

FAR: Fourier Aerial Video Recognition

Divya Kothandaraman¹[0000–0002–6276–4968], Tianrui
Guan¹[0000–0002–6892–9778], Xijun Wang¹, Shuowen Hu², Ming
Lin¹[0000–0003–3736–6949], and Dinesh Manocha¹[0000–0001–7047–9801]

¹ University of Maryland College Park, United States

² DEVCOM Army Research Laboratory, United States

dkr@umd.edu

<https://gamma.umd.edu/far>

Abstract. We present an algorithm, Fourier Activity Recognition (FAR), for UAV video activity recognition. Our formulation uses a novel Fourier object disentanglement method to innately separate out the human agent (which is typically small) from the background. Our disentanglement technique operates in the frequency domain to characterize the extent of temporal change of spatial pixels, and exploits convolution-multiplication properties of Fourier transform to map this representation to the corresponding object-background entangled features obtained from the network. To encapsulate contextual information and long-range space-time dependencies, we present a novel Fourier Attention algorithm, which emulates the benefits of self-attention by modeling the weighted outer product in the frequency domain. Our Fourier attention formulation uses much fewer computations than self-attention. We have evaluated our approach on multiple UAV datasets including UAV Human RGB, UAV Human Night, Drone Action, and NEC Drone. We demonstrate a relative improvement of **8.02% - 38.69%** in top-1 accuracy and up to **3** times faster over prior works.

1 Introduction

Deep learning techniques have been widely used for activity recognition [21, 7, 5]. Video analysis of scenes captured using UAV cameras [43, 52] is much harder than activity recognition in ground-camera datasets [7, 63]. In these UAV videos, the object of interest, *i.e.* the human actor (any individual appearing in the video performing scripted or non-scripted actions), is typically much smaller in terms of number of pixels or the area than the corresponding background, and thus provides less knowledge than a front view capture. Moreover, it is harder to capture and label UAV videos. Overall, there are fewer and smaller labeled datasets of aerial videos, as compared to ground videos. For instance, ground-camera datasets like Kinetics-400 [7] contain 306,245 videos while the recent UAV-Human [43] database has 22,476 videos.

Given that the size of the human actor in UAV videos is much smaller than the corresponding background, a neural network trained on these datasets may

² The second and third authors contributed equally

learn to infer more from the background [41] than the human actor. While both background and context are important [12], the network must learn to first identify the human actor and the corresponding action, and then deduce relations of the human actor with the surroundings in a judicious manner. In the absence of annotated detection boxes that can demarcate the human actor, the network needs to be able to differentiate the moving human actor from the background in an intrinsic manner. One approach is to detect the object of interest via object detection [57]. However, action recognition models that heavily rely on localization of the human actor require near to perfect object detection accuracy [82]. While it is practically not feasible to annotate all datasets for object detection, object detectors trained on ground camera datasets will not generalize well UAV videos due to domain gap issues [70,8,4]. Domain adaptation solutions do not lead to perfect generalization yet.

On the other hand, traditional optical flow [3] techniques require hundreds of optimization iterations each frame, and split the network into RGB and motion streams which increases computation and model parameters [54]. Low computation alternatives such as deep learning based optical flow [58,33,16], motion feature networks [39] and ActionFlowNet [50] are inferior in performance compared to optical flow. Techniques such as background subtraction [53] and motion segmentation [77] are not very promising either [60,19]. Thus, the network needs to learn to automatically *disentangle* [81,18] object feature representations from the corresponding entangled state containing both the object and background information.

In addition to object-background separation, it is important for the network to acquire knowledge [5] about the context, relationships between the object, and background and intra-pixels correspondences as well. Self-attention [69,78] can model this information by capturing long-range dependencies within an image/video. Prior work on attention based video activity recognition [5,44,1] has seen two classes of self attention networks by either directly applying self-attention on convolutional layers or using self-attention as the building block. Mathematically, the core step in the computation of self-attention is matrix multiplication, which makes it computationally expensive.

1.1 Main contributions

We present a novel method, FAR, for UAV video action recognition. The design of FAR in the frequency domain is motivated by the fact that frequency spectrums contain knowledge about a signals' characteristics that are not easily interpretable in the time domain. Our novel components include:

- We propose a novel Fourier Object Disentanglement method (FO) to bestow the network with the ability to *intrinsically* recognize the moving human actor from the background. FO operates in the frequency domain dictated by the spectrum of the Fourier transform corresponding to the temporal dimensions of the video. It characterizes the motion of the human actor based on the magnitude and rate of temporal change of feature maps that encode information about the spatial pixels of the video. The amplitudes at

each spatial-temporal location of the feature maps are innately representative of dynamic salient, static salient, dynamic non-salient and static non-salient regions, in the same order of relevance. This also empowers the network to handle videos with moving background pixels and dynamic cameras.

- We present Fourier Attention (FA) to encapsulate context and long range space-time dependencies within a video. Fourier attention works in the frequency domain corresponding to the space-time dimensions of the video, and emulates the benefits of self-attention. The time complexity of FA is $O(n^2 \log n)$ as opposed to $O(n^3)$ for traditional self-attention, and the accuracy of Fourier-attention approximates that of self-attention.

Moreover, such a representation promotes global mixing. FAR has multiple benefits. (i) It elegantly exploits the mathematical properties of Fourier transform to achieve the desired objectives of object background separation and context encoding by performing fewer computations than traditional methods. (ii) It is parameter-less, i.e., it does not have any learnable layers/ parameters. (iii) FAR can be embedded within any 3D action recognition network such as I3D [7,21] to achieve state-of-the-art performance. (iv) FAR converges faster than the corresponding 3D action recognition backbone.

We experimentally demonstrate that FAR outperforms prior work by 8.02%–38.69% performance across multiple UAV datasets including UAV Human RGB [43], UAV Human Night [43], Drone Action [52], and NEC Drone [13]. We compare with the state-of-the-art Fourier method, efficient attention method and self-attention based transformer methods and demonstrate accuracy, computation and memory benefits.

2 Related Work

Action Recognition: Action recognition is a well studied topic in computer vision. The emergence of large-scale ground-camera videos datasets [7,63,49] has led to development of deep learning techniques for action recognition. We refer the reader to [11] for a survey on action recognition. Broadly speaking, three classes of network architectures have been proposed for action recognition. The first [64,23,26,30,73] builds on the two-stream theory in cognition to model space and time separately. The second [22,21,7,5,67,32] models space and time jointly via 3D CNNs. The third class includes transformer-based architectures [55,27,46,5,72]. These transformer-based solutions are built on self-attention [69,78] and have high computational complexity. In the interest of optimizing GPU memory, frame sampling strategies [28,80,37,29] for video action recognition have been proposed. The above solutions are focused on challenges pertaining to action recognition in ground-camera videos. However, UAV video action recognition is much more difficult.

UAV Action Recognition: UAV video databases [2,43,13,52] have been used to develop solutions [65,51,15,68] for UAV action recognition. However, these solutions are directly based off techniques designed for ground-camera datasets [7,49], where the size of the object is comparable to the background. Moreover, for

ground camera videos, an auxilliary guidance factor based on object detection [57] is a viable option. However, these assumptions do not hold true in UAV videos [17,48].

FFT and Deep Learning FFT has been immensely used in traditional image [6,56] and video processing [79,14] applications. Fast Fourier Transform (FFT) has been recently used in deep learning methods. One of its first applications was to accelerate convolution operations [38]. Incorporating FFT between NN layers [10,75,66] instead of CNNs to transform the feature space to the frequency domain, and aid global mixing of knowledge, has been used to improve accuracy for image classification, detection and ground-camera action recognition. An interesting application of FFT includes image stylization [76] as a guiding factor for domain adaptation. Most recently, FFT was used to naively replace self-attention layers [69] for NLP applications.

Efficient attention Methods to improve memory efficiency of transformers include modifications in matrix multiplication [61], low rank approximations [71], kernel modifications [35] for linear time complexity [74,59,36,42]. While most of these solutions are focused on NLP and image-based computer vision tasks, EA [61] demonstrates results on temporal action localization and STAR [62] performs skeleton action recognition. The former can be regarded as a localization task w.r.t. the temporal dimension while the latter uses pose information making the task of classification easier. None of these solutions are customized to UAV action recognition which brings forth different challenges.

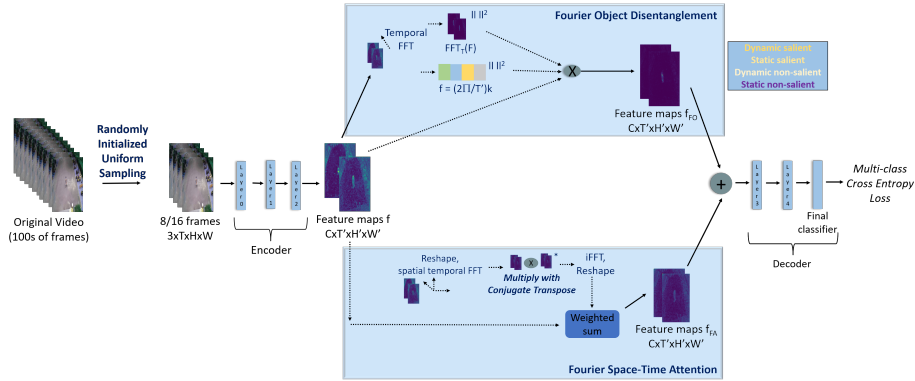


Fig. 1: Fourier Object Disentanglement (FO) and Space-Time Fourier Attention (FA): FO empowers the network to intrinsically separate out the moving human agent from the background, without the need for any annotated object detection bounding boxes. This enables our network to explicitly focus on the low resolution human agent performing action, and not just learn from background cues. FO inherently characterizes salient and non-salient, and static and dynamic regions of the scene via the amplitudes of the feature maps it computes. FA elegantly exploits the mathematical properties of the Fourier transform to imbibe the properties of self-attention and capture contextual knowledge and long-range space-time dependencies at a much lower computational complexity.

3 Fourier Disentangled Space Time Attention

In this section, we describe our approach. We design two novel methods to decipher the human actor performing action, and encode context. Fourier Object Disentanglement (FO) disentangles the object from the background in an automatic manner. Fourier Space-Time Attention (FA) imbibes the properties of self-attention to capture long range space-time relationships at a lower computational cost. These modules can be embedded within any state-of-the-art 3D video recognition backbone such as I3D [7] or X3D [21] for improved action recognition. We now describe the methods in detail.

3.1 Fourier Object Disentanglement

We present a Fourier Object Disentanglement (FO) method to automatically separate the human actor from the background. Movement of the human actor in the scene can be characterised by temporal change of feature maps encoding spatial pixels (across space dimensions $H \times W$) in the video frames. The rate and magnitude of change of a signal can be quantified by amplitude of a signal at different frequencies. Thus, to identify the movement, we first transform the feature maps to a temporal frequency space. We perform this computation using 1D Fourier transform along the temporal dimension. Specifically, let $f(c, t, h, w) \in C \times T' \times H' \times W'$ denote the feature maps on which FO is applied, where C is the number of channels and T' and $(H' \times W')$ denote the temporal and spatial dimensions of the feature maps, respectively. The amplitude of the temporal Fourier transform at the frequency $-2\pi k/N$ is:

$$\mathcal{F}_T(f)(k) = \sum_{n=0}^{n=T'} f(c, t, h, w) \times e^{-2\pi kn/N}, \quad (1)$$

which can be computed efficiently using the FFT algorithm [24]. $\mathcal{F}_T(f)(k)$ mathematically represents the amplitude of the temporal signal at every spatial and channel location of the feature map f , at various frequencies. Intuitively, high frequency in the temporal dimension corresponds to the movement, and low frequency represents static regions of the scene. Therefore, regions corresponding to the moving human actor should have higher amplitude of Fourier transform at high frequencies. To infer the presence of the moving human actor at various spatial locations, we encapsulate the relationships between amplitudes and frequencies by multiplying the L2-norm-square of the amplitude at each frequency with the L2-norm-square of the frequency itself. L2-norm ensures that frequencies and amplitudes are positive. L2-norm-square amplifies high amplitudes of the Fourier transform of the signal at high frequencies and suppress low amplitudes at low frequencies for disentangling dynamic regions of the scene. The frequencies, in order, are: $fr_k = [e^{-2\pi k/N}]$, $k = 1 \dots T'$.

Note that the frequencies are independent of the input video. Thus, the dynamic mask \mathcal{M}_{FO} can be represented as

$$M_{FO} = \|\mathcal{F}_T(f)(k)\|_2^2 \times \|fr_k\|_2^2, \quad (2)$$

where $|a|_2^2$ is L2 norm-square of a vector $|a|$. M_{FO} disentangles (or amplifies) parts of the scene corresponding to moving pixels. This may include moving background (and camera motion) in addition to moving human actor. Our next task is to use M_{FO} to demarcate moving object pixels from **moving background pixels**.

To further separate out only the moving actor, we capitalize on the activation maps f computed by the model. While not perfect, the activations at salient regions of the scene are higher than those at the non-salient regions. Hence, the final object disentangled representations can be represented as a dot product of M_{FO} and network features f , which amplifies dynamic, salient regions of the scene. Mathematically,

$$F_{FO} = f \odot M_{FO}. \quad (3)$$

According to this formulation, dynamic salient regions are amplified the most, and static non-salient regions are heavily suppressed. The amplitude at static salient regions and dynamic non-salient regions is lower than the amplitude at dynamic salient regions. Due to the l_2 operation in the computation of M_{FO} and linear application of f in Equation 3, static salient regions have a higher amplitude than the dynamic non-salient regions. Thus, the ordering of amplitudes that is formed as: dynamic-salient > static-salient > dynamic-non-salient > static-non-salient, in concordance with the relevance for decision making for action recognition. Thus, static as well as dynamic background regions have lower amplitudes than static and dynamic regions of the object executing action.

Time complexity: The time complexity of FO depends on the time complexity of 1D FFT, which is $n \log(n)$, for an n -element input vector. Consider the classical case [31] where the temporal and spatial dimensions at the mid level feature representations is half and one-fourth of the number of frames sampled and spatial dimensions of the image respectively. The number of FFTs that need to be computed is $C \times (H/8) \times (W/8)$ where C, H, W correspond to the number of channels at the mid-level, and spatial dimensions of the image. Therefore, the total time complexity is $C \times (H/8) \times (W/8) \times (T/2) \log(T/2)$.

3.2 Space-Time Fourier Attention

Consider a scene that depicts a human actor swimming in a swimming pool. Here, it is important to decipher the relationship between the human actor and the pool. While explicit modeling of correspondences between different pixels illustrating pose, orientation, and joint movements may not be necessary, it is crucial for the neural network to inherently capture this knowledge. Space-time self-attention for video action recognition [5,44,1] is capable for extracting this knowledge, but comes at the cost for expensive matrix multiplications.

We propose Fourier Space-Time Attention (FA) for acquiring knowledge about the long-range space-time relations within a video. Fourier attention approximates self-attention in an elegant fashion at a reduced computational cost. To understand the mechanics of Fourier attention, we first succinctly present self-attention [78]. The inputs to self-attention are key, query and value vectors which are representations obtained by 1×1 convolutions using a common input

feature map. Vaswani et al. [69] describe the computation of self-attention as “a weighted sum of the values, where the weight (or sub-attention) assigned to each value is computed by a compatibility function of the query with the corresponding key.” Key, query and value are 1×1 convolution layers transforming the input feature maps. Mathematically, with \mathbf{x} representing the input feature maps, and \odot denoting matrix multiplication,

$$\text{Attention} = \text{Value}(\mathbf{x}) \odot [\text{Query}(\mathbf{x})^T \odot \text{Key}(\mathbf{x})]^T \quad (4)$$

Our space-time Fourier attention method proceeds as follows. The first step is to obtain a representation equivalent to the key-query computation, termed *Fourier sub-attention*. Fourier sub-attention is motivated by autocorrelation, which is the correlation coefficient between different parts of the same signal. We define *Fourier sub-attention* as the element-wise product of the Fourier transform of feature maps with the conjugate transpose of the Fourier transform of these feature maps (Equation 6). To compute this space-time Fourier sub-attention, we reshape the video feature maps f to a 3D representation $C \times T' \times (HW)$, which are transformed to the frequency domain via 2D Fourier transform along the space and time axes as follows:

$$\mathcal{F}_{ST}(f)(m, n) = \sum_{h, w} f(c, t, h, w) e^{-2\pi m h / M} e^{-2\pi n w / N}. \quad (5)$$

computed efficiently using the FFT algorithm [24]. FFT is a representation of the signal as a whole at a wide spectrum of frequencies, and enables inherent and exhaustive global mixing between various spatial and temporal regions of the video. The space-time Fourier sub-attention \mathcal{A}_{ST} in the Fourier domain is simply the element wise multiplication between \mathcal{F}_{ST} and its complex conjugate \mathcal{F}_{ST}^* :

$$\mathcal{A}_{ST} = \mathcal{F}_{ST} \times \mathcal{F}_{ST}^* \quad (6)$$

Next, we compute the inverse FFT (\mathcal{IF}) of \mathcal{A}_{ST} to obtain the correlations in the time domain, and reshape to $C \times T' \times H' \times W'$. These sub-attention “weights” are then used in a dot product (or element wise multiplication) with the input feature maps f to compute the final space-time Fourier attention maps f_{FA} . A scaling factor λ_{FA} , chosen empirically to be 0.01, scales these Fourier attention maps, which are then sum-fused with the input feature maps. Mathematically,

$$f_{FA} = F + \lambda_{FA} \times \mathcal{IF}(\mathcal{A}_{ST}), \quad (7)$$

Time complexity: Traditional self-attention [69] requires the model to perform two matrix multiplications. In the first matrix multiplication of self attention, we multiply the query matrix ($THW \times C$) with the key matrix ($C \times THW$). The time complexity is $O(C \times THW \times THW)$. In the second matrix multiplication, we multiply the value matrix ($C \times HWT$) with the attention matrix ($HWT \times HWT$). The complexity of this stage is $O(C \times HWT \times HWT)$. Hence, the overall time complexity of space-time self-attention [5] is $O(HWT \times HWT \times C)$.

In contrast, our Fourier attention solves the problem via one 2D FFT and one 2D iFFT. 2D FFT is computed on a matrix of dimensions $HW \times T$. The number of 2D FFTs that need to be computed is equal to the number of channels (C). Hence, the complexity is $O(C \times HWT \log(HWT))$. The complexity of 2D FFT and 2D iFFT are the same. Therefore, the overall time complexity of Fourier attention is $O(C \times HWT \log(HWT))$. Clearly, Fourier attention is much more efficient than self attention. In terms of accuracy, space-time Fourier attention is comparable to space-time self-attention [5].

3.3 Mathematical Analysis

Lemma 1. *Given an input matrix A , Fourier attention as well self-attention [69, 5] encapsulate long-range relationships for global mixing by computing outer products.*

Proof: We refer the reader to the supplementary material for the detailed proof. We present a concise version here. Without loss of generality, let $[a_{ij}]$ denote the elements of a square matrix A (with dimensions N) in 2D. f, g, h represent 1×1 convolutions for key, query, value computations in self-attention. The self-attention matrix S_{mn} is:

$$S_{mn} = \sum_{l=1}^N h a_{ml} \sum_{k=1}^N [g a_{lk} \times f a_{kn}] \quad (8)$$

Fourier attention F_{mn} is:

$$F_{mn} = \sum_{b=1}^N \sum_{c=1}^N \overbrace{\exp(-2\pi mc/N) \exp(-2\pi nb/N) a_{mn}}^{h_{mn}(b,c)} \times \left\{ \underbrace{\sum_{j=1}^N \sum_{i=1}^N \exp(-2\pi j(b-c)/N) a_{ij}}_{f_{mn}(b,c)} \times \underbrace{\exp(-2\pi i(c-b)/N) a_{ij}}_{g_{mn}(b,c)} \right\} \quad (9)$$

f, g, h in Equation 8 are 1×1 convolutions, and that the exponential terms span the entire spectrum of frequencies lets us define f, g, h for Fourier attention as shown in Equation 9. Thus, the equation for Fourier attention can be simplified as:

$$F_{mn} = \sum_{b=1}^N \sum_{c=1}^N h_{mn}(b, c) a_{mn} \times \left\{ \sum_{j=1}^N \sum_{i=1}^N f_{mn}(b, c) a_{ij} \times g_{mn}(b, c) a_{ij} \right\}$$

In self-attention, f, g, h are learnable. In contrast, in Fourier attention, f, g, h are pre-defined by the Fourier spectrum. Nonetheless, they exhaustively cover the Fourier spectrum. Moreover, the terms involved and the structure of computations (multiplications followed by summation) in Equations 8 and 10 are similar, both promote global mixing and encapsulate long-range relationships.

4 FAR: Activity Recognition in UAVs

We present FAR, a network for video action recognition in UAVs (Figure 1). FAR samples 8-16 frames from the input video by using randomly initialized uniform sampling, described in Section 4.1. These frames are passed through the first few layers of the 3D backbone network (or encoder) to generate feature maps f . These features contain entangled object and background information along the space-time dimensions. The choice of this intermediate layer in the backbone network that extracts feature maps f is a careful trade-off between the spatial-temporal resolutions needed for FAR to work well and the amount of knowledge contained in the networks' layers. We describe this choice in detail in this section, as well as present ablation experiments in 5.3 to justify our choice.

The *Fourier Object Disentanglement* module (Section 3.1), and the *Fourier Space-Time Attention* module (Section 3.2) act on f , in parallel, to generate f_{FO} and f_{FA} , respectively. f_{FO} and f_{FA} are sum fused and passed through the remaining layers of the neural network to generate the final action classification probability distribution, used in a multi-class cross entropy loss term with the ground-truth label for back-propagation.

Incorporating FO within the 3D backbone: Typically, to encapsulate temporal movement at each spatial location, we need to ensure that the spatial temporal dimensions of the feature map is not too small. Thus, it is useful to perform this operation using mid-level features (output from the middle layer of the network, as shown in Figure 1) that strike a fine balance between generic features that capture context, and focused high level features (at output layer).

Incorporating FA within the 3D backbone: After FO, the network does not contain any background signal. Hence, Fourier attention needs to be applied either before FO or in parallel with FO. FO is applied on mid-level features. Applying FA at a high level is not very effective because the extracted features do not have sufficient information. Hence we apply FA on the mid-level features as well, in parallel with the Fourier object disentanglement module.

4.1 Randomly Initialized Uniform Sampling

It is computationally expensive to use all the frames in a video. In traditional uniform sampling, T frames are sampled at uniform intervals. The standard way of uniform sampling under-utilizes [80,37] the knowledge that can be gained from the original video, which adds to the pre-existing issue of limited data. We use a variation of uniform sampling to improve the variance of the network and hence boost accuracy. First, we compute the step size as the ratio of total number of frames in the video and number of frames that we desire to sample. Next, we generate a random number between 0 and step size, and correspondingly designate the first frame to be sampled. This is followed by uniformly sampling video frames at step size intervals from the designated first frame.

5 Experiments and Results

We will make all code and trained models publicly available.

5.1 Datasets

In this section, we briefly describe the UAV datasets used for evaluating FAR. **UAV Human RGB** [43] is the largest UAV-based human behavior understanding dataset. Split 1 contains 15172 and 5556 images for training and testing respectively captured under various adversities including illumination, time of day, weathers, etc. **UAV Human Night Camera** [43] contains videos similar to UAV Human RGB captured using a night-vision camera. The night vision camera captures videos in color mode in the daytime, and grey-scale mode in the nighttime. **Drone Action** [52] is an outdoor drone video dataset captured using a free flying drone. It has 240 HD RGB videos across 13 human actions. **NEC Drone** [13] is an indoor UAV video dataset with 16 human actions, performed by human subjects in an unconstrained manner.

5.2 Implementation Details

Backbone network architecture: We benchmark our models using two state-of-the-art video recognition backbone architectures (i) I3D [7] (CVPR 2017) (ii) X3D-M [21] (CVPR 2020). For both X3D and I3D, we extract mid-level features after the second layer.

Training details: Our models were trained using NVIDIA GeForce 1080 Ti GPUs, and NVIDIA RTX A5000 GPUs. Initial learning rates were $\{0.01, \text{ and } 0.001\}$ across datasets. We use the Stochastic Gradient Descent (SGD) optimizer with weight decay of 0.0005 and momentum of 0.9, and cosine/ poly annealing for learning rate decay. The final softmax predictions of all our models were constrained using multi-class cross entropy loss.

Evaluation: We report top-1 and top-5 accuracies.

Table 1: **Results on UAV Human RGB. Table (a):** FAR can be embedded within any 3D action recognition backbone to achieve state-of-the-art performance. Pretraining with Kinetics boosts performance, and large input sizes work better since FA and FO are designed to capture global, as well as local knowledge. FAR imparts improvements of 2.20%-38.69% over 3D action recognition backbones across training configurations. **Table (b) - Ablation experiments:** We demonstrate that each component of FAR imparts substantial improvement in top-1 accuracy by upto **8%**.

Backbone	FAR	Frames	Input Size	Init.	Top-1	Top-5
(i) I3D	✗	8	540×960	Kinetics	21.06	40.81
(ii) I3D	✓	8	540×960	Kinetics	29.21	50.27
(iii) X3D-M	✗	16	224×224	None	27.0	44.2
(iv) X3D-M	✓	16	224×224	None	27.6	44.1
(v) X3D-M	✗	16	224×224	Kinetics	30.6	50.3
(vi) X3D-M	✓	16	224×224	Kinetics	31.9	50.3
(vii) X3D-M	✗	8	540×540	Kinetics	36.6	57.1
(viii) X3D-M	✓	8	540×540	Kinetics	38.6	59.2

(a)

FO	FA	Sampling	Top-1
✗	✗	✗	21.06
✓	✗	✗	25.89
✗	✓	✗	24.15
✓	✓	✗	27.00
✓	✓	✓	29.21

(b)

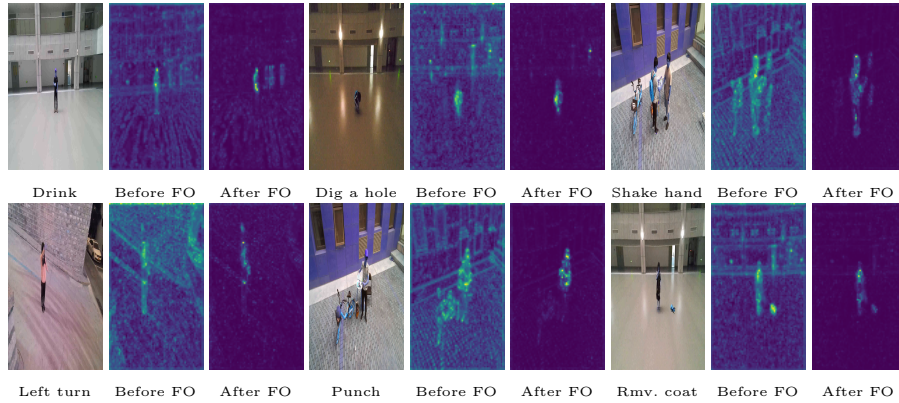


Fig. 2: **Qualitative results on UAV Human RGB.** We show the effect of our Fourier Object Disentanglement (FO) method. In each sample, the images, in order, correspond to a frame from the video, feature representation before disentanglement and the feature representation after disentanglement respectively. Notice the effectiveness of FO in scenes with light noise (Row 1 Image 2, Row 2 Image 3), dim light (Row 1 Image 2), dynamic camera and dynamic background (Row 1 Image 1). Regions of the scene corresponding to moving human actor (or salient dynamic) are amplified most (solid yellow). Static background is completely suppressed (solid purple). Static salient regions are slightly amplified (e.g. lower body of human actor in Row 2 Image 3 - yellow), and dynamic backgrounds are suppressed to a great extent (pale yellow in Row 1 Image 1). We show more results in the supplementary material.

5.3 Main Results: UAV Human RGB

Benchmarking FAR FAR can be embedded within any 3D action recognition backbone to achieve state-of-the-art performance. In Table 1a, we show results on UAVHuman RGB at different frame rates, input sizes, backbone network architectures and pre-trained weights initialization. In experiment (i) and (ii), we use the I3D backbone, and initialize the network with pretrained Kinetics weights. Spatially, we downsample the input video by a factor of 2, and sample 8 frames per video. This configuration gives the network full access to the spatial portions of the video at every stage of training and testing. FAR imparts a relative improvement of **38.69%** and **23.18%** in top-1 and top-5 accuracy, respectively.

In the subsequent experiments, we use X3D-M, as the backbone. Many vision-based papers crop the original video into small patches of resolution 224×224 . We explore this in experiments (iii)-(vi) under two settings: without initializing with Kinetics pretrained weights, and by initializing with Kinetics pretrained weights. Concurrent with our intuition, initializing with Kinetics pretrained weights results in better performance than without initializing with Kinetics pretrained weights. In both cases, with small crop size, FAR improves performance over the corresponding baselines by 2.2% – 4.24%. At a resolution of 224×224 , there is a slight decrease in Top-5 accuracy (0.1%).

Video action recognition is a global level task. Hence, it is important for the network to see a larger spatial region of the video to understand context, and get a better view of the human actor. Moreover, since the design of FAR




Fig. 3: FAR converges much faster than the state-of-the-art action recognition method X3D-M [21]. In the left curve, we show the top-1 train accuracy as a function of the networks’ training iterations. In the right figure, we plot the training loss curve. We demonstrate that FAR imparts convergence benefits over prior work, under the same hyperparameter and GPU configurations.

is specifically inspired of challenges pertaining to object background separation, and context encoding, the margin of improvement imparted by FAR to the backbone architecture is larger when the crop size is higher. At a crop size of 540×540 , FAR improves top-1 and top-5 accuracies by 5.46% and 3.67% respectively, over the corresponding baselines.

Table 2: **Comparisons with state-of-the-art self-attention based transformer methods on UAV Human.** We initialize all our models with Kinetics pre-trained weights. We observe higher accuracy and computational benefits (up to **3x**) with FAR.

Method	Param M	Top-1 %	FLOPs GFlops/video	GPU GB/video	Inference Time (sec)/video
I. Baseline non-attention models					
X3D-M	3.8	36.6	14.39	3	0.08
I3D-M	12	21.06	346.55	10	0.1
II. Comparisons with the state-of-the-art Fourier-based method					
I3D + FNet [40] (2021)	12	24.39	346.56	10	0.1
I3D + FAR (Ours)	12	29.21	346.6	10	0.2
III. Comparison with the state-of-the-art efficient-attention method					
I3D + Efficient Attention [61] (WACV 2021)	12	21.13	462	13.3	0.12
I3D + Fourier Attention (Ours)	12	24.15	346.57	10	0.19
IV. Comparisons with the state-of-the-art self-attention based transformer methods					
ViT-B-TimeFormer (ICML 2021) [5]	121.4	33.9	2380	32	0.27
MViT (ICCV 2021) [20]	36.6	24.3	70.8	9	0.16
X3D-M + FAR (Ours)	3.8	38.6	14.41	3	0.09

Ablation Experiments FAR ablations. We present ablation experiments on the components of FAR in Table 1b. We use the I3D backbone [7], and sample 8 frames per video. We initialize with Kinetics pretrained weights, and spatially downsample (and then feed in the entire frame) the video by a factor of 2. In the first four experiments, we uniformly sample 8 frames from frame 0. The first row is the baseline experiment with neither Fourier object disentanglement nor Fourier attention. We observe in the experiment corresponding to Row 2 that object disentanglement improves performance by 22.9% over the baseline. FO is a high pass filter (L2). The usage of a linear (L1) high pass filter results in an accuracy of 25.56%. Thus, we used the L2 high-pass filter as it results in a higher accuracy. Next, we determine the importance of context and long-range

space-time relationships by using only Fourier attention, and demonstrate an improvement of 14.67% in Row 3.

FO and FA complement each other - the former disentangles object from the background, while the latter decodes contextual information and inter-pixel, inter-frame relations. Using FO and FA in parallel, and sum-fusing the resultant feature maps cumulatively improves performance by 28.2% over the baseline. Finally, we incorporate the sampling scheme, vis-a-vis, randomly initialized uniform sampling along with FO and FA. This results in a final accuracy of 29.21%, 38.69% over the baseline.

FA ablations. We conduct ablation experiments on our proposed Fourier Attention (FA). We use the I3D backbone, and a video resolution of $540 \times 960 \times 8$. In all these experiments, FO is applied on level 2. In the first experiment, we extend FA to channels [25] in addition to space-time, at level 2, the accuracy is 26.48. In the second experiment, we apply channel FA at level 4 [25] while retaining space-time FA at level 2, the accuracy further degrades to 25.77. In contrast, the accuracy with space-time Fourier attention is 27.00. Thus, we find that global mixing at channel level does not contribute to improvement in performance.

Next, we explore the notion of multi-level FA, where FA is applied at multiple levels and not just level 2, the accuracy is 29.16. In contrast, FAR’s accuracy is 29.21. Our conclusion is that FA extracts knowledge prerequisite to learning long-range space-time relationships at level 2, to its maximum capacity, and applying it at more layers is redundant.

State-of-the-art comparisons We report state-of-the-art comparisons in Table 2. For all our experiments, we set the temporal and spatial resolutions at 8 frames and upto 540 (short side) respectively. We establish the baseline accuracies using non-attention networks vis-a-vis, I3D [7] and X3D [21], in experiment I.

Comparisons with FNet. In Table 2 experiment II, we report the accuracy using FNet, which is the state-of-the-art Fourier transform based self-attention method. FNet naively replaces every self-attention layer with the Fourier transform of the feature maps at that level. The motivation is to "mix" different parts of the feature representation and thus gain global information. Originally designed for NLP, it achieves 92 – 97% of the accuracy of BERT counterparts on the GLUE counterparts. However, when applied to video activity recognition on UAV Human RGB with the I3D backbone, we find that the accuracy is just 24.39%. In contrast, with the same backbone and hyperparameter settings, we demonstrate that FAR achieves 29.21%, an improvement of 19.76%.

Comparisons with efficient attention methods. We compare with the current state-of-the-art efficient attention method [61] in experiment III. For fair comparisons, we use the I3D backbone in both cases, at a video resolution of $540 \times 960 \times 8$. We demonstrate better accuracies with our Fourier Attention formulation at lower FLOPs and GPU memory.

Comparisons with transformer/self-attention methods. We compare against self-attention based transformer methods in Table 2. Specifically, we compare against (i) TimesFormer [5] (ICML 2021) - a self-attention video recognition

method based on joint space-time self attention, and (ii) MViT (ICCV 2021) - a transformer based method that combines multi-scale feature hierarchies. We demonstrate much better performance at lower number of FLOPs, GPU memory and inference time. Another benefit is that FAR does not add any new parameters to the neural network and uses the same number of parameters as its backbone network. In contrast, MVIT and TimeSformer use much higher number of parameters.

Table 3: **Results on more UAV datasets.** We demonstrate that FAR improves the state-of-the-art accuracy by **8.02%-17.61%** across popular UAV datasets.

Method	Frames	Input Size	Init.	Top-1
(i) Dataset: UAV Human Night [43]				
I3D [7]	8	480×640	Kinetics	28.72
FAR	8	480×640	Kinetics	33.78
(ii) Dataset: Drone Action [52]				
HLPF [34]	All	1920×1080	None	64.36
PCNN [9]	-	1920×1080	None	75.92
X3D-M [21]	16	224×224	Kinetics	83.4
FAR	16	224×224	Kinetics	92.7
(iii) NEC Drone [13]				
X3D-M [21]	8	960×540	Kinetics	66.15
FAR	8	960×540	Kinetics	71.46

5.4 Results: More UAV Datasets

We demonstrate the effectiveness of FAR on multiple UAV benchmarks in Table 3. We demonstrate that FAR outperforms prior work by **17.61%, 11.15% and 8.02%** on UAV Human Night, Drone Action, and NEC Drone respectively.

6 Conclusions, Limitations and Future Work

We present a new method for UAV Video Action Recognition. Our method exploits the mathematical properties of Fourier transform to automatically disentangle object from the background, and to encode long-range space-time relationships in a computationally efficient manner. We demonstrate benefits in terms of accuracy, computational complexity and training time on multiple UAV datasets. Our method has a few limitations. The sampling strategy based on randomly initialization is a naive method to span all video frames. It might be interesting to explore the usage of better video sampling strategies [80,37]. Next, we assume that the input videos contain only one human agent performing action. Multi-action videos could be a potential extension of our method. Moreover, we believe that FAR can be extended to other tasks such as video object segmentation and video generation, front-camera action recognition, graphics and rendering [45,47].

Acknowledgements: We thank Rohan Chandra for reviewing the paper. This research has been supported by ARO Grants W911NF1910069, W911NF2110026 and Army Cooperative Agreement W911NF2120076.

References

1. Arnab, A., Dehghani, M., Heigold, G., Sun, C., Lučić, M., Schmid, C.: Vivit: A video vision transformer. arXiv preprint arXiv:2103.15691 (2021) [2](#), [6](#)
2. Barekatain, M., Martí, M., Shih, H.F., Murray, S., Nakayama, K., Matsuo, Y., Prendinger, H.: Okutama-action: An aerial view video dataset for concurrent human action detection. In: Proceedings of the IEEE conference on computer vision and pattern recognition workshops. pp. 28–35 (2017) [3](#)
3. Beauchemin, S.S., Barron, J.L.: The computation of optical flow. ACM computing surveys (CSUR) **27**(3), 433–466 (1995) [2](#)
4. Benjdira, B., Bazi, Y., Koubaa, A., Ouni, K.: Unsupervised domain adaptation using generative adversarial networks for semantic segmentation of aerial images. Remote Sensing **11**(11), 1369 (2019) [2](#)
5. Bertasius, G., Wang, H., Torresani, L.: Is space-time attention all you need for video understanding? arXiv preprint arXiv:2102.05095 (2021) [1](#), [2](#), [3](#), [6](#), [7](#), [8](#), [12](#), [13](#)
6. Buijs, H., Pomerleau, A., Fournier, M., Tam, W.: Implementation of a fast fourier transform (fft) for image processing applications. IEEE Transactions on Acoustics, Speech, and Signal Processing **22**(6), 420–424 (1974) [4](#)
7. Carreira, J., Zisserman, A.: Quo vadis, action recognition? a new model and the kinetics dataset. In: proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6299–6308 (2017) [1](#), [3](#), [5](#), [10](#), [12](#), [13](#), [14](#)
8. Chen, Y., Li, W., Sakaridis, C., Dai, D., Van Gool, L.: Domain adaptive faster r-cnn for object detection in the wild. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3339–3348 (2018) [2](#)
9. Chéron, G., Laptev, I., Schmid, C.: P-cnn: Pose-based cnn features for action recognition. In: Proceedings of the IEEE international conference on computer vision. pp. 3218–3226 (2015) [14](#)
10. Chi, L., Jiang, B., Mu, Y.: Fast fourier convolution. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M.F., Lin, H. (eds.) Advances in Neural Information Processing Systems. vol. 33, pp. 4479–4488. Curran Associates, Inc. (2020), <https://proceedings.neurips.cc/paper/2020/file/2fd5d41ec6cfab47e32164d5624269b1-Paper.pdf> [4](#)
11. Choi, J.: Action recognition list of papers. In: <https://github.com/jinwchoi/awesome-action-recognition> [3](#)
12. Choi, J., Gao, C., Messou, J.C., Huang, J.B.: Why can’t i dance in the mall? learning to mitigate scene bias in action recognition. arXiv preprint arXiv:1912.05534 (2019) [2](#)
13. Choi, J., Sharma, G., Chandraker, M., Huang, J.B.: Unsupervised and semi-supervised domain adaptation for action recognition from drones. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 1717–1726 (2020) [3](#), [10](#), [14](#)
14. Chun, B.T., Bae, Y., Kim, T.Y.: Automatic text extraction in digital videos using fft and neural network. In: FUZZ-IEEE’99. 1999 IEEE International Fuzzy Systems. Conference Proceedings (Cat. No. 99CH36315). vol. 2, pp. 1112–1115. IEEE (1999) [4](#)
15. Ding, M., Li, N., Song, Z., Zhang, R., Zhang, X., Zhou, H.: A lightweight action recognition method for unmanned-aerial-vehicle video. In: 2020 IEEE 3rd International Conference on Electronics and Communication Engineering (ICECE). pp. 181–185. IEEE (2020) [3](#)

16. Dosovitskiy, A., Fischer, P., Ilg, E., Hausser, P., Hazirbas, C., Golkov, V., Van Der Smagt, P., Cremers, D., Brox, T.: FlowNet: Learning optical flow with convolutional networks. In: Proceedings of the IEEE international conference on computer vision. pp. 2758–2766 (2015) [2](#)
17. Du, D., Qi, Y., Yu, H., Yang, Y., Duan, K., Li, G., Zhang, W., Huang, Q., Tian, Q.: The unmanned aerial vehicle benchmark: Object detection and tracking. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 370–386 (2018) [4](#)
18. Dundar, A., Shih, K.J., Garg, A., Pottorf, R., Tao, A., Catanzaro, B.: Unsupervised disentanglement of pose, appearance and background from images and videos. arXiv preprint arXiv:2001.09518 (2020) [2](#)
19. Ellenfeld, M., Moosbauer, S., Cardenes, R., Klauck, U., Teutsch, M.: Deep fusion of appearance and frame differencing for motion segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4339–4349 (2021) [2](#)
20. Fan, H., Xiong, B., Mangalam, K., Li, Y., Yan, Z., Malik, J., Feichtenhofer, C.: Multiscale vision transformers. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6824–6835 (2021) [12](#)
21. Feichtenhofer, C.: X3d: Expanding architectures for efficient video recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 203–213 (2020) [1](#), [3](#), [5](#), [10](#), [12](#), [13](#), [14](#)
22. Feichtenhofer, C., Fan, H., Malik, J., He, K.: Slowfast networks for video recognition. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 6202–6211 (2019) [3](#)
23. Feichtenhofer, C., Pinz, A., Zisserman, A.: Convolutional two-stream network fusion for video action recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1933–1941 (2016) [3](#)
24. Frigo, M., Johnson, S.G.: Fftw: An adaptive software architecture for the fft. In: Proceedings of the 1998 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP'98 (Cat. No. 98CH36181). vol. 3, pp. 1381–1384. IEEE (1998) [5](#), [7](#)
25. Fu, J., Liu, J., Tian, H., Li, Y., Bao, Y., Fang, Z., Lu, H.: Dual attention network for scene segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3146–3154 (2019) [13](#)
26. Gammulle, H., Denman, S., Sridharan, S., Fookes, C.: Two stream lstm: A deep fusion framework for human action recognition. In: 2017 IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 177–186. IEEE (2017) [3](#)
27. Girdhar, R., Carreira, J., Doersch, C., Zisserman, A.: Video action transformer network. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 244–253 (2019) [3](#)
28. Gowda, S.N., Rohrbach, M., Sevilla-Lara, L.: Smart frame selection for action recognition. arXiv preprint arXiv:2012.10671 (2020) [3](#)
29. Griffin, B.A., Corso, J.J.: BubbleNets: Learning to select the guidance frame in video object segmentation by deep sorting frames. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8914–8923 (2019) [3](#)
30. Hara, K., Kataoka, H., Satoh, Y.: Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet? In: Proceedings of the IEEE conference on Computer Vision and Pattern Recognition. pp. 6546–6555 (2018) [3](#)
31. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016) [6](#)

32. Hussein, N., Gavves, E., Smeulders, A.W.: Timeception for complex action recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 254–263 (2019) [3](#)
33. Ilg, E., Mayer, N., Saikia, T., Keuper, M., Dosovitskiy, A., Brox, T.: FlowNet 2.0: Evolution of optical flow estimation with deep networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2462–2470 (2017) [2](#)
34. Jhuang, H., Gall, J., Zuffi, S., Schmid, C., Black, M.J.: Towards understanding action recognition. In: Proceedings of the IEEE international conference on computer vision. pp. 3192–3199 (2013) [14](#)
35. Katharopoulos, A., Vyas, A., Pappas, N., Fleuret, F.: Transformers are rnns: Fast autoregressive transformers with linear attention. In: International Conference on Machine Learning. pp. 5156–5165. PMLR (2020) [4](#)
36. Kim, Y.J., Awadalla, H.H.: Fastformers: Highly efficient transformer models for natural language understanding. arXiv preprint arXiv:2010.13382 (2020) [4](#)
37. Korbar, B., Tran, D., Torresani, L.: Scsampler: Sampling salient clips from video for efficient action recognition. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6232–6242 (2019) [3](#), [9](#), [14](#)
38. Lavin, A., Gray, S.: Fast algorithms for convolutional neural networks. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4013–4021. IEEE Computer Society, Los Alamitos, CA, USA (jun 2016). <https://doi.org/10.1109/CVPR.2016.435>, <https://doi.ieeecomputersociety.org/10.1109/CVPR.2016.435> [4](#)
39. Lee, M., Lee, S., Son, S., Park, G., Kwak, N.: Motion feature network: Fixed motion filter for action recognition. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 387–403 (2018) [2](#)
40. Lee-Thorp, J., Ainslie, J., Eckstein, I., Ontanon, S.: Fnet: Mixing tokens with fourier transforms (2021) [12](#)
41. Li, K., Wu, Z., Peng, K.C., Ernst, J., Fu, Y.: Tell me where to look: Guided attention inference network. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 9215–9223 (2018) [2](#)
42. Li, R., Su, J., Duan, C., Zheng, S.: Linear attention mechanism: An efficient attention for semantic segmentation. arXiv preprint arXiv:2007.14902 (2020) [4](#)
43. Li, T., Liu, J., Zhang, W., Ni, Y., Wang, W., Li, Z.: Uav-human: A large benchmark for human behavior understanding with unmanned aerial vehicles. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16266–16275 (2021) [1](#), [3](#), [10](#), [14](#)
44. Liu, Z., Ning, J., Cao, Y., Wei, Y., Zhang, Z., Lin, S., Hu, H.: Video swin transformer. arXiv preprint arXiv:2106.13230 (2021) [2](#), [6](#)
45. Lloyd, D.B., Govindaraju, N.K., Quammen, C., Molnar, S.E., Manocha, D.: Logarithmic perspective shadow maps. ACM Transactions on Graphics (TOG) **27**(4), 1–32 (2008) [14](#)
46. Mazzia, V., Angarano, S., Salvetti, F., Angelini, F., Chiaberge, M.: Action transformer: A self-attention model for short-time human action recognition. arXiv preprint arXiv:2107.00606 (2021) [3](#)
47. Mitchell, D.P., Netravali, A.N.: Reconstruction filters in computer-graphics. ACM Siggraph Computer Graphics **22**(4), 221–228 (1988) [14](#)
48. Mittal, P., Singh, R., Sharma, A.: Deep learning-based object detection in low-altitude uav datasets: A survey. Image and Vision computing **104**, 104046 (2020) [4](#)
49. Monfort, M., Andonian, A., Zhou, B., Ramakrishnan, K., Bargal, S.A., Yan, T., Brown, L., Fan, Q., Gutfreund, D., Vondrick, C., et al.: Moments in time dataset:

- one million videos for event understanding. *IEEE transactions on pattern analysis and machine intelligence* **42**(2), 502–508 (2019) [3](#)
50. Ng, J.Y.H., Choi, J., Neumann, J., Davis, L.S.: Actionflownet: Learning motion representation for action recognition. In: 2018 IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 1616–1624. IEEE (2018) [2](#)
 51. Peng, H., Razi, A.: Fully autonomous uav-based action recognition system using aerial imagery. In: International Symposium on Visual Computing. pp. 276–290. Springer (2020) [3](#)
 52. Perera, A.G., Law, Y.W., Chahl, J.: Drone-action: An outdoor recorded drone video dataset for action recognition. *Drones* **3**(4), 82 (2019) [1](#), [3](#), [10](#), [14](#)
 53. Piccardi, M.: Background subtraction techniques: a review. In: 2004 IEEE International Conference on Systems, Man and Cybernetics (IEEE Cat. No. 04CH37583). vol. 4, pp. 3099–3104. IEEE (2004) [2](#)
 54. Piergiovanni, A., Ryoo, M.S.: Representation flow for action recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9945–9953 (2019) [2](#)
 55. Plizzari, C., Cannici, M., Matteucci, M.: Spatial temporal transformer network for skeleton-based action recognition. *arXiv preprint arXiv:2012.06399* (2020) [3](#)
 56. Reddy, B.S., Chatterji, B.N.: An fft-based technique for translation, rotation, and scale-invariant image registration. *IEEE transactions on image processing* **5**(8), 1266–1271 (1996) [4](#)
 57. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems* **28**, 91–99 (2015) [2](#), [4](#)
 58. Ren, Z., Yan, J., Ni, B., Liu, B., Yang, X., Zha, H.: Unsupervised deep learning for optical flow estimation. In: Thirty-First AAAI Conference on Artificial Intelligence (2017) [2](#)
 59. Schlag, I., Irie, K., Schmidhuber, J.: Linear transformers are secretly fast weight programmers. In: International Conference on Machine Learning. pp. 9355–9366. PMLR (2021) [4](#)
 60. Sengupta, S., Jayaram, V., Curless, B., Seitz, S.M., Kemelmacher-Shlizerman, I.: Background matting: The world is your green screen. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2291–2300 (2020) [2](#)
 61. Shen, Z., Zhang, M., Zhao, H., Yi, S., Li, H.: Efficient attention: Attention with linear complexities. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 3531–3539 (2021) [4](#), [12](#), [13](#)
 62. Shi, F., Lee, C., Qiu, L., Zhao, Y., Shen, T., Muralidhar, S., Han, T., Zhu, S.C., Narayanan, V.: Star: Sparse transformer-based action recognition. *arXiv preprint arXiv:2107.07089* (2021) [4](#)
 63. Sigurdsson, G.A., Gupta, A., Schmid, C., Farhadi, A., Alahari, K.: Charades-ego: A large-scale dataset of paired third and first person videos. *arXiv preprint arXiv:1804.09626* (2018) [1](#), [3](#)
 64. Simonyan, K., Zisserman, A.: Two-stream convolutional networks for action recognition in videos. *arXiv preprint arXiv:1406.2199* (2014) [3](#)
 65. Sultani, W., Shah, M.: Human action recognition in drone videos using a few aerial training examples. *Computer Vision and Image Understanding* **206**, 103186 (2021) [3](#)
 66. Tancik, M., Srinivasan, P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J., Ng, R.: Fourier features let networks learn high frequency functions in low dimensional domains. In:

- Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M.F., Lin, H. (eds.) Advances in Neural Information Processing Systems. vol. 33, pp. 7537–7547. Curran Associates, Inc. (2020), <https://proceedings.neurips.cc/paper/2020/file/55053683268957697aa39fba6f231c68-Paper.pdf> 4
67. Tran, D., Wang, H., Torresani, L., Feiszli, M.: Video classification with channel-separated convolutional networks. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5552–5561 (2019) 3
 68. Ulhaq, A., Yin, X., Zhang, Y., Gondal, I.: Action-02mcf: A robust space-time correlation filter for action recognition in clutter and adverse lighting conditions. In: International Conference on Advanced Concepts for Intelligent Vision Systems. pp. 465–476. Springer (2016) 3
 69. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. In: Advances in neural information processing systems. pp. 5998–6008 (2017) 2, 3, 4, 7, 8
 70. Wang, M., Deng, W.: Deep visual domain adaptation: A survey. Neurocomputing **312**, 135–153 (2018) 2
 71. Wang, S.e.a.: Linformer: Self-attention with linear complexity. arXiv:2006.04768 (2020) 4
 72. Wang, X., Girshick, R., Gupta, A., He, K.: Non-local neural networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7794–7803 (2018) 3
 73. Wang, X., Jabri, A., Efros, A.A.: Learning correspondence from the cycle-consistency of time. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2566–2576 (2019) 3
 74. Xiong, Y., Zeng, Z., Chakraborty, R., Tan, M., Fung, G., Li, Y., Singh, V.: Nystömformer: A nystöm-based algorithm for approximating self-attention. In: Proceedings of the... AAAI Conference on Artificial Intelligence. AAAI Conference on Artificial Intelligence. vol. 35, p. 14138. NIH Public Access (2021) 4
 75. Xu, K., Qin, M., Sun, F., Wang, Y., Chen, Y., Ren, F.: Learning in the frequency domain. In: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1737–1746. IEEE Computer Society, Los Alamitos, CA, USA (jun 2020). <https://doi.org/10.1109/CVPR42600.2020.00181>, <https://doi.ieeecomputersociety.org/10.1109/CVPR42600.2020.00181> 4
 76. Yang, Y., Soatto, S.: Fda: Fourier domain adaptation for semantic segmentation. 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 4084–4094 (2020) 4
 77. Zappella, L., Lladó, X., Salvi, J.: Motion segmentation: A review. Artificial Intelligence Research and Development pp. 398–407 (2008) 2
 78. Zhang, H., Goodfellow, I., Metaxas, D., Odena, A.: Self-attention generative adversarial networks. In: International conference on machine learning. pp. 7354–7363. PMLR (2019) 2, 3, 6
 79. Zhang, Z., Zhao, J., Zhang, D., Qu, C., Ke, Y., Cai, B.: Contour based forest fire detection using fft and wavelet. In: 2008 International conference on computer science and software engineering. vol. 1, pp. 760–763. IEEE (2008) 4
 80. Zhi, Y., Tong, Z., Wang, L., Wu, G.: Mgsampler: An explainable sampling strategy for video action recognition. arXiv preprint arXiv:2104.09952 (2021) 3, 9, 14
 81. Zhu, Y., Deng, C., Cao, H., Wang, H.: Object and background disentanglement for unsupervised cross-domain person re-identification. Neurocomputing **403**, 88–97 (2020) 2
 82. Zou, Z., Shi, Z., Guo, Y., Ye, J.: Object detection in 20 years: A survey. arXiv preprint arXiv:1905.05055 (2019) 2