

DiffFAR: Differentiable Frequency-based Disentanglement for Aerial Video Action Recognition

Divya Kothandaraman, Ming Lin and Dinesh Manocha
University of Maryland College Park
<https://gamma.umd.edu/diffar>

Abstract—We present a learning algorithm, DiffFAR, for human activity recognition in videos. Our approach is designed for UAV videos, which are mainly acquired from obliquely placed dynamic cameras that contain a human actor along with background motion. Typically, the human actors occupy less than one-tenth of the spatial resolution. DiffFAR simultaneously harnesses the benefits of frequency domain representations, a classical analysis tool in signal processing, and data driven neural networks. We build a differentiable static-dynamic frequency mask prior to model the salient static and dynamic pixels in the video, crucial for the underlying task of action recognition. We use this differentiable mask prior to enable the neural network to intrinsically learn disentangled feature representations via an identity loss function. Our formulation empowers the network to inherently compute disentangled salient features within its layers. Further, we propose a cost-function encapsulating temporal relevance and spatial content to sample the most important frame within uniformly spaced video segments. We conduct extensive experiments on the UAV Human dataset and the NEC Drone dataset and demonstrate relative improvements of 5.72% – 13.00% over the state-of-the-art and 14.28% – 38.05% over the corresponding baseline model.

I. INTRODUCTION

UAVs are being increasingly used for surveillance, security, inspection, search and rescue, agriculture and videography. Action recognition [1] for videos captured by these UAVs is a crucial video analysis task. While there has been immense work on videos obtained using front-view cameras, research on UAV videos is limited. Deep neural networks for human activity recognition [2], [3] have been widely used for ground-camera scenes [4], [5], where the human actors take a high fraction of the pixels in the video scenes. While these methods work well on ground-camera videos, their accuracy on UAV videos can be low. This is mainly because UAV videos consist of human actors that occupy less than 10% of the spatial resolution of the video. Moreover, the action captured by these cameras may be low-resolution, and there is background and camera movement.

Moving objects in videos can be detected using object detection and tracking [6], optical flow [7], motion segmentation [8], etc. However, many issues arise in terms of using them on UAV videos: unavailability of ground-truth object detection labels, domain gap between scenes with access to object localization and scenes on which action recognition needs to be performed [9], [10], [11], the need for a separate learning technique for modeling motion [12], etc. Thus, the network needs to automatically disentangle regions of the feature map, which are vital for video recognition, from




Fig. 1: Human activity recognition in videos recorded using dynamic cameras at oblique and aerial angles typically consist of human actors that occupy less than 10% of the spatial resolution of the video. Background motion, noise, and low-resolution activities add to the complexities faced by video recognition systems. Our method achieves upto **41.86%** and **80.75%** top-1 accuracies on UAV Human and NEC Drone, respectively.

entangled representations encompassing all aspects of the video.

The design of a differentiable method to compute disentangled representations enables the neural network to intrinsically learn disentangled regions, from motion cues and learned saliency information. A differentiable formulation between the neural network and motion representations establishes a bridge between classical and learning-based techniques, harnessing the power of both. Differentiable programming has been used in diverse tasks such as simulation [13] and control [14]. However, it has not been explored in action recognition literature.

A. Main Contributions

We present a novel method, DiffFAR, for activity recognition in UAV videos with dynamic backgrounds and moving cameras. Our approach is designed to handle *oblique* and *aerial* camera views with *low actor-background pixel ratios*, which are common characteristics in videos captured using aerial cameras. We take advantage of frequency domain representations [15] of the feature maps that provide an alternate complementary formulation to holistically model the time-evolution of pixels corresponding to various entities in the scene. We propose a differentiable static-dynamic frequency mask that (i) empowers the network to learn disentangled feature maps via a differentiable loss function and (ii) samples crucial frames from the video at test time. The novel components of our work include:

- 1) **Differentiable static-dynamic mask:** “Dynamic” saliency mask denotes important regions of the scene corresponding to moving pixels. Similarly, “static” saliency mask represents important regions of the scene corresponding to stationary pixels. We propose a *differentiable* frequency-based “static” saliency mask that complements the frequency-based “dynamic” saliency mask. While the “dynamic” mask highlights dynamic salient regions of the scene and suppresses static non-salient regions, our “static” mask amplifies static salient regions and suppresses dynamic non-salient regions. Our approach uses a combination of *both masks to handle dynamic and static salient regions* of the scene (including *moving human actors*) and suppresses dynamic backgrounds, noise, and static non-salient regions.
- 2) **Differentiable disentanglement:** We bestow the network with the ability to inherently “learn” disentangled feature representations that are crucial for the underlying task of action recognition, via an identity loss function using the “*differentiable*” *static-dynamic frequency mask*. This empowers layers of the network to serve as the disentangler and directly compute salient features in its forward pass.
- 3) **Mask sampling:** We propose a frame sampling method using our *static-dynamic frequency mask*, to select the most informative frame within uniformly spaced non-overlapping video segments. These frames serve as a complementary descriptor to frames chosen using uniform sampling.

DiffFAR can be embedded within any 3D action recognition backbone. It is parameter-free i.e. it does not add any learnable layers or parameters to the backbone neural network, which simplifies initialization and training. We conduct extensive experiments on two UAV datasets: UAV Human [16] and NEC Drone [17]. These datasets are captured at varying altitudes, angles and lighting conditions. We demonstrate a considerable improvement of **2.84%** – **13.05%** in Top-1 accuracy over the corresponding baselines. We demonstrate state-of-the-art performance on UAV Human and NEC Drone by at least **3.26%** and **7.74%**, respectively, over the state-of-the-art methods. Moreover, DiffFAR is on par with the state-of-the-art transformer architectures while using far fewer computations (up to **61.8×**) and network parameters (up to **31.94×**).

II. RELATED WORK

a) Aerial Video Recognition.: The availability of UAV datasets [16], [18], [19] has fostered aerial video research [20] pertaining to person reidentification [21], human detection [22], tracking [23], [24], pose estimation [25], few-shot learning [26], drone detection [27] and path planning [28]. Many architectures have been proposed to specifically tackle aerial video action recognition [29], [30], [31], [32], in addition to generic action recognition [2], [3]. Recently, FAR [15] proposed a frequency-based method to disentangle moving objects by modulating feature maps. However, FAR does not place any explicit emphasis on static

regions of the scene, a complementary descriptor. Moreover, FAR applies the mask as a transformation over feature maps, denying the network the opportunity to directly ‘learn’ disentangled feature maps in a differentiable manner.

b) Differentiability and Deep Learning.: Differentiability has emerged as a powerful tool towards making neural networks learn tailored feature representations for diverse tasks such as simulation [13], control [14], video structure from motion [33], and physics [34]. In this paper, we present a differentiable frequency-domain method for aerial video action recognition.

c) Motion Representation Methods.: A popular motion representation method is optical flow [35], which is very expensive. Deep learning based optical flow computation [36], [37], [7] has a lower computational complexity, but requires a two-stream NN [38] for action recognition that increases the network complexity. Object tracking methods [39] require ground-truth bounding boxes that are difficult to obtain. Other alternatives such as motion feature networks [40] and ActionFlowNet [41] are less accurate optical flow. Techniques such as background subtraction [42] and motion segmentation [8] are not ideal for action recognition methods that rely heavily on object localization [43], [44]. In contrast, we present a computationally efficient frequency-based disentanglement technique to automatically highlight regions of the scene salient for action recognition.

d) Frame sampling.: Frame sampling [45] has mainly been explored in the context of untrimmed videos [46] using Reinforcement learning [47], saliency information [48], audio information [49], etc. Most of these methods require the neural network to be retrained using the selected frames which is expensive. Moreover, CNN methods such as MG Sampler [50] use only prior frames, while it is beneficial to use future frames as well in a non-online setting.

III. DIFFFAR

We present a novel method, DiffFAR, for video recognition in UAV videos. These videos, recorded using dynamic cameras placed at oblique angles obstructing the field of view of the action being performed, contain human actors occupying fewer than 10% of the pixels, along with background motion and light noise. Our method has multiple benefits: (i) It can be embedded within any 3D action recognition backbone without modifying the structure of the underlying neural network. (ii) It is parameter-free, i.e., it does not add any learnable layers or parameters to the backbone neural network. These properties allow us to initialize the neural network with weights pre-trained on large-scale ground-camera datasets such as Kinetics [4] and simplifies training.

We present DiffFAR in Figure 2. We sample T frames from the video using uniform sampling [15], where T is the frame rate. The architecture involves a 3D CNN backbone such as I3D [4] or X3D [3]. 3D CNNs typically consist of an initial set of convolution layers, followed by 4 blocks of ResNet-type convolution layers. This is followed by a final prediction layer. Our architecture retains the structure of this backbone. We propose a differentiable static-dynamic

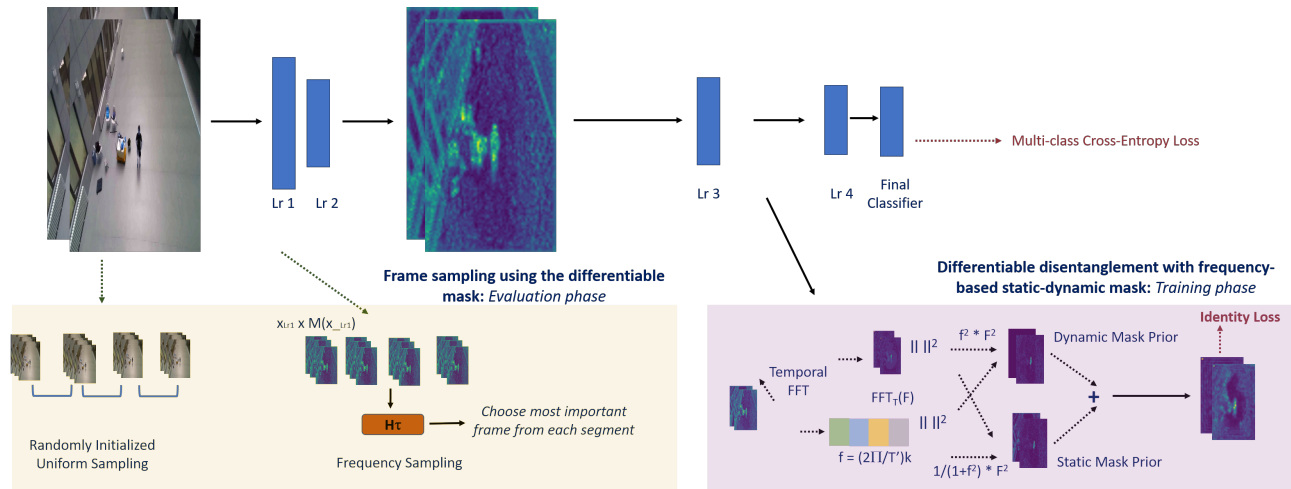


Fig. 2: We present a differentiable frequency-based method for aerial video recognition. Our differentiable static-dynamic frequency mask provides a prior for disentangled regions relevant to action recognition. This mask is used to guide the learning of disentangled features within the layers of the neural network using an identity function. Furthermore, we propose a frame sampling strategy that chooses the best frame within each uniform video segment at test time, using the static-dynamic frequency mask and temporal difference.

Notation	Term
$M_{dynamic}$	Dynamic mask
M_{static}	Static mask
M	$M_{dynamic} + M_{static}$
f	Frequencies
F	Fourier transform
x	Feature representations
C_p	Prior frames cost
C_f	Future frames cost
T	Frame sampling rate

TABLE I: Symbols used in the paper

frequency mask to highlight regions of the scene relevant for action prediction. This is computed using the outputs at the second block of the neural network. The neural network is trained using traditional multi-class cross entropy loss for action recognition and an identity loss. The identity loss function is applied at the outputs of the third block of the neural network, and bestows the network with the ability to learn disentangled feature maps in a differentiable manner via an identity loss function. We also propose a method for frame sampling using the differentiable static-dynamic frequency mask, which samples frames from the video at the evaluation phase. We now describe our method in detail.

A. Background: Differentiability

Differentiable programming allows for computing derivatives of functions or gradient operators to learn representations for the underlying task. For feature disentanglement, differentiable programming can be a more powerful tool than applying transformations over features [15]. This is because blocks of the network are optimized to directly compute disentangled features, where the entire set of layers serves as the disentanglement module.

Frequency domain representations contain an abundance

of information for disentanglement that can improve the performance of action recognition neural networks. One way to do so is to apply the corresponding frequency-domain masks as a transformation over neural network features. However, enabling the neural network to learn disentangled representations within its layers, using frequency-domain information as well as training data, creates an amalgamation of techniques from traditional signal processing and deep learning. This has the potential to empower the neural network with the benefits of both worlds.

However, the training of neural networks requires loss functions to be differentiable in order to perform backpropagation. This raises questions pertaining to frequency domain representations that can be used for training, appropriate loss function and the layers of the neural network that should be used for learning robust disentangled features.

B. Differentiable Static-Dynamic Frequency Mask

Consider a video depicting the action of human diving. The focus of the camera shifts from the diving springboard to the water beneath as the action is executed. Moreover, the springboard might move as the person jumps into the water, and there may be other humans moving around the swimming pool. In the absence of any external localization ground-truth, the neural network needs to innately disentangle the human actor performing the action of diving from the rest of the scene. In general, a video scene consists of four types of regions - dynamic salient regions, static salient regions, dynamic non-salient regions and static non-salient regions. The first two are very important for the underlying task of action recognition. A robust feature disentanglement technique should highlight pixels corresponding to static salient regions and dynamic salient regions and suppress the non-salient patches.

The frequency domain of a signal is a powerful representation that can quantify the temporal change of spatial pixels, and hence can be used for predicting static and dynamic regions of the scene. To identify dynamic pixels, we can formulate the “dynamic” mask $M_{dynamic}$ [15] as the dot product of the temporal FFT F_T of the feature maps and the corresponding frequencies f , over the frequency spectrum $k = 1 \dots T$. For features x , mathematically,

$$M_{FAR}(x) = \underbrace{\|F_T(x)\|_2^2 \times \|f\|_2^2}_{M_{dynamic}} \odot x, \quad (1)$$

The formulation represented by Equation 1 amplifies high amplitudes at high frequencies, and suppresses low amplitudes at low frequencies. The final dot product with the feature representation ensures that only salient moving regions are amplified, and moving background regions are suppressed. Since this equation only models dynamic regions of the scene, it causes inconsistencies in static parts of the scene corresponding to salient regions; it may also incorrectly disentangle dynamic backgrounds and noise. To alleviate this issue, we propose modeling the salient static parts of the video through a static mask M_{static} .

High amplitudes at low frequencies correspond to static regions of the scene. Given a feature representation x , we mathematically model this relation as follows:

$$M_{static} = \|F_T(x)\|_2^2 \odot (1/(1 + \|f\|_2^2)), \quad (2)$$

The inverse shifted quadratic formulation of the frequency component ensures that the computation of the mask is stable. A dot product of M_{static} and network features amplifies static, salient regions and heavily suppresses dynamic non-salient regions. M_{static} , in conjunction with the dynamic mask $M_{dynamic}$, forms a static-dynamic frequency mask prior M . Thus, the net mask encapsulating static, and dynamic parts of the scene is $M = M_{static} + M_{dynamic}$. M guides the network to *learn* object disentangled low-level feature representations, as described in the following section.

C. Differentiable Disentanglement

Our next goal is to use the differentiable static-dynamic frequency masks to help the neural network learn disentangled feature representations. We guide the learning of disentangled feature representations at the penultimate layer of the network using the static-dynamic mask prior M from mid-level features. Since mid-level features provide a fine balance between encoding video features and maintaining the space-time resolution of the video, we use mid-level features x_{mid} to derive the disentanglement mask $x_{mid,M}$, computed by applying M on x_{mid} . We constrain the features at the penultimate layer, x_p to resemble disentangled feature representations using an identity loss function as follows:

$$L_{mask} = \lambda_{mask} \cdot MSE\{x_p, M(x_{mid}) \odot x_p\}, \quad (3)$$

where MSE is mean squared error. In the above equation, the dot product of mask $M(x_{mid})$ (M applied on x_{mid}) and x_p transforms x_p to the corresponding disentangled feature representation. Further, the application of the identity

loss enables the network to intrinsically learn high-quality disentangled feature representations at the penultimate layer. During backpropagation, gradients flow from x_p as well as x_{mid} (through M) ensuring that neural network layers between the mid-level and penultimate layer learn disentangled features. The hyperparameter λ_{mask} , typically set to 0.1 in our benchmarks, scales the identity loss function. We use a combination of the classical multi-class cross entropy loss function [4] applied on the final softmax layer of the neural network, along with L_{mask} , to train the action recognition neural network.

D. Frame Sampling

Uniform sampling divides the video into uniform segments, equal to the number of frames N to be sampled, and selects the k_{th} from each segment where k is an integer lesser than or equal to the number of frames in each segment. Such a sampling strategy is commonly adopted during the training of action recognition architectures on trimmed videos and works well due to its ability to span the duration of the video. However, samples selected from each video segment using uniform sampling are not always the most salient frames in the corresponding segment. We present a method to choose the best frame within uniform video segments. Our solution can be applied at the evaluation phase and does not require the neural network to be retrained.

A good frame selection strategy should consider frames prior to the video segment under consideration as well as frames succeeding the video segment. Frames that provide maximal information w.r.t. prior as well as future frames can be regarded as salient frames. For the task of action recognition, movement or difference of feature representations between frames provides this information. Corresponding to segment i ($i = 1 \dots T$), where T is the number of frames, for frame j 's ($j = 1 \dots s$), where s is the number of frames within each segment (or total number of frames in the video divided by sampling rate), we denote the feature representations $x_{i,j}$. Since movement is best encapsulated by the salient regions, which are useful for predictions, we compute the saliency mask representations for features $x_{i,j}$ as $M \odot x_{i,j}$. We denote uniform frames for segment i as x_{U-i} , and the corresponding saliency mask representations as $M \odot x_{U-i}$. Weighted sum of mean squared difference between $M \odot x_{i,j}$ and prior and future frames gives the cost w.r.t. prior C_p and future frames C_f . The weights are designed to be inversely proportional to the distance between the video segment under consideration for frame sampling and the video segment in the equation for computation for prior and future costs. Mathematically,

$$C_p = \sum_{k=1}^{i-1} \frac{T-i+k}{T} \times (M \odot x_{i,j} - M \odot x_{U-k})^2 \quad (4)$$

$$C_f = \sum_{k=i+1}^T \frac{T+i-k}{T} \times (M \odot x_{i,j} - M \odot x_{U-k})^2. \quad (5)$$

The net cost is a weighted sum of C_p and C_f where the weights, w_p and w_f , are proportional to the costs and time between the video segment under consideration for frame

Model	Data Aug.	Input	Frames	Top-1 (%)
I3D Experiments				
I3D	✗	(540, 960)	8	23.86
Ours (I3D Backbone)	✗	(540, 960)	8	32.94
X3D Experiments				
X3D	✗	(540, 960)	16	35.71
Ours (X3D backbone)	✗	(540, 960)	16	40.81
X3D + Data Augmentation Experiments				
X3D	✓	(540, 540)	8	36.6
X3D + FD	✓	(540, 560)	8	41.86

TABLE II: We evaluate our method on UAV Human. Across two 3D CNN backbones (I3D, X3D) and datasets, we demonstrate relative improvements of **14.28%** - **38.05%** in the Top-1 accuracy over the corresponding baselines.

sampling, the beginning of the video, and the end of the video. Mathematically, $w_p = \{C_p / (C_p + C_f)\} \times i$ and $w_f = \{C_p / (C_p + C_f)\} \times (T - i)$.

a) *Ensembling the predictions.*: We use the set of uniform frames F_U to sample the set of salient frames F_S . We use the trained action recognition neural network to compute the predictions for F_U as well as F_S . These predictions are summed to obtain the final prediction, which is used to determine the predicted action class.

IV. EXPERIMENTS

A. Network and Training Details

a) *Backbone network architecture.*: We benchmark our models using two state-of-the-art 3D CNN based action recognition architectures :- I3D (CVPR 2017) and X3D (CVPR 2020). We use level 2 features (mid-level) of the underlying 3D ResNet-101 to derive object disentanglement masks. This mask is used to guide the learning (using Equation 3) of disentangled feature representations obtained after level 3. The cost functions for frequency-based frame sampling are computed using features from level 1.

b) *Training details.*: All our models were trained using NVIDIA GeForce 1080 Ti GPUs and NVIDIA RTX A5000 GPUs. Initial learning rates were set at 0.01 or 0.1. We use the Stochastic Gradient Descent (SGD) for optimization, with 0.0005 weight decay, and 0.9 momentum. We use cosine/poly annealing for learning rate decay and multi-class cross entropy loss to constrain the final softmax predictions of the neural network. We report top-1 accuracies for all our models.

c) *Datasets.*: We benchmark our models on two UAV datasets:- UAV Human [16] and NEC Drone [17]. UAV Human is an outdoor activities dataset, containing low-resolution videos taken under adverse lighting and weather conditions. It contains 155 actions, many of which are similar and hard to distinguish. Moreover, the videos contain dynamic backgrounds, camera motions and noise. NEC Drone is an indoor dataset with single-human and two-person actions. It contains 16 actions performed in an unconstrained manner.

Experiment	Top-1 (%)
$M_{dynamic}$	28.1
$M_{dynamic} + L_{mask}$	29.68
$M_{static} + M_{dynamic} + L_{mask}$	31.95

TABLE III: We demonstrate the effectiveness of each component of FD with ablation experiments on UAV Human. All our experiments use RIUS for training as well as testing, and I3D backbone, and a spatial and temporal resolution of 540×960 and 8 frames respectively.

Model	Backbone	Data Aug.	Input	Frames	Params (M)	FLOPs (G)	Top-1 (%)
UAV Human							
I3D (CVPR 2017) [4]	-	✗	(540, 960)	8	28.0	108	21.06
MVIT (ICCV 2021) [51]	-	✓	(224, 224)	16	36.6	70.8	24.3
X3D-M (CVPR 2020) [3]	-	✓	(224, 224)	8	3.8	6.2	30.6
X3D-M (CVPR 2020) [3]	-	✓	(540, 540)	8	3.8	36.03	36.6
X3D-M (CVPR 2020) [3]	-	✗	(540, 960)	16	3.8	36.03	x
FNet (arXiv 2021) [52]	I3D	✗	(540, 960)	8	28.0	108	24.39
FAR (ECCV 2022) [15]	I3D	✗	(540, 960)	8	28.0	108	29.21
FAR (ECCV 2022) [15]	X3D-M	✓	(540, 540)	8	3.8	36.6	38.6
Ours	X3D-M	✓	(540, 540)	8	3.8	36.6	41.86
Ours	X3D-M	✗	(540, 960)	16	3.8	130.13	40.81

TABLE IV: We compare our method against prior methods on UAV Human. We demonstrate state-of-the-art performance by at least 5.72%. Our method outperforms transformer architectures while using far fewer computations (upto **61.8** \times) and network parameters (upto **31.94** \times). Our method imposes minimal memory overhead over the backbone neural network.

d) *Evaluation.*: We evaluate all action recognition models using the top-1 accuracy score, which is the ratio of the number of correctly classified samples to the total number of samples in the evaluation set. We also report top-5 accuracy which is the ratio of samples for which the top-5 classes (in terms of probability of predictions) contains the ground-truth class. We report all improvements in terms of relative performance. The relative improvement of method 1 over method 2, where the absolute accuracies of method 1 and method 2 are $x_1\%$ and $x_2\%$ respectively, is $(x_1 - x_2) / x_2 \times 100$

B. Results: UAV Human

a) *Our method can be embedded within any state-of-the-art 3D CNN to improve the performance.*: We demonstrate the effectiveness of our method over various backbone architectures in Table II. We report relative improvements of 38.05% and 14.28% using the I3D and X3D backbones respectively. The large improvements in top-1 accuracy are indicative of the ability of FD to inherently localize the human actor from the background and identify actions from obscured top-down and oblique views in the presence of dynamic UAV cameras and noise.

All models use uniform sampling (randomly initialized) [15] for training. In the first two sets of experiments, we do not perform any data augmentation. The differentiable static-dynamic frequency mask and identity loss are used for training and the ensemble frequency-based frame sampling strategy is used at evaluation for our method. In the third set of experiments, the X3D model is trained with spatial data augmentations, consistent with the publicly available official code for X3D [3] in Meta AI’s SlowFast GitHub repository.

Model	Input	Frames	Top-1 (%)	Top-5 (%)
I3D Experiments				
w/o sampling	(540, 960)	8	31.75	52.92
with sampling	(540, 960)	8	32.94	54.40
X3D Experiments				
w/o sampling	(540, 960)	16	40.12	59.08
with sampling	(540, 960)	16	40.81	59.94

TABLE V: We present ablation experiments on the frame sampling method. None of our models use data augmentation.

While we use the differentiable static-dynamic frequency mask and identity loss for training, the frame sampling and evaluation scheme are consistent with the SlowFast repository codebase i.e. our ensemble frame selection strategy is not used for testing.

b) Ablation experiments on the differentiable static-dynamic frequency mask and identity loss.: We present ablation experiments on UAV Human in Table III. We use the I3D backbone at spatial and temporal resolutions of 540×960 and 8 frames respectively. All experiments use uniform sampling (randomly initialized) for training as well as testing.

- In the first experiment, we apply only the dynamic mask $M_{dynamic}$ (Equation 1), as a transformation over mid-level features, as in FAR. The accuracy is 28.1%. [15].
- In the second experiment, we quantitatively demonstrate the benefits of differentiability. We use $M_{dynamic}$ to train the network to learn disentangled features in a differentiable manner using L_{mask} (Equation 3). The absolute improvement over the first experiment is 1.58%.
- In the third experiment, we show the effects of the static-dynamic mask. We use M_{static} (Equation 2) along with the setup in the second experiment. The absolute improvement over the second experiment is 2.27%.

c) Effectiveness of the ensemble frame sampling method.: We present experiments on the effectiveness of the frame sampling method in Table V. On I3D, the frame sampling method improves top-1 and top-5 accuracies by (relative) 3.74% and 2.79% respectively. On X3D, the frame sampling method improves top-1 and top-5 accuracies by (relative) 1.71% and 1.45% respectively.

d) State-of-the-art comparisons.: We present comparisons against the state-of-the-art in Table IV. We compare against 3D CNNs (I3D [4], X3D [3]), transformer models (MVIT [51]) and frequency-based methods (FNet [52], FAR [15]). Our method has benefits in terms of top-1 accuracy as well as computation (FLOPs, number of parameters). It has minimal memory overheads over the corresponding backbone neural network. We demonstrate (relative) improvements of at least 5.72% over the state-of-the-art method.

e) Ablation experiments on the frame sampling method.: All experiments are conducted on the X3D backbone, without data augmentation and a spatial-temporal reso-

Model	Backbone	Input	Frames	Params (M)	FLOPs (G)	Top-1 (%)
X3D-M (CVPR 2020) [3]	-	(960, 540)	8	3.8	64.05	66.15
FAR (ECCV 2022) [15]	X3D-M	(960, 540)	8	3.8	65.08	71.46
Ours	X3D-M	(960, 540)	8	3.8	65.08	80.75

TABLE VI: We present experiments on NEC Drone. We demonstrate improvements of **13%** - **22%** over prior work.

lution of $16 \times 540 \times 960$. The purpose of the disentanglement mask in frame sampling is to ensure samples are chosen based on the relevance of only salient features. When applied on a model trained for learning disentangled features, a good way to verify if the network has actually learnt disentangled features is to remove the mask in the calculation of cost functions for frame sampling. This results in an accuracy of 40.9%, proving that our differentiable method is indeed successful. Setting $w_p = w_f = 1$ results in an accuracy of 40.70%.

C. Results: NEC Drone

We present results on NEC-Drone in Table VI. We use uniform sampling [15] to train all models. In the experiment corresponding to our method, we initialize the neural network with weights corresponding UAV Human. Following prior work [15], while testing, we select the most frequent class prediction among the networks' result using all sets of Randomly Initialized Uniform Sampling frames.

We obtain a top-1 accuracy of 80.75%. We compare against 3D CNNs (X3D [3]) and frequency-based methods (current state-of-the-art for aerial video recognition - FAR [15]) and show relative improvements of **13.00%** - **22.07%**.

V. CONCLUSION, LIMITATIONS AND FUTURE WORK

In this paper, we presented a differentiable disentanglement method that simultaneously exploits the potential of learning-based neural networks and traditional signal processing. The method includes a static-dynamic frequency mask and identity loss for learning the disentanglement, along with a frame selection method at the evaluation stage. Our method has a few limitations. First, it is completely supervised. This means it requires labeled data and annotation is a laborious process. Unsupervised, self-supervised and transfer learning methods can alleviate this issue. Second, while our method is more effective than transformers in terms of accuracy, speed and memory, it is still expensive due to 3D convolutions and requires multiple GPUs and iterations to train. Future work on methods that can be run on mobile GPUs will be of practical importance. In our experiments and datasets, we assumed that a single human-agent is performing the action. Our method can be incorporated within any 3D neural network to improve performance and it might be interesting to explore the usage of our method within multi-agent neural networks.

Acknowledgements: This research has been supported by ARO Grants W911NF2110026 and Army Cooperative Agreement W911NF2120076

REFERENCES

- [1] I. Rodríguez-Moreno, J. M. Martínez-Otzeta, B. Sierra, I. Rodríguez, and E. Jauregi, "Video activity recognition: State-of-the-art," *Sensors*, vol. 19, no. 14, p. 3160, 2019.
- [2] C. Feichtenhofer, H. Fan, J. Malik, and K. He, "Slowfast networks for video recognition," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 6202–6211.
- [3] C. Feichtenhofer, "X3d: Expanding architectures for efficient video recognition," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 203–213.
- [4] J. Carreira and A. Zisserman, "Quo vadis, action recognition? a new model and the kinetics dataset," in *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 6299–6308.
- [5] M. Monfort, A. Andonian, B. Zhou, K. Ramakrishnan, S. A. Bargal, T. Yan, L. Brown, Q. Fan, D. Gutfreund, C. Vondrick, *et al.*, "Moments in time dataset: one million videos for event understanding," *IEEE transactions on pattern analysis and machine intelligence*, vol. 42, no. 2, pp. 502–508, 2019.
- [6] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," *Advances in neural information processing systems*, vol. 28, pp. 91–99, 2015.
- [7] A. Dosovitskiy, P. Fischer, E. Ilg, P. Hausser, C. Hazirbas, V. Golkov, P. Van Der Smagt, D. Cremers, and T. Brox, "FlowNet: Learning optical flow with convolutional networks," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 2758–2766.
- [8] L. Zappella, X. Lladó, and J. Salvi, "Motion segmentation: A review," *Artificial Intelligence Research and Development*, pp. 398–407, 2008.
- [9] M. Wang and W. Deng, "Deep visual domain adaptation: A survey," *Neurocomputing*, vol. 312, pp. 135–153, 2018.
- [10] Y. Chen, W. Li, C. Sakaridis, D. Dai, and L. Van Gool, "Domain adaptive faster r-cnn for object detection in the wild," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 3339–3348.
- [11] B. Benjdira, Y. Bazi, A. Koubaa, and K. Ouni, "Unsupervised domain adaptation using generative adversarial networks for semantic segmentation of aerial images," *Remote Sensing*, vol. 11, no. 11, p. 1369, 2019.
- [12] H. Gammulle, S. Denman, S. Sridharan, and C. Fookes, "Two stream lstm: A deep fusion framework for human action recognition," in *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 2017, pp. 177–186.
- [13] Y. Qiao, J. Liang, V. Koltun, and M. Lin, "Differentiable simulation of soft multi-body systems," *Advances in Neural Information Processing Systems*, vol. 34, pp. 17 123–17 135, 2021.
- [14] T. Takahashi, J. Liang, Y.-L. Qiao, and M. C. Lin, "Differentiable fluids with solid coupling for learning and control," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 7, 2021, pp. 6138–6146.
- [15] D. Kothandaraman, T. Guan, X. Wang, S. Hu, M. Lin, and D. Manocha, "Fourier disentangled space-time attention for aerial video recognition," *arXiv preprint arXiv:2203.10694*, 2022.
- [16] T. Li, J. Liu, W. Zhang, Y. Ni, W. Wang, and Z. Li, "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*, 2021, pp. 16 266–16 275.
- [17] J. Choi, G. Sharma, M. Chandraker, and J.-B. Huang, "Unsupervised and semi-supervised domain adaptation for action recognition from drones," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2020, pp. 1717–1726.
- [18] P. Zhu, L. Wen, D. Du, X. Bian, H. Ling, Q. Hu, Q. Nie, H. Cheng, C. Liu, X. Liu, *et al.*, "Visdrone-det2018: The vision meets drone object detection in image challenge results," in *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, 2018, pp. 0–0.
- [19] I. Bozcan and E. Kayacan, "Au-air: A multi-modal unmanned aerial vehicle dataset for low altitude traffic surveillance," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 8504–8510.
- [20] K. Nguyen, C. Fookes, S. Sridharan, Y. Tian, X. Liu, F. Liu, and A. Ross, "The state of aerial surveillance: A survey," *arXiv preprint arXiv:2201.03080*, 2022.
- [21] M. C. Fysh and M. Bindemann, "Person identification from drones by humans: insights from cognitive psychology," *Drones*, vol. 2, no. 4, p. 32, 2018.
- [22] J. Portmann, S. Lynen, M. Chli, and R. Siegwart, "People detection and tracking from aerial thermal views," in *2014 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2014, pp. 1794–1800.
- [23] Y. He, C. Fu, F. Lin, Y. Li, and P. Lu, "Towards robust visual tracking for unmanned aerial vehicle with tri-attentional correlation filters," in *2020 IEEE/RJSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 1575–1582.
- [24] M. Mueller, G. Sharma, N. Smith, and B. Ghanem, "Persistent aerial tracking system for uavs," in *2016 IEEE/RJSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2016, pp. 1562–1569.
- [25] A. Benini, M. J. Rutherford, and K. P. Valavanis, "Real-time, gpu-based pose estimation of a uav for autonomous takeoff and landing," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 3463–3470.
- [26] W. Sultani and M. Shah, "Human action recognition in drone videos using a few aerial training examples," *Computer Vision and Image Understanding*, vol. 206, p. 103186, 2021.
- [27] M. W. Ashraf, W. Sultani, and M. Shah, "Dogfight: Detecting drones from drones videos," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 7067–7076.
- [28] S. Hayat, E. Yanmaz, T. X. Brown, and C. Bettstetter, "Multi-objective uav path planning for search and rescue," in *2017 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2017, pp. 5569–5574.
- [29] K. Kotecha, D. Garg, B. Mishra, P. Narang, and V. K. Mishra, "Background invariant faster motion modeling for drone action recognition," *Drones*, vol. 5, no. 3, p. 87, 2021.
- [30] H. Peng and A. Razi, "Fully autonomous uav-based action recognition system using aerial imagery," in *International Symposium on Visual Computing*. Springer, 2020, pp. 276–290.
- [31] M. Ding, N. Li, Z. Song, R. Zhang, X. Zhang, and H. Zhou, "A lightweight action recognition method for unmanned-aerial-vehicle video," in *2020 IEEE 3rd International Conference on Electronics and Communication Engineering (ICECE)*. IEEE, 2020, pp. 181–185.
- [32] A. Ulhaq, X. Yin, Y. Zhang, and I. Gondal, "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*. Springer, 2016, pp. 465–476.
- [33] Z. Teed and J. Deng, "Deepv2d: Video to depth with differentiable structure from motion," *arXiv preprint arXiv:1812.04605*, 2018.
- [34] J. Degraeve, M. Hermans, J. Dambre, *et al.*, "A differentiable physics engine for deep learning in robotics," *Frontiers in neurorobotics*, p. 6, 2019.
- [35] S. S. Beauchemin and J. L. Barron, "The computation of optical flow," *ACM computing surveys (CSUR)*, vol. 27, no. 3, pp. 433–466, 1995.
- [36] Z. Ren, J. Yan, B. Ni, B. Liu, X. Yang, and H. Zha, "Unsupervised deep learning for optical flow estimation," in *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- [37] E. Ilg, N. Mayer, T. Saikia, M. Keuper, A. Dosovitskiy, and T. Brox, "FlowNet 2.0: Evolution of optical flow estimation with deep networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2462–2470.
- [38] K. Simonyan and A. Zisserman, "Two-stream convolutional networks for action recognition in videos," *Advances in neural information processing systems*, vol. 27, 2014.
- [39] Y. Zhang, T. Wang, K. Liu, B. Zhang, and L. Chen, "Recent advances of single-object tracking methods: A brief survey," *Neurocomputing*, vol. 455, pp. 1–11, 2021.
- [40] M. Lee, S. Lee, S. Son, G. Park, and N. Kwak, "Motion feature network: Fixed motion filter for action recognition," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 387–403.
- [41] J. Y.-H. Ng, J. Choi, J. Neumann, and L. S. Davis, "Actionflownet: Learning motion representation for action recognition," in *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 2018, pp. 1616–1624.
- [42] M. Piccardi, "Background subtraction techniques: a review," in *2004 IEEE International Conference on Systems, Man and Cybernetics (IEEE Cat. No. 04CH37583)*, vol. 4. IEEE, 2004, pp. 3099–3104.
- [43] S. Sengupta, V. Jayaram, B. Curless, S. M. Seitz, and I. Kemelmacher-Shlizerman, "Background matting: The world is your green screen," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 2291–2300.

- [44] M. Ellenfeld, S. Moosbauer, R. Cardenes, U. Klauck, and M. Teutsch, "Deep fusion of appearance and frame differencing for motion segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 4339–4349.
- [45] Z. Wu, C. Xiong, C.-Y. Ma, R. Socher, and L. S. Davis, "Adaframe: Adaptive frame selection for fast video recognition," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 1278–1287.
- [46] W. Wu, D. He, X. Tan, S. Chen, and S. Wen, "Multi-agent reinforcement learning based frame sampling for effective untrimmed video recognition," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 6222–6231.
- [47] S. N. Gowda, M. Rohrbach, and L. Sevilla-Lara, "Smart frame selection for action recognition," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 2, 2021, pp. 1451–1459.
- [48] B. Korbar, D. Tran, and L. Torresani, "Scsampler: Sampling salient clips from video for efficient action recognition," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 6232–6242.
- [49] R. Gao, T.-H. Oh, K. Grauman, and L. Torresani, "Listen to look: Action recognition by previewing audio," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 10 457–10 467.
- [50] Y. Zhi, Z. Tong, L. Wang, and G. Wu, "Mgsampler: An explainable sampling strategy for video action recognition," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 1513–1522.
- [51] H. Fan, B. Xiong, K. Mangalam, Y. Li, Z. Yan, J. Malik, and C. Feichtenhofer, "Multiscale vision transformers," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 6824–6835.
- [52] J. Lee-Thorp, J. Ainslie, I. Eckstein, and S. Ontanon, "Fnet: Mixing tokens with fourier transforms," 2021.