) = \begin{bmatrix} \cos \theta_t & -\sin \theta_t \\ \sin \theta_t & \cos \theta_t \end{bmatrix} \quad (12)$$

We localize the human actor at the start frame and enlarge the region by about 10% of its height to obtain the reference  $F_r$  [16]. We conduct ablation studies on the size of  $F_r$  in Section. C. Considering the human actor may perform actions that have large vertical changes like stretching

Method	Backbone	Frames Number	Input Size	Initialization	Top-1 Acc. (%) $\uparrow$
X3D-M [10]	-	16	$224 \times 224$	None	27.0
X3D-L [10]	-	16	$224 \times 224$	None	27.6
FAR [26]	X3D-M	16	$224 \times 224$	None	27.6
<b>Ours (MITFAS)</b>	X3D-M	16	$224 \times 224$	None	<b>40.2</b>
FAR [26]	X3D-M	8	$540 \times 540$	None	28.8
<b>Ours (MITFAS)</b>	X3D-M	8	$540 \times 540$	None	<b>38.4</b>
I3D [5]	ResNet-101	8	$540 \times 960$	Kinetics	21.1
FNet [29]	I3D	8	$540 \times 960$	Kinetics	24.3
FAR [26]	I3D	8	$540 \times 960$	Kinetics	29.2
FAR [26]	X3D-M	8	$620 \times 620$	Kinetics	39.1
<b>Ours (MITFAS)</b>	X3D-M	8	$620 \times 620$	Kinetics	<b>46.6</b>
X3D-M [10]	-	16	$224 \times 224$	Kinetics	30.6
MViT [9]	-	16	$224 \times 224$	Kinetics	24.3
FAR [26]	X3D-M	16	$224 \times 224$	Kinetics	31.9
<b>Ours (MITFAS)</b>	X3D-M	16	$224 \times 224$	Kinetics	<b>50.8</b>

Table 2. Benchmarking UAV Human and comparisons with prior arts.. For  $224 \times 224$  resolution and 16 frames input, when training from scratch, our approach achieves a 13.2% improvement over the baseline X3D-M and 12.6% over the current state-of-the-art FAR. For  $520 \times 520$  resolution and 8 frames input, MITFAS overperforms the current state-of-the-art FAR by 9.6% when training from scratch. For  $224 \times 224$  resolution and 16 frames input, when initializing with Kinetics pre-trained weights, MITFAS improves the top-1 accuracy over baseline by 20.2% and over SOTA method by 18.9%. For resolution over  $620 \times 620$  and 8 frames input, when initializing with Kinetics pretrained weights, MITFAS overperforms the current state-of-the-art FAR by 7.5%. Our method obtains better performance in all settings, which illustrates the effectiveness of our proposed MITFAS.

arms, we add 15% height as the margin on the top of  $F_r$  to ensure all the information about the action are included and crop the region as our final reference image. Therefore, we enlarge the region by 25% vertically and 10% horizontally to get  $F_r$ .

We use the sliding window strategy with scalable window sizes to find the aligned regions or features in all the frames. To make the process more efficient, we do not apply sliding window search over the entire frame. Instead, once we compute  $\omega_t^*$  at time  $t$ , we use the same operation at  $t+1$  to obtain the region  $L_{\omega_t^*}(F_{t+1})$ . We expand  $L_{\omega_t^*}(F_{t+1})$  by 25% as the searching area at  $t+1$ . In this way, we could significantly decrease the overall mutual information computations by only searching in the searching area which is a subset of  $F_{t+1}$ . In order to improve the reliability, we occasionally re-perform localization to update the searching area. More ablation studies on the impact of searching area size is given in the supplementary.

Once all the  $\omega_t^*$  are found for all time  $t$ , well-aligned frames are obtained by the transformation. We use our frame sampling method illustrated in Section. 3.2 to generate a sequence of 8 or 16 frames for model training. We will randomly pick a start frame as  $F_0$ , denoting the index of  $F_0$  in the sequence as  $k_0$ . To maintain the randomness in our sampling strategy, we set a randomly generated stride

$r_1$  when sampling  $F_1$ . We compute our candidate pool  $C_1$  by a set of all the frames that have index greater than  $k_0$ , but not exceed  $k_0 + r_1$ . Next, we find the most informative frame in the candidate pool using Eq. 11 and use it as  $F_1$ . We follow the same strategy to sample all the subsequent frames.

After obtaining all the sampled frames, we use a temporal inference backbone network to extract and learn from spatial-temporal features from the human actions. We employ X3D [10] as the backbone in our method for its efficiency and performance on video tasks. However, our method could be combined with any action recognition models for better behavior understandings on UAV videos.

## 4. Results

In this section, we describe our implementation and present the results. We compare the performance with other state-of-the-art video action recognition methods on 3 UAV datasets. The implementation and training details are shown in Appendix. A.

### 4.1. Results on UAV Human

UAV Human is currently the largest UAV-based human behavior understanding dataset. It contains scenarios captured from both indoor and outdoor environments with dif-

Method	Frames	Input Size	Init.	Top-1
HLPF	All	1920 × 1080	None	64.3
PCNN	-	1920 × 1080	None	75.9
X3D-M	16	224 × 224	Kinetics	83.4
FAR	16	224 × 224	Kinetics	92.7
<b>Ours</b>	16	224 × 224	Kinetics	<b>100.0</b>

Table 3. Results on Drone Action. Our method achieves 100% top-1 accuracy, 16.6% over the baseline method X3D-M [10], outperforming current state-of-the-art method FAR [26] by 7.3% under same configuration. (HLPF [22], PCNN [6])

Method	Frames	Input Size	Init.	Top-1
X3D-M	8	960 × 540	Kinetics	66.1
FAR	8	960 × 540	Kinetics	71.4
<b>Ours</b>	8	540 × 540	Kinetics	<b>78.6</b>

Table 4. Results on NEC Drones. Our method shows an improvement of 12.5% on top-1 accuracy against the baseline X3D-M [10], 7.2% over current state-of-the-art FAR [26].

ferent lighting and weather conditions. The videos are captured in dynamic backgrounds with different UAV motions and flying altitudes. It has 155 annotated actions, many of which are hard to distinguish such as squeeze and yawn.

We compare our method against prior state-of-the-art methods on UAV Human. As shown in Table 2, we implement our method and compare the performance with other state-of-the-art methods in various configurations in terms of the backbone network, frame rates, frame input sizes, and weights initialization. We use X3D-M as the backbone of our method with two different initialization settings. One of them is training from scratch and the other is initialized with Kinetics pretrained weights.

First, when using the same configuration (frames, input size, initialization), our method outperforms all the prior methods by a large margin. When training from scratch, we achieve a 12.6% improvement over current state-of-the-art methods. We get an 18.9% improvement when using Kinetics pretrained weights. This indicates the effectiveness of our method, which reduces the information redundancy and makes the model learn more from the motion changes rather than background variations.

## 4.2. Results on NEC Drone

NEC Drone is an indoor dataset that contains 5,250 videos with 16 actions performed by 19 actors. The videos are captured using a UAV flying at a low altitude on a basketball court. Compare to UAV Human, NEC Drone has more consistent lighting conditions while bringing more noises caused by light reflections.

We present the results on NEC Drone in Table 4. We obtain a Top-1 accuracy of 78.6%. We compare our method

Sampling Method	Top-1	Sampling Method	Top-1
Random	23.8	TFA + Random	39.8
Uniform	25.8	TFA + Uniform	42.2
MG Sampler	28.1	TFA + MG Sampler	45.5
MIS	<b>28.7</b>	TFA + MIS	<b>46.2</b>

Table 5. Temporal Feature Alignment (TFA) and Mutual Information Sampling (MIS) ablation studies on UAV-Human-Subset. The baseline is vanilla X3D with random [11] and uniform sampling [28], and we add our methods TFA and MIS step by step. From our experiments, TFA boost the accuracy by 16-17.5%. MIS outperforms the random sampling, uniform sampling, and MG Sampler [53].

Method	UAV-Human	Drone Action
Bounding box tracking [8]	47.4	95.9
Spatial-temporal action detection [31]	47.9	95.9
<b>TFA(ours)</b>	<b>50.8</b>	<b>100.0</b>

Table 6. Comparison with other methods [8,31].

against the baseline X3D-M and shows an improvement of 12.5%. Our approach outperforms the current SOTA FAR on NEC Drone by 7.2%. Note that, the improvement we achieved is obtained with 1/2 input frame size, which further demonstrates the advantage of our method.

## 4.3. Results on Drone Action

Drone Action is an outdoor video dataset that was captured using a free-flying UAV in low altitude and low speed. It contains 240 videos across 13 human actions performed by 10 human actors. Drone Action is the smallest dataset we used, but it is collected using a free-flying UAV that results in continuous position changes of the human actor.

As shown in Table 3, we achieve 100% Top-1 accuracy which outperforms current SOTA by 7.3% under the same configuration, which further illustrates the benefits of our proposed MITFAS.

## 4.4. Ablation Experiments

In this subsection, we mainly show the results of ablation experiments to demonstrate the effectiveness of the two components of our approach: Temporal Feature Alignment(TFA) and Mutual Information Sampling(MIS). More ablation studies are given in Appendix.

We randomly pick 30% videos for each action label in UAV-Human and conduct the ablation experiments on this UAV-Human subset. We use X3D-M [10] as the backbone network. All results are generated by using a sequence of 16 frames with a resolution of 224 × 224.

**Effectiveness of Temporal Feature Alignment** For Temporal Feature Alignment (TFA), our objective is to solve the small resolution corresponding to the human actor and viewpoint changes in the UAV videos. Our TFA finds and aligns the region that contains dominant information about the action for each frame in the video. As shown

Sampling Method	Alpha	Beta	UAV-Human Top-1	DroneAction Top-1
X3D + TFA + MIS	1.0	0.0	45.3	94.5
X3D + TFA + MIS	0.0	1.0	45.7	95.9
X3D + TFA + MIS	1.0	0.5	45.5	97.2
X3D + TFA + MIS	1.0	1.0	<b>46.2</b>	<b>100</b>

Table 7. Mutual Information Sampling (MIS) ablation studies on UAV-Human-subset and Drone Action. The baseline is vanilla X3D with TFA, we test the MITFAS Sampling in terms of two hyperparameters for mutual information and joint mutual information,  $\alpha$  and  $\beta$  respectively. From our experiments, MITFAS obtains the best accuracy when  $\alpha = 1.0$  and  $\beta = 1.0$ .

Similarity Measure	Top-1 Acc
Euclidean Distance	42.1
Cosine Similarity	39.5
Peak Signal-to-Noise Ratio	43.4
Structural Similarity Index Measure	44.8
<b>Mutual Information</b>	<b>46.2</b>

Table 8. Comparison with other similarity measures on UAV-Human Subset. Compared to other similarity measures, mutual information achieves the best accuracy.

in Table. 5, our TFA improves the top-1 accuracy by 16 - 17.5% when it is integrated with X3D and different sampling methods.

We also compare our TFA with other methods in Table 6. The bounding box tracking method [8] applies the person detector for foreground patch detection on all the temporal frames and then extracts the foreground patch based on the bounding boxes. Standard spatial-temporal action detection pipeline [31] integrates the detection and tracking algorithms to generate the proposal for the feature extraction. The results are generated using X3D and uniform sampling with the same configurations. As shown in Table. 6, our method improves the top-1 accuracy over other two methods by 2.9% on UAV-Human and 4.1% on Drone Action. Such improvement is attributed to our proposed TFA can not only extract the foreground patches but also align all the patches so that the main body of the human actor is well-matched in the temporal domain. Unlike the other two methods, our method does not align the bounding boxes. The bounding boxes are only used to locate the human actor. As illustrated in Section 3.1, the alignment in our method is performed at the pixel-level, ensuring that the generated frame closely resembles the preceding one. Therefore, the model could focus on the pixels corresponding to the parts of the human body that contribute most to the actions during training. Moreover, our method does not require any training procedures and could be utilized in any scenarios without domain issues.

#### Effectiveness of Mutual Information Sampling For

Mutual Information Sampling (MIS), our goal is to sample the informative frames that better represent the video for the action recognition methods. We compare it with three other sampling methods. First, we compare with two baseline methods: (1) Random sampling [11] where frames are randomly picked (2) Uniform sampling [28] where frames are sampled uniformly given a randomly generated start and end point. Then, we compare with the current state-of-the-art MG Sampler [53] which uses an adaptive sampling strategy based on temporal consistency between adjacent frames. As shown in Table. 5, compared with other sampling methods, MIS results in 0.6 - 6.4% improvement in Top-1 accuracy for UAV videos, which demonstrates the effectiveness of our proposed method.

**Hyperparameters** We evaluate our Mutual Information Sampling in terms of two hyperparameters for mutual information and joint mutual information,  $\alpha$  and  $\beta$  in Eq. 11 respectively. The baseline is vanilla X3D with TFA. As shown in Table 7, from our experiments, MITFAS obtains the best accuracy when  $\alpha = 1.0$  and  $\beta = 1.0$ . This demonstrates that both items in Eq.11 are equally important in discriminating the more informative frames.

**Comparison of other similarity measures** We compare the result of using mutual information with other similarity measures in Table.8. The results demonstrate that mutual information is a better criterion for measuring the similarity between images for UAV videos.

## 5. Conclusion, Limitations and Future Work

We propose a novel approach for video action recognition on UAVs. Our approach is designed to handle the varying and small resolution of the human, large changes in the positions of the human actor between frames, and partially occluded key points of the actions caused by continuous movement of the UAVs. We present a mutual information-based feature alignment to obtain and align the action features in the temporal domain. Our method is efficient and works well on UAV videos. We also present a novel frame sampling method to find the most informative frames in the video. We compare with prior approaches and demonstrate improvements in Top-1 accuracy on 3 UAV datasets. Our approach has a few limitations. First, we assume there does not exist a long-range spatial relationship between the human actor and the background. Second, we assume the input videos contain only one scripted human agent performing some action. We would like to explore the possibility of extending our method to multi-human or multi-action videos.

**Acknowledgement** This work was supported in part by ARO Grants W911NF2110026, W911NF2310046, W911NF2310352 and Army Cooperative Agreement W911NF2120076



## References

- [1] Oscar L Barbed, Pablo Azagra, Lucas Teixeira, Margarita Chli, Javier Civera, and Ana C Murillo. Fine-grained pointing recognition for natural drone guidance. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1040–1041, 2020. [3](#)
- [2] Mohammadamin Barekatin, Miquel Martí, Hsueh-Fu Shih, Samuel Murray, Kotaro Nakayama, Yutaka Matsuo, and Helmut Prendinger. Okutama-action: An aerial view video dataset for concurrent human action detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 28–35, 2017. [3](#)
- [3] Gavin Brown, Adam Pocock, Ming-Jie Zhao, and Mikel Luján. Conditional likelihood maximisation: A unifying framework for information theoretic feature selection. *Journal of Machine Learning Research*, 13(2):27–66, 2012. [5](#)
- [4] Kaidi Cao, Jingwei Ji, Zhangjie Cao, Chien-Yi Chang, and Juan Carlos Niebles. Few-shot video classification via temporal alignment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10618–10627, 2020. [2](#)
- [5] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6299–6308, 2017. [1](#), [3](#), [6](#)
- [6] Guilhem Chéron, Ivan Laptev, and Cordelia Schmid. P-cnn: Pose-based cnn features for action recognition. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3218–3226, 2015. [7](#)
- [7] Jinwoo Choi, Gaurav Sharma, Manmohan Chandraker, and Jia-Bin Huang. Unsupervised and semi-supervised domain adaptation for action recognition from drones. In *IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 1717–1726, 2020. [1](#), [2](#), [3](#)
- [8] Ugur Demir, Yogesh S Rawat, and Mubarak Shah. Tinyviral: Low-resolution video action recognition. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 7387–7394. IEEE, 2021. [3](#), [7](#), [8](#)
- [9] Haoqi Fan, Bo Xiong, Karttikeya Mangalam, Yanghao Li, Zhicheng Yan, Jitendra Malik, and Christoph Feichtenhofer. Multiscale vision transformers. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6804–6815, 2021. [6](#)
- [10] Christoph Feichtenhofer. X3d: Expanding architectures for efficient video recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 203–213, 2020. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [11](#), [12](#)
- [11] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981. [7](#), [8](#)
- [12] Harshala Gammulle, Simon Denman, Sridha Sridharan, and Clinton Fookes. Predicting the future: A jointly learnt model for action anticipation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5562–5571, 2019. [3](#)
- [13] Weihao Gao, Sreeram Kannan, Sewoong Oh, and Pramod Viswanath. Estimating mutual information for discrete-continuous mixtures. In *NIPS*, 2017. [5](#)
- [14] Ruben Geraldès, Artur Gonçalves, Tin Lai, Mathias Villerabel, Wenlong Deng, Ana Salta, Kotaro Nakayama, Yutaka Matsuo, and Helmut Prendinger. Uav-based situational awareness system using deep learning. *IEEE Access*, 7:122583–122594, 2019. [3](#)
- [15] Shreyank N. Gowda, Marcus Rohrbach, and Laura Sevilla-Lara. Smart frame selection for action recognition. In *AAAI*, 2021. [2](#)
- [16] Irtiza Hasan, Shengcai Liao, Jinpeng Li, Saad Ullah Akram, and Ling Shao. Generalizable pedestrian detection: The elephant in the room. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11328–11337, 2021. [5](#), [11](#)
- [17] Jiun Tian Hoe, Kam Woh Ng, Tianyu Zhang, Chee Seng Chan, Yi-Zhe Song, and Tao Xiang. One loss for all: Deep hashing with a single cosine similarity based learning objective. *Advances in Neural Information Processing Systems*, 34:24286–24298, 2021. [2](#)
- [18] Alain Hore and Djemel Ziou. Image quality metrics: Psnr vs. ssim. In *2010 20th international conference on pattern recognition*, pages 2366–2369. IEEE, 2010. [2](#)
- [19] Shihua Huang, Zhichao Lu, Ran Cheng, and Cheng He. Fapn: Feature-aligned pyramid network for dense image prediction. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 864–873, 2021. [2](#)
- [20] Zilong Huang, Yunchao Wei, Xinggang Wang, Wenyu Liu, Thomas S Huang, and Humphrey Shi. Alignseg: Feature-aligned segmentation networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(1):550–557, 2021. [2](#)
- [21] Mihir Jain, Hervé Jégou, and Patrick Bouthemy. Better exploiting motion for better action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2555–2562, 2013. [3](#)
- [22] Hueihan Jhuang, Juergen Gall, Silvia Zuffi, Cordelia Schmid, and Michael J Black. Towards understanding action recognition. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3192–3199, 2013. [7](#)
- [23] Xu Ji, João F. Henriques, and Andrea Vedaldi. Invariant information distillation for unsupervised image segmentation and clustering. *ArXiv*, abs/1807.06653, 2018. [2](#)
- [24] Yu-Gang Jiang, Qi Dai, Xiangyang Xue, Wei Liu, and Chong-Wah Ngo. Trajectory-based modeling of human actions with motion reference points. In *European Conference on Computer Vision*, pages 425–438. Springer, 2012. [3](#)
- [25] Stefan Klein, Marius Staring, and Josien P. W. Pluim. Evaluation of optimization methods for nonrigid medical image registration using mutual information and b-splines. *IEEE Transactions on Image Processing*, 16:2879–2890, 2007. [2](#)
- [26] Divya Kothandaraman, Tianrui Guan, Xijun Wang, Shuowen Hu, Ming Lin, and Dinesh Manocha. Far: Fourier aerial video recognition. In Shai Avidan, Gabriel Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner, editors, *Computer Vision – ECCV 2022*, pages 657–676. Cham, 2022. Springer Nature Switzerland. [1](#), [3](#), [6](#), [7](#), [11](#)

- [27] Divya Kothandaraman, Xijun Wang, Tianrui Guan, Sean Hu, Ming Lin, and Dinesh Manocha. Frequency-based aerial video recognition. In *Open Architecture/Open Business Model Net-Centric Systems and Defense Transformation 2023*, volume 12544, pages 127–129. SPIE, 2023. 3
- [28] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017. 7, 8
- [29] James Lee-Thorp, Joshua Ainslie, Ilya Eckstein, and Santiago Ontanon. Fnet: Mixing tokens with fourier transforms. *arXiv preprint arXiv:2105.03824*, 2021. 6
- [30] Tianjiao Li, Jun Liu, Wei Zhang, Yun Ni, Wenqian Wang, and Zhiheng Li. Uav-human: A large benchmark for human behavior understanding with unmanned aerial vehicles. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16266–16275, 2021. 1, 2, 3
- [31] Wenhe Liu, Guoliang Kang, Po-Yao Huang, Xiaojun Chang, Yijun Qian, Junwei Liang, Liangke Gui, Jing Wen, and Peng Chen. Argus: Efficient activity detection system for extended video analysis. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision Workshops*, pages 126–133, 2020. 7, 8
- [32] Zhenguang Liu, Runyang Feng, Haoming Chen, Shuang Wu, Yixing Gao, Yunjun Gao, and Xiang Wang. Temporal feature alignment and mutual information maximization for video-based human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11006–11016, 2022. 2
- [33] Hao Lu, Yutong Dai, Chunhua Shen, and Songcen Xu. Indices matter: Learning to index for deep image matting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3266–3275, 2019. 2
- [34] Frederik Maes, André M. F. Collignon, Dirk Vandermeulen, Guy Marchal, and Paul Suetens. Multimodality image registration by maximization of mutual information. *IEEE Transactions on Medical Imaging*, 16:187–198, 1997. 2
- [35] Yang Mi. *Moving-Camera Video Content Analysis Via Action Recognition and Homography Transformation*. PhD thesis, University of South Carolina, 2020. 3
- [36] Balmukund Mishra, Deepak Garg, Pratik Narang, and Vipul Mishra. Drone-surveillance for search and rescue in natural disaster. *Computer Communications*, 156:1–10, 2020. 3
- [37] Hazar Mliki, Fatma Bouhlel, and Mohamed Hammami. Human activity recognition from uav-captured video sequences. *Pattern Recognition (PR)*, 100:107140, 2020. 3
- [38] Mathew Monfort, Bolei Zhou, Sarah Adel Bargal, Alex Andonian, Tom Yan, Kandan Ramakrishnan, Lisa M. Brown, Quanfu Fan, Dan Gutfreund, Carl Vondrick, and Aude Oliva. Moments in time dataset: One million videos for event understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42:502–508, 2020. 1
- [39] Lichao Mou, Yuansheng Hua, Pu Jin, and Xiao Xiang Zhu. Event and activity recognition in aerial videos using deep neural networks and a new dataset. In *IGARSS 2020-2020 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, pages 952–955. IEEE, 2020. 3
- [40] Kien Nguyen, Clinton Fookes, Sridha Sridharan, Yingli Tian, Xiaoming Liu, Feng Liu, and Arun Ross. The state of aerial surveillance: A survey. *arXiv preprint arXiv:2201.03080*, 2022. 3
- [41] Asanka G Perera, Yee Wei Law, and Javaan Chahl. Drone-action: An outdoor recorded drone video dataset for action recognition. *Drones*, 3(4):82, 2019. 1, 2, 3
- [42] Asanka G Perera, Yee Wei Law, Titilayo T Ogunwa, and Javaan Chahl. A multiviewpoint outdoor dataset for human action recognition. *IEEE Transactions on Human-Machine Systems*, 50(5):405–413, 2020. 3
- [43] Josien P. W. Pluim, J. B. Antoine Maintz, and Max A. Viergever. Mutual-information-based registration of medical images: a survey. *IEEE Transactions on Medical Imaging*, 22:986–1004, 2003. 2
- [44] Jian Ren, Xiaohui Shen, Zhe L. Lin, and Radomír Mech. Best frame selection in a short video. *2020 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 3201–3210, 2020. 2
- [45] De Rosal Igantius Moses Setiadi. Psnr vs ssim: imperceptibility quality assessment for image steganography. *Multimedia Tools and Applications*, 80(6):8423–8444, 2021. 2
- [46] Waqas Sultani and Mubarak Shah. Human action recognition in drone videos using a few aerial training examples. *Computer Vision and Image Understanding*, 206:103186, 2021. 3
- [47] Paul A. Viola and William M. Wells. Alignment by maximization of mutual information. *International Journal of Computer Vision*, 24:137–154, 1995. 2, 4, 11
- [48] Xijun Wang, Ruiqi Xian, Tianrui Guan, Celso M de Melo, Stephen M Nogar, Aniket Bera, and Dinesh Manocha. Aztr: Aerial video action recognition with auto zoom and temporal reasoning. *arXiv preprint arXiv:2303.01589*, 2023. 3
- [49] Peter Washington, Aaron Kline, Onur Cezmi Mutlu, Emilie Leblanc, Cathy Hou, Nate Stockham, Kelley Paskov, Brianna Chrisman, and Dennis Wall. Activity recognition with moving cameras and few training examples: applications for detection of autism-related headbanging. In *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–7, 2021. 3
- [50] Zuxuan Wu, Caiming Xiong, Chih-Yao Ma, Richard Socher, and Larry S. Davis. Adaframe: Adaptive frame selection for fast video recognition. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1278–1287, 2019. 2
- [51] Shanshan Zhang, Rodrigo Benenson, and Bernt Schiele. Citypersons: A diverse dataset for pedestrian detection. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4457–4465, 2017. 11
- [52] Yuxiang Zhao, Khurram Shafique, Zeeshan Rasheed, and Maoxu Li. Janusnet: Detection of moving objects from uav platforms. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3899–3908, 2021. 1
- [53] Yuan Zhi, Zhan Tong, Limin Wang, and Gangshan Wu. Mgsampler: An explainable sampling strategy for video action recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1513–1522, 2021. 2, 7, 8