

Is Appearance Free Action Recognition Possible?

Filip Ilic¹, Thomas Pock¹, and Richard P. Wildes²

¹ Graz University of Technology, Austria

² York University, Canada

Abstract. Intuition might suggest that motion and dynamic information are key to video-based action recognition. In contrast, there is evidence that state-of-the-art deep-learning video understanding architectures are biased toward static information available in single frames. Presently, a methodology and corresponding dataset to isolate the effects of dynamic information in video are missing. Their absence makes it difficult to understand how well contemporary architectures capitalize on dynamic vs. static information. We respond with a novel Appearance Free Dataset (AFD) for action recognition. AFD is devoid of static information relevant to action recognition in a single frame. Modeling of the dynamics is necessary for solving the task, as the action is only apparent through consideration of the temporal dimension. We evaluated 11 contemporary action recognition architectures on AFD as well as its related RGB video. Our results show a notable decrease in performance for all architectures on AFD compared to RGB. We also conducted a complimentary study with humans that shows their recognition accuracy on AFD and RGB is very similar and much better than the evaluated architectures on AFD. Our results motivate a novel architecture that revives explicit recovery of optical flow, within a contemporary design for best performance on AFD and RGB.

Keywords: Action Recognition, Action Recognition Dataset, Deep Learning, Static and Dynamic Video Representation, Human Motion Perception

1 Introduction

1.1 Motivation

Action recognition from video has been subject of significant effort in developing and improving both algorithms and datasets [30,63]. This interplay between algorithm and dataset advances has led to paradigm shifts in both. Algorithms have evolved from primarily hand-crafted mathematically, physically, and heuristically driven approaches to methods based on deep-learning architectures. Datasets have evolved from relatively small, carefully selected videos to

Corresponding author: filip.ilic@tugraz.at

Code & Data: [f-ilic.github.io/AppearanceFreeActionRecognition](https://github.com/f-ilic/AppearanceFreeActionRecognition)

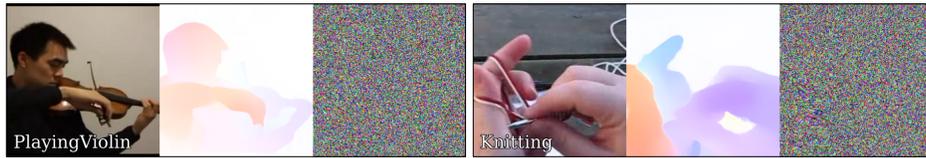


Fig. 1. Appearance Free Dataset (AFD) for Action Recognition. Within each set of three images we show: a single RGB frame (left); corresponding optical flow in Middlebury colour coding [4] (middle); corresponding appearance free frame (right). When viewed as video the AFD reveals the motion in the original video even as any single frame provides no relevant discriminative information.

massive, web-crawled collections. As a result, current state-of-the-art algorithms for action recognition achieve impressive levels of performance on challenging datasets. In contrast, the internal representations learned by the architectures remain under explored [24]. This lack of understanding is unsatisfying both scientifically as well as pragmatically. From a scientific perspective, such detailed analysis is integral to understanding how a system operates and can guide further improvements. From a pragmatic perspective, multiple jurisdictions are beginning to require explainability as a precondition for deployment of artificial intelligence technologies [1,2]; technologies lacking such documentation may not see real-world application.

Some evidence suggests that contemporary deep-learning architectures for video understanding can perform well on the task of action recognition with little to no regard for the actual actors [59,23,9]. These studies suggest that static visual information available in a single frame (*e.g.* colour, spatial texture, contours, shape) drives performance, rather than dynamic information (*e.g.*, motion, temporal texture). While some methodologies have been aimed at understanding the representations learned by these architectures [39,16,61], none provide an approach to completely disentangle static vs. dynamic information. In response, we have developed an Appearance Free Dataset (AFD) for evaluating action recognition architectures when presented with purely dynamic information; see Fig. 1. AFD has been rendered by animating spatial noise, historically known as Random Dot Cinematographs RDCs [29], using image motion extracted from the UCF101 dataset [51] by a state-of-the-art optical flow estimator [53]. The resulting videos show no pattern relevant to action recognition in any single frame; however, when viewed as video reveal the underlying motion. We have produced AFD for the entire UCF101 dataset, evaluated a representative sample of contemporary action recognition architectures and used our results to drive development of a novel architecture that enhances performance when confronted with purely dynamic information, as present in AFD.

The ability of action recognition (and other) video architectures to work in absence of static appearance information is not only an academic exercise. Real-world deployment scenarios may require it, such as in the presence of camou-

flage. Two examples: Video surveillance for security should be able to cope with nefarious actors who artificially camouflage their activities; video-based wildlife monitoring must be robust to the natural camouflage that many animals possess to hide their presence.

1.2 Related work

Datasets A wide variety of datasets for development and evaluation of automated approaches of video understanding are available [3]. For action recognition, in particular, there is a large body ranging from a few dozen classes [36,51,38,39,31] to massively crawled datasets with classes in the hundreds [31,8,19,20]. Some work has specifically focused on curating videos where temporal modeling is particularly important for action recognition [33,48,19,38]; however, they still contain strong cues from single frame static appearance (*e.g.* colour, shape, texture). Overall, no action recognition dataset completely disentangles single frame static information from multiframe dynamic information.

Camouflaged actors would provide a way to evaluate recognition systems based primarily on dynamic information. While there are camouflage video understanding dataset available, they are of animals in the wild as selected for object segmentation (*e.g.* [37,5]). In contrast, it is unlikely that a non-trivial number of real-world videos of camouflaged actors can be found. In response, our Appearance Free Dataset (AFD) provides synthetic camouflaged action recognition videos as coherently moving patterns of spatial noise, historically known as Random Dot Cinematograms (RDCs) [29]. These patterns are defined via animation of an initial random spatial pattern, where the pixelwise intensity values of the pattern are randomly selected. By design, they reveal no information relevant to the motion in any single frame; however, when viewed as video the motion is apparent; examples are shown in the right columns of Fig. 1. Such videos have a long history in the study of motion processing in both biological [6,44] and machine [58,62] vision systems. Interestingly, humans can understand complex human body motion solely from sparse dots marking certain body points, even while merely a random dot pattern is perceived in any single frame [28].

Synthetic video, both striving for photorealism [42,7,11,46,45] and more abstracted [52,40], is a common tool in contemporary computer vision research that allows for careful control of variables of interest. Our effort adds to this body with its unique contribution of a synthetic camouflage action recognition dataset for probing a system’s ability to recognize actions purely based on dynamic information. Notably, while our texture patterns are synthetic, they are animated by the motion of real-world actions [51].

Models and action recognition architectures A wide range of action recognition architectures have been developed [3]. Most contemporary architectures can be categorized as single stream 3D (x, y, t) convolutional, two-stream and attention based methods. Single stream approaches are motivated by the great success of 2D convolutional architectures on single image tasks [35,50,26,22]. In essence, the 3D architectures extend the same style of processing by adding temporal

support to their operations [27,55,8,56,13]. Two-stream architectures have roots in biological vision systems [25,18]. The idea of having separate pathways for **static** video content and **dynamic** information has been variously adopted. A key distinction in these designs is whether the pathway for dynamic information explicitly relies on optical flow estimation [49,15] or internally computed features that are thought to emulate flow like properties [14]. Attention-based approaches rely on various forms of non-local spatiotemporal data association as manifested in transformer architectures [60,12]. For evaluation of architectures on the new AFD, we select a representative sampling from each of these categories.

Interpretability Various efforts have addressed the representational abilities of video understanding architectures ranging from dynamic texture recognition [21], future frame selection [17] and comparing 3D convolutional vs. LSTM architectures [41]. Other work centering on action recognition focused on visualization of learned filters in convolutional architectures [16,61], or trying to remove scene biases for action recognition [39,38]. Evidence also suggests that optical flow is useful exactly because it contains invariances related to single frame appearance information [47]. More closely related to our study is work that categorized various action classes as especially dependent on temporal information; however, single frame static information still remains [48]. Somewhat similarly, an approach tried to tease apart the bias of various architectures and datasets towards static vs. dynamic information by manipulating videos through stylization; however, single frame static information remained a confounding factor [34]. While insights have been gained, no previous research has been able to completely disentangle the ability of these models to capitalize on single frame static information vs. dynamic information present across multiple frames. We concentrate on the ability of these systems to capture dynamic information, with an emphasis on action recognition.

1.3 Contributions

In light of previous work, we make the following three major contributions.

1. We introduce the Appearance Free Dataset (AFD) for video-based action recognition. AFD is a synthetic derivative of UCF101 [51] having *no static* appearance cues, however, revealing real-world motion as video.
2. We evaluate 11 contemporary architectures for action recognition on AFD; *i.e.* dynamic information alone. We conduct a psychophysical study with humans on AFD, and show significantly better performance of humans than networks. These results question the ability of the tested networks to use dynamic information effectively.
3. We provide a novel improved action recognition architecture with enhanced performance on AFD, while maintaining performance on standard input.

We make AFD, associated code and our novel architecture publicly available.

2 Appearance free dataset

Our proposed appearance free dataset, AFD101, is built from the original UCF101 [51] by employing a state-of-the-art optical flow estimator [53] to generate the corresponding framewise flow that is used to animate spatial noise patterns. The resulting AFD101, consisting of 13,320 videos, depicts realistic motion even while any individual frame is devoid of static appearance information that is relevant for action recognition. Figure 1 and the project webpage provide illustrative examples. We use UCF101 from many possible choices [3] as the basis for AFD because it was widely used in action recognition evaluation and is large enough to support training, yet small enough to facilitate numerous experiments.

2.1 Dataset generation methodology

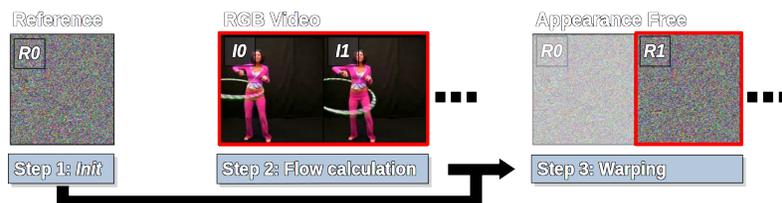


Fig. 2. Appearance free videos are created by 1) initializing noise R_0 , 2) calculating optical flow between input frame pairs, and 3) using the optical flow to warp R_0 . Steps 2 and 3 are repeated for each frame-pair in the input video.

Generation of AFD from RGB video follows three key steps; see Fig. 2.

1. Initialization: Generate a single frame of noise, with same spatial dimensions as the frames in the input video. The pixel values are sampled uniformly i.i.d. as three-channel RGB values. Let this frame be denoted as R_0 .
2. Flow calculation: Generate interframe optical flow for each temporally adjacent pair of input RGB frames, $I_{t-1}, I_t, t \in [1, T]$, yielding the flow $F_{t-1,t}$, with T being the number of frames in the input video. We use RAFT [53], for the extraction of optical flow.
3. Warping: We warp the initial noise frame R_0 using $F_{t-1,t}, t \in [1, T]$, to generate the next noise frame, R_t . The output frames, R_t , are an appearance free version of the input RGB video, where at each frame all that is seen is i.i.d. noise, but whose interframe motion corresponds to that of the original video; this output video has the same spatial and temporal dimension as the input video.

2.2 Implementation details

The details of our proposed methodology are particularly important, as all *static* discriminatory single frame information must be removed; to this end near-

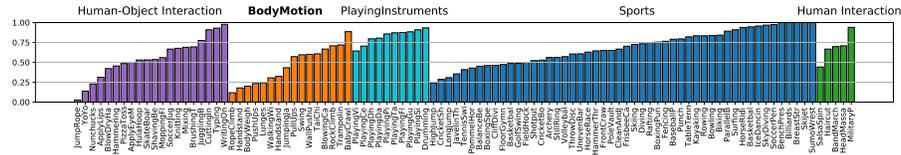


Fig. 3. ResNet50 Top-1 Accuracy on single frame UCF101 by *Action Groups*.

est neighbor interpolation during warping is performed. For each warp, newly occluded pixels are overwritten by the new value moved to that location; de-occluded pixels are filled with R_0 values to ensure that the spatial extent of the noise is constant. Along the border where the warping leads to undefined values (*e.g.* pixel at that location moves elsewhere and no new value is warped into that location) are treated as de-occlusions as well. In preliminary experiments, we found that other choices of implementation, *e.g.* bilinear interpolation as opposed to nearest neighbor interpolation, or always estimating flow with respect to I_0 always yielded inferior results and unwanted artifacts. In the worst cases these artifacts led to revealing motion boundaries of actors in single frames.

2.3 Is optical flow appearance free?

The concept of appearance consists of multiple intertwined qualities, including texture, contour and shape. While optical flow is indeed free of texture it is not free of contour and shape induced by motion of the camera or the actor itself; this fact can be observed qualitatively in Fig. 1: The flow visualizations alone reveal violin playing and knitting. We quantify this observation as follows. We train a strong single frame recognition architecture, ResNet50 [22], on single frames from three datasets: UCF101 (the original), UCF101Flow, comprised of optical flow frames generated from UCF101 by RAFT [53], and AFD101. Top-1 recognition results for UCF101, UCF101Flow and AFD101 are 65.6%, 29.4% and 1.1%, respectively. These results show that only AFD101 contains no discriminatory information in a single frame, as only 1.1% top-1 accuracy is in line with random guessing.

2.4 AFD5: A reduced subset suitable for small scale exploration

To evaluate the performance of state-of-the-art architectures in the context of human performance, a reduced subset of the classes in UCF101 is needed, simply because it is not feasible to conduct an experiment with humans with the original 101 categories. We chose to have five classes as it is possible for humans to hold five action classes in working memory for a prolonged duration [43]. In the following, we refer to the reduced size version of AFD101 as AFD5 and the corresponding reduction of UCF101, *i.e.* the RGB video, as UCF5.

A defining property of the subset is that it discourages recognition from static information in single frames even within the original RGB input. To this

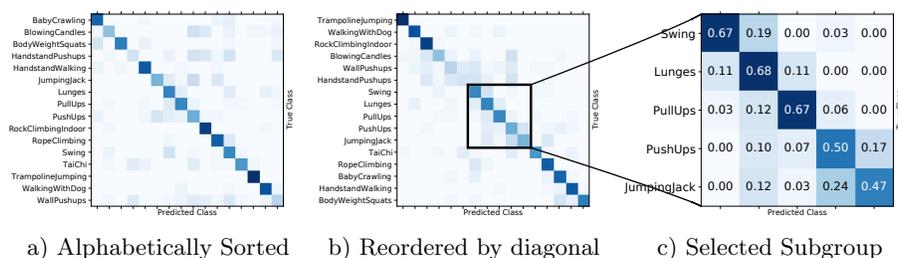


Fig. 4. Singleframe ResNet50: After finetuning on the *BodyMotion* action group of UCF101, we a) inspect the resulting confusion matrix b) reorder it by its diagonal weight so that classes that are often confused among each other get grouped together, and c) select the most prominent subgroup.

end we train a ResNet50 [22], used by many state-of-the-art architectures as inspiration for their backbone [3], on the entire UCF101 dataset in a single frame classification manner, *i.e.*, we randomly sample a frame from UCF101. By design, this procedure isolates the ability of any given class to support recognition on the basis of a single frame. UCF101 is subdivided into five action groups [51]: *Human Object Interaction*, *BodyMotion*, *PlayingInstruments*, *Sports* and *Human Human Interaction*; Fig. 3 has single frame recognition accuracy by group.

Our selection criterion is to choose videos with similar appearance; therefore, we select a subset of actions within a group. To select the action group we consider the following points. The *Body Motion* group is particularly attractive because it is the only category that has five or more classes with below 25% single frame recognition accuracy. The *Human Object Interaction* group is the next closest in terms of having several (albeit fewer than five) categories with such low single frame accuracy. However, its classes tend to be distinguished by featuring a prominent object and we do not want to promote categorization based on single frame object recognition. The *PlayingInstruments* group also features prominent objects and likewise is a poor choice. The remaining groups (*Sports* and *Human Human Interaction*) have generally higher single frame recognition accuracies; so, also are excluded. These considerations lead us to *Body Motion* as the group of interest for present concerns.

With the selection of *Body Motion*, we finetune the previously trained ResNet on that particular group. Given that finetuning, we perform confusion matrix reordering [54] and select the five classes that are most confused among each other in the group. More precisely we find the permutation matrix that has the highest interclass entropy, excluding classes with less than 50% accuracy; see Fig 4. This results in the choice of the following five classes for AFD5: *Swing*, *Lunges*, *PushUps*, *PullUps*, *JumpingJack*. AFD5 consists in total of 583 videos with approximately 116 videos per class.

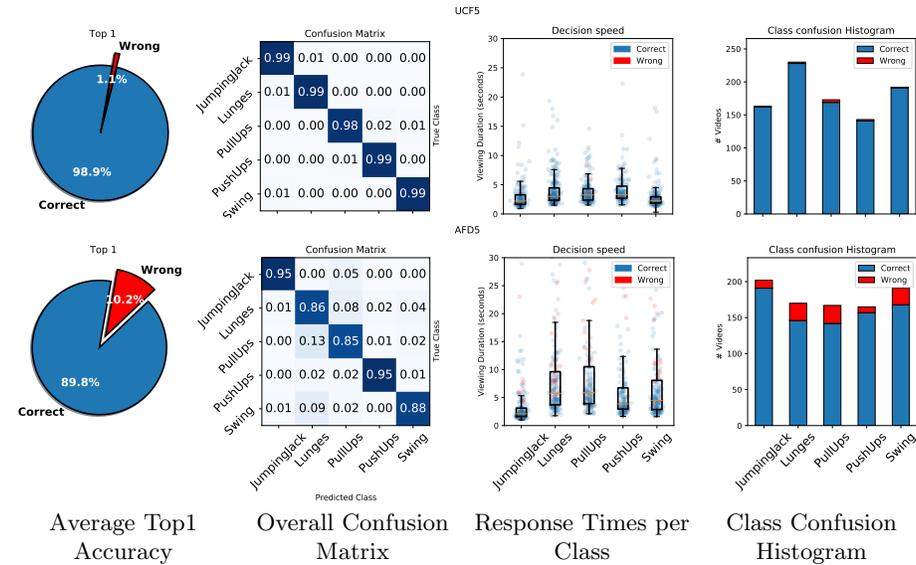


Fig. 5. Results of psychophysical study. Top row: Performance on standard RGB video, UCF5. Bottom row: Performance on appearance free data, AFD5.

3 Psychophysical study: Human performance on AFD5

To set a baseline for Appearance Free Data we perform a psychophysical study. The ability of humans to recognize actions from dynamic information alone, *i.e.* without any single frame static information, has been documented previously [28,10,57]. These studies are conducted with sparse dots, typically at major body joints, on otherwise blank displays. Our study appears to be unique in its use of dense noise patterns *i.e.* dense random dots.

3.1 Experiment design

A session in the experiment consisted of the following seven phases. (i) The participant is presented with a brief slide show explaining the task. The participant is told that they will see a series of videos as well as a menu on the same display from which to indicate their perception of a human action that is depicted. (ii) Training is provided on UCF5, *i.e.*, RGB video. During this phase, four training videos from each class, totaling 20 videos, are presented with correct responses indicated in the menu. The videos are shown in randomized order. (iii) Testing is performed on 90 UCF5 test videos. Each video repeats until the participant makes their selection on the menu. (iv) A rest period of five minutes is provided. (v) Training is provided on AFD5, *i.e.* appearance free video. During this phase, four training videos from each class are presented with correct responses indicated in the menu. The videos also are shown in randomized order.

(vi) Testing is performed on 90 AFD5 test videos. Each video repeats until the participant makes their selection on the menu. (vii) Participants are given a questionnaire to complete, with questions including visual impairments (all had normal or corrected vision), familiarity with computer vision, action recognition, as well as their impression of task difficulty. During both test phases, action category choices and response times were recorded. Ten human participants were recruited as volunteers to participate in the experiment. All experiments were conducted in the same environment, *i.e.* computer/display, lighting and seating distance from the display (0.5 meters). Participants were allowed to view freely from their seat and take as much time in making their responses as they like. All participants completed their session within approximately 30 minutes, including the questionnaire. Participants were informed and consented to the data gathered in written and signed manner, and were provided with sweets and thanks at task completion. The sample size of ten participants is in our view sufficient to show that various people with different backgrounds are *able* to solve AFD data without much trouble and to establish a baseline for comparison to action recognition architectures.

3.2 Human results

Results of the psychophysical study are shown in Fig. 5. On average, our participants perform with an accuracy of 89.8% on AFD5 and 98.9% on UCF5. Participants report that mistakes on UCF5, the RGB videos, resulted from accidental clicks when making selections. Our analysis of AFD5 samples that were most frequently misclassified by the participants reveals that those videos typically have sudden and high speed camera movement, *e.g.* as induced by quick panning and zooming, which in turn creates very large optical flow displacement. The plots of response time show that participants take longer to make their choice on AFD5, with the exception of *JumpingJack* which remains largely unaffected. The extended time taken for AFD5 suggests that appearance free recognition requires more effort and repeated playback of the video is necessary, even though ultimate performance is high.

Overall, the psychophysical results document the strong ability of humans to recognize actions on the basis of AFD, as well as standard RGB videos.

4 Computational study: Model performances on AFD

We evaluate 11 state-of-the-art action recognition architectures on UCF101 and our new AFD101. We select representative examples from the currently most widely used categories of action recognition architectures, *cf.* Sec. 1.2: single stream 3D convolutional [13,8,56], two-stream convolutional [49,14] and attention-based [12,60]. We also evaluate on a standard 2D convolutional architecture (ResNet50 [22]) to verify that our AFD does not support classification based on single frame static information. Notably, our results from the psychophysical study allow us to compare the performance of the network architectures to that of humans on both UCF5 and AFD5.

4.1 Training procedure

To guarantee a fair comparison all reported results are obtained by training and testing with our procedure, described below. Furthermore, since we want to investigate architectural benefits on Appearance Free Data, all results (except for one architecture, see below) are obtained without pretraining - concurrently often done on Kinetics400 [8] - as there is no AFD for Kinetics. Our goal is not to show how to achieve best performance on UCF101; it is to show the difference in performance of each network when presented with *only* dynamic information. We also ran preliminary experiments with pretraining on Kinetics and finetuning on AFD, but saw no significant improvement during AFD testing compared to no pretraining. Nevertheless, since all the evaluated architectures, save one, employ a ResNet-like backbone we use ImageNet pretrained versions of those to initialize our training. The exception is MViT, which needed to be initialized with Kinetics400 weights, as its performance was not increasing from baseline with the aforementioned scheme. However the Kinetics weights were only used as an initializer and we allowed for retraining of the entire backbone to make the results comparing UCF and AFD as comparable as possible. We acknowledge that our training strategy does not allow for the architectures to show state-of-the-art performance on RGB input in our experiments; however, our approach is necessary to allow for fair comparison of performance on RGB vs. AFD and measure the respective performance between the modalities.

4.2 Training details

We train with Adam [32] using an early stopping scheme, and a base learning rate of $3e^{-4}$. The batch size is chosen as large as possible on two GPUs with 24GB memory each. Data augmentation, a staple of modern deep learning, is kept the same for all evaluated models and datasets; uniform temporal sampling with jitter, random start frame selection, and contrast, hue, and brightness jitter are applied. Random horizontal flipping of entire videos is used, as actions in UCF101 do not depend on handedness. A quite aggressive data augmentation scheme that changes over the course of every 20 epochs to a less aggressive one, and capping at 80 epochs is used. We find that this aggressive data augmentation technique is needed to achieve reasonable results when no pretraining is used. Spatial center crops of 224×224 are used as the final input to all networks with the exception of X3D XS and S, which use a spatial resolution of 160×160 . The reported results are the averages of the three standard splits used on UCF101, now applied to both the original RGB as well as AFD. The results on UCF5 and AFD5 were obtained by finetuning the networks previously trained on their 101 class counterparts and using the same splits.

4.3 Architecture results

Table 1 shows top-1 accuracy results for all evaluated architectures on UCF101, AFD101, UCF5 and AFD5. Note that the results for the single frame ResNet on

UCF101 essentially are the same as those shown in Fig. 3. The results in Table 1 also validate our claim of having removed all *static*, single frame information relevant to recognition, as the performance is equal to random choice.

For most of the action recognition architectures, it is seen that obtaining $\approx 70\text{-}80\%$ top-1 accuracy is possible. This fact makes the sizable drop of performance on AFD data, the Δ columns in Table 1, especially striking. The average performance drops by approximately 30% almost evenly across all architectures - with two notable exceptions: The Fast stream of SlowFast and I3D OF a common I3D architecture with an optical flow estimator before the network input. These two architectures are noteworthy because their ability to maintain performance across regular and appearance free video apparently stems from their design to prioritize temporal information over spatial; although, this focus on temporal information leads them to be less competitive on the standard RGB input in comparison to the other evaluated architectures. The comparison to human performance shows X3D variations to be the best architectures, competitive with human performance on UCF5; *however*, the best architectures on AFD5 (X3D M and SlowFast) are *considerably* below human top-1 accuracy (18% below). These results show that there is room for improvement by enhancing the ability of current architectures to exploit dynamic information for action recognition. It also suggests that optical flow is *not* what is learned by these networks, when no explicit representation of flow is enforced. Importantly, TwoStream [49] and I3D [8] OF were evaluated with RAFT optical to allow for a fair comparison with each other and with our new algorithm, introduced in the next section. The summary plots below Table 1 show learnable architecture parameter counts and ordering by AFD accuracy. These plots further emphasize the fact that explicit representation of dynamic information (*e.g.* as in I3D OF and to some extent Fast) is best at closing the gap between performance on RGB and AFD. X3D M is among the top performing architectures on RGB; further, among those top RGB performers, it shows the smallest Δ . Also, it has a relatively small number of learnable parameters, with only Fast notably smaller. These observations motivate the novel architecture that we present in the next section. It improves performance on AFD, while also slightly improving performance on RGB. This architecture is given as E2S-X3D in Table 1. Moreover, it performs on par with humans on UCF5 and AFD5.

5 Two-stream strikes back

The results of our evaluation of architectures in Sec. 4 show the importance of explicit representation of dynamic information for strong performance on AFD. This result motivates us to revive the two-stream appearance plus optical flow design [49]. While state of the art when introduced, compared to more recent action recognition architectures, *e.g.* those evaluated in Sec. 4, it is no longer competitive [63]; so, we incorporate two-streams into the top performing architecture on RGB input with the smallest drop on AFD, *i.e.* X3D.

	$f \times r$	UCF101	AFD101	Δ	UCF5	AFD5	Δ
Single Image Input							
ResNet50 [22]	1×1	65.6	1.1	64.5	74.2	20.1	54.1
Video Input							
I3D [8]	8×8	67.8	43.3	24.5	95.6	55.1	40.5
TwoStream [49]	$\binom{1 \times 1}{10 \times 1}$	76.1	29.3	46.8	78.1	62.3	15.8
C2D [60]	8×8	77.4	42.6	34.8	84.5	69.0	15.5
R2+1D [56]	16×4	68.2	28.9	39.3	80.6	67.9	12.7
Slow[14]	8×8	73.3	37.6	35.7	89.9	40.6	49.3
Fast [14]	32×2	50.5	45.9	4.6	71.4	65.7	5.7
SlowFast [14] $\binom{\text{Slow}}{\text{Fast}}$	$\binom{8 \times 8}{32 \times 2}$	82.9	55.4	27.5	91.0	70.2	20.8
I3D[8] OF	8×8	54.5	44.4	10.1	74.4	59.7	14.7
X3D XS [13]	4×12	79.8	43.2	36.6	95.2	57.0	38.2
X3D S [13]	13×6	80.8	56.9	23.9	97.3	69.8	27.5
X3D M [13]	16×5	81.5	58.2	23.3	98.8	71.5	27.3
MViT [12]	16×4	82.9	39.8	43.1	95.1	22.2	72.9
E2S-X3D Ours $\binom{M}{S}$	$\binom{16 \times 5}{13 \times 6}$	85.7	73.9	11.8	99.4	90.8	8.6
Human Average					98.9	89.8	9.1

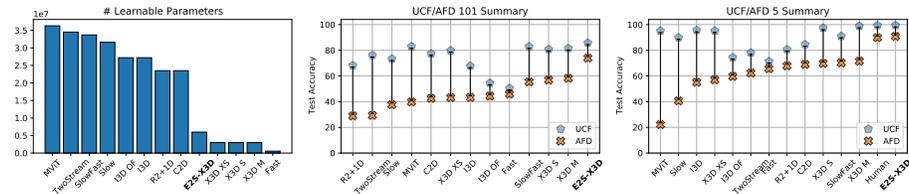


Table 1. Top-1 accuracy for the evaluated action recognition architectures. The second column indicates the number of frames and the temporal sampling rate, $f \times r$, of each network. For two-stream approaches the notation $\binom{\text{Appearance}}{\text{Motion}}(f \times r)$ is used. A drop in performance across all networks on AFD data is observed. Plots below the table show parameter counts as well as summarize the accuracy findings, sorted by AFD performance. Absolute performance on UCF is of secondary importance; it merely gives a point of comparison with respect to AFD. The Δ columns shows the difference between UCF and AFD performance, which highlights the relative *dynamic* recognition capability of architectures.

5.1 E2S-X3D: Design of a novel action recognition architecture

The design of our novel architecture, ES2-X3D, is shown in Fig. 6. Initially, optical flow fields and RGB images are processed in separate 3D convolutional streams. Both streams are versions of the X3D architecture [13], where we use the S variant for optical flow and the M variant for RGB. These two architectures differ in their spatiotemporal resolution of the input and as a consequence the activation maps. X3D M uses larger spatial and temporal extents, whereas S has smaller spatial and temporal supports. Since the input flow field already

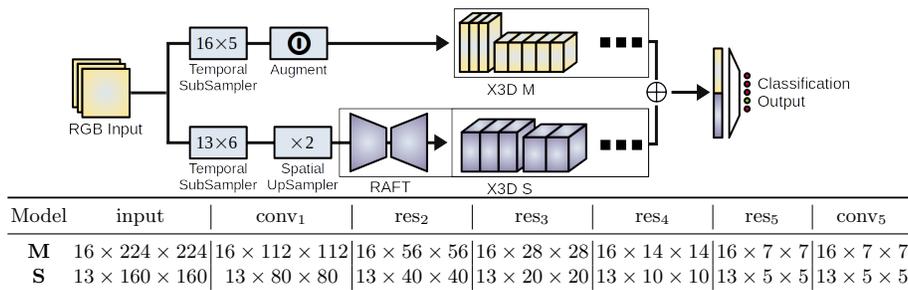


Fig. 6. Our novel ES2-X3D operates with two parallel streams for processing of RGB (top stream) and optical flow (bottom stream). Late fusion via concatenation is followed by a fully connected and Softmax layer. The table shows the output sizes of each employed X3D architecture in the format $T \times H \times W$.

represents the dynamic information with detail, the additional spatiotemporal extent is not needed; in contrast, the RGB stream without the optical flow benefits from larger spatial and temporal support, see Fig. 6 (bottom). We validate these choices below via ablation. The dimensionality of the layers is shown in Fig. 6. Following the separate parallel processing, late fusion is performed. We fuse the outputs via simple concatenation of the outputs of the two streams. While more sophisticated fusion schemes might be considered, *e.g.* [15], we leave that for future work. Finally, a fully connected layer followed by a Softmax layer yields the classification output, which is used with a cross-entropy loss to train the network. For optical flow input, we use a state-of-the-art optical flow estimator, RAFT [53] pretrained on Sintel [7]. We do not further train the flow extractor and only use it in its inference mode. To capture fine grained motion and to adjust for the different spatial resolution RAFT was trained with, we add an up-sampling layer to the RGB images prior to RAFT. Data augmentation, as described in Sec. 4, is only used in the appearance stream, as indicated in Fig. 6. The training, validation and testing of the network follows the same procedure as the other networks; described in Sec. 4.

UCF5					AFD5				
Architecture	Input				Architecture	Input			
	$\binom{RGB}{-}$	$\binom{-}{RAFT}$	$\binom{RGB}{RAFT}$			$\binom{RGB}{-}$	$\binom{-}{RAFT}$	$\binom{RGB}{RAFT}$	
XS	4 × 12	95.2	82.1	—	XS	4 × 12	57.0	77.3	—
S	13 × 6	97.3	89.5	—	S	13 × 6	69.8	80.3	—
M	16 × 5	98.8	86.8	—	M	16 × 5	71.5	78.9	—
$\binom{M}{S}$	$\binom{16 \times 5}{13 \times 6}$	—	—	99.4	$\binom{M}{S}$	$\binom{16 \times 5}{13 \times 6}$	—	—	90.8

Table 2. Performance of X3D architecture configurations, XS, S and M for UCF5 and AFD5. Ablation is performed for both optical flow and RGB inputs.

5.2 E2S-X3D: Empirical evaluation

E2S-X3D outperforms all competing architectures on AFD101, with the closest competitor (singlestream X3D) trailing by 15.7%; see Table 1. It also improves performance on UCF101, albeit slightly; 2.8% points gained on MViT (MViT has around $6\times$ more parameters). To rule out an ensembling effect, we compared our network in its default configuration (RGB in one stream, optical flow in the other) with an alternative where it inputs RGB to both streams. We find a performance drop of 10.6% on UCF101 and 30.7% on AFD101, validating our hypothesis that explicit optical flow is crucial. Moreover, it is on par with human performance on both AFD5 and UCF5. These results document the advantage of explicit modeling of motion for action recognition. Table 2 shows results of ablations across various X3D configurations for both AFD5 and UCF5. It shows that for both input modalities, best performance is achieved when the S configuration is used for optical flow and the M configuration for the RGB stream. This set of experiments empirically validates our final E2S-X3D design.

6 Conclusions

We introduce an extension to a widely used action recognition dataset that disentangles static and dynamic video components. In particular, no single frame contains any *static* discriminatory information in our appearance free dataset (AFD); the action is only encoded in the temporal dimension. We show, by means of a psychophysical study with humans, that solving AFD is possible with a high degree of certainty ($\sim 90\%$ Top1 Accuracy). This result is especially interesting as 11 current state-of-the-art action recognition architectures show much weaker performance as well as a steep drop in performance when comparing standard RGB to AFD input. These results lend to the interpretability of the evaluated architectures, by documenting their ability to exploit dynamic information. In particular, this shows that optical flow, or a similarly descriptive representation, is not an emergent property of *any* of the tested network architectures.

We propose a novel architecture incorporating the explicit computation of optical flow and use insights from recent action recognition research. This explicit form of modeling *dynamics* allows our approach to outperform all competing methods, and compete with human level performance. Given the strong performance of our architecture in our evaluation, future work could consider training and investigation of performance on larger datasets [8,19] as well as application to other tasks where non-synthetic data is available and dynamic information plays a crucial role, *e.g.* camouflaged animal detection [37]. Alternative fusion strategies are also a potential for improvement of our present architecture.

References

1. Laying down harmonised rules on artificial intelligence (artificial intelligence act) and amending certain union legislative acts. European Commission (2021)
2. Regulating AI: Critical issues and choices. Law Council of Ontario (2021)
3. Aafaq, N., Mian, A., Liu, W., Gilani, S.Z., Shah, M.: Video description: A survey of methods, datasets, and evaluation metrics. *ACM Computing Surveys* **52**(6), 1–37 (2019)
4. Baker, S., Scharstein, D., Lewis, J., Roth, S., Black, M.J., Szeliski, R.: A database and evaluation methodology for optical flow. *International Journal of Computer Vision* **92**(1), 1–31 (2011)
5. Bideau, P., Learned-Miller, E.: It’s moving! A probabilistic model for causal motion segmentation in moving camera videos. In: *Proceedings of the European Conference on Computer Vision* (2016)
6. Braddick, O.J.: Low-level and high-level processes in apparent motion. *Philosophical Transactions of the Royal Society of London. B, Biological Sciences* **290**(1038), 137–151 (1980)
7. Butler, D.J., Wulff, J., Stanley, G.B., Black, M.J.: A naturalistic open source movie for optical flow evaluation. In: *Proceedings of the European Conference on Computer Vision* (2012)
8. Carreira, J., Zisserman, A.: Quo vadis, action recognition? A new model and the kinetics dataset. In: *Proceedings of the Conference on Computer Vision and Pattern Recognition* (2017)
9. Choi, J., Gao, C., Messou, C.E.J., Huang, J.B.: Why can’t I dance in the mall? Learning to mitigate scene bias in action recognition. In: *Proceedings of the Conference on Advances in Neural Information Processing Systems* (2019)
10. Dittrich, W.H.: Action categories and the perception of biological motion. *Perception* **22**(1), 15–22 (1993)
11. 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 International Conference on Computer Vision* (2015)
12. Fan, H., Xiong, B., Mangalam, K., Li, Y., Yan, Z., Malik, J., Feichtenhofer, C.: Multiscale vision transformers. In: *Proceedings of the International Conference on Computer Vision* (2021)
13. Feichtenhofer, C.: X3D: Expanding architectures for efficient video recognition. In: *Proceedings of the Conference on Computer Vision and Pattern Recognition* (2020)
14. Feichtenhofer, C., Fan, H., Malik, J., He, K.: SlowFast networks for video recognition. In: *Proceedings of the International Conference on Computer Vision* (2019)
15. Feichtenhofer, C., Pinz, A., Wildes, R.P.: Spatiotemporal multiplier networks for video action recognition. In: *Proceedings of the Conference on Computer Vision and Pattern Recognition* (2017)
16. Feichtenhofer, C., Pinz, A., Wildes, R.P., Zisserman, A.: Deep insights into convolutional networks for video recognition. *International Journal of Computer Vision* **128**(2), 420–437 (2020)
17. Ghodrati, A., Gavves, E., Snoek, C.G.M.: Video time: Properties, encoders and evaluation. In: *British Machine Vision Conference* (2018)
18. Goodale, M.A., Milner, A.D.: Separate visual pathways for perception and action. *Trends in Neurosciences* **15**(1), 20–25 (1992)

19. Goyal, R., Ebrahimi Kahou, S., Michalski, V., Materzynska, J., Westphal, S., Kim, H., Haenel, V., Fruend, I., Yianilos, P., Mueller-Freitag, M., Hoppe, F., Thureau, C., Bax, I., Memisevic, R.: The “something something” video database for learning and evaluating visual common sense. In: Proceedings of the International Conference on Computer Vision (2017)
20. Gu, C., Sun, C., Ross, D.A., Vondrick, C., Pantofaru, C., Li, Y., Vijayanarasimhan, S., Toderici, G., Ricco, S., Sukthankar, R., Schmid, C., Malik, J.: AVA: A video dataset of spatio-temporally localized atomic visual actions. In: Proceedings of the Conference on Computer Vision and Pattern Recognition (2018)
21. Hadji, I., Wildes, R.P.: A new large scale dynamic texture dataset with application to convnet understanding. In: Proceedings of the European Conference on Computer Vision (2018)
22. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the Conference on Computer Vision and Pattern Recognition (2016)
23. He, Y., Shirakabe, S., Satoh, Y., Kataoka, H.: Human action recognition without human. In: Proceedings of the European Conference on Computer Vision (2016)
24. Hiley, L., Preece, A., Hicks, Y.: Explainable deep learning for video recognition tasks: A framework & recommendations. arXiv preprint arXiv:1909.05667 (2019)
25. Hubel, D.H., Wiesel, T.N.: Receptive fields of single neurones in the cat’s striate cortex. *The Journal of Physiology* **148**(3), 574–591 (1959)
26. Iandola, F.N., Han, S., Moskewicz, M.W., Ashraf, K., Dally, W.J., Keutzer, K.: SqueezeNet: AlexNet-level accuracy with 50x fewer parameters and < 0.5 mb model size. arXiv preprint arXiv:1602.07360 (2016)
27. Ji, S., Xu, W., Yang, M., Yu, K.: 3D convolutional neural networks for human action recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **35**(1), 221–231 (2012)
28. Johansson, G.: Visual perception of biological motion and a model for its analysis. *Perception & Psychophysics* **14**(2), 201–211 (1973)
29. Julesz, B.: *Foundations of Cyclopean Perception*. U. Chicago Press (1971)
30. Kang, S.M., Wildes, R.P.: Review of action recognition and detection methods. arXiv preprint arXiv:1610.06906 (2016)
31. Karpathy, A., Toderici, G., Shetty, S., Leung, T., Sukthankar, R., Fei-Fei, L.: Large-scale video classification with convolutional neural networks. In: Proceedings of the Conference on Computer Vision and Pattern Recognition (2014)
32. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
33. Kong, Y., Jia, Y., Fu, Y.: Interactive phrases: Semantic descriptions for human interaction recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **36**(9), 1775–1788 (2014)
34. Kowal, M., Siam, M., Islam, A., Bruce, N.D.B., Wildes, R.P., Derpanis, K.G.: A deeper dive into what deep spatiotemporal networks encode: Quantifying static vs. dynamic information. In: Proceedings of the Conference on Computer Vision and Pattern Recognition (2022)
35. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. *Proceedings of the Advances in Neural Information Processing Systems* (2012)
36. Kuehne, H., Jhuang, H., Garrote, E., Poggio, T., Serre, T.: HMDB: A large video database for human motion recognition. In: Proceedings of the International Conference on Computer Vision (2011)

37. Lamdouar, H., Yang, C., Xie, W., Zisserman, A.: Betrayed by motion: Camouflaged object discovery via motion segmentation. In: Proceedings of the Asian Conference on Computer Vision (2020)
38. Li, Y., Li, Y., Vasconcelos, N.: Diving48 dataset. <http://www.svc1.ucsd.edu/projects/resound/dataset.html>
39. Li, Y., Li, Y., Vasconcelos, N.: Resound: Towards action recognition without representation bias. In: Proceedings of the European Conference on Computer Vision (2018)
40. Mahmood, N., Ghorbani, N., Troje, N., Pons-Moll, G., Black, M.: AMASS: Archive of Motion Capture as Surface Shapes. In: Proceedings of the International Conference on Computer Vision (2019)
41. Manttari, J., Broomé, S., Folkesson, J., Kjellstrom, H.: Interpreting video features: A comparison of 3D convolutional networks and convolutional LSTM networks. In: Proceedings of the Asian Conference on Computer Vision (2020)
42. Mayer, N., Ilg, E., Hausser, P., Fischer, P., Cremers, D., Dosovitskiy, A., Brox, T.: A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In: Proceedings of the Conference on Computer Vision and Pattern Recognition (2016)
43. Miller, G.A.: The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological review* **63**(2), 81 (1956)
44. Nishida, S., Kawabe, T., Sawayama, M., Fukiage, T.: Motion perception: From detection to interpretation. *Annual review of Vision Science* **4**, 501–523 (2018)
45. Richter, S.R., Hayder, Z., Koltun, V.: Playing for benchmarks. In: Proceedings of the International Conference on Computer Vision (2017)
46. Ros, G., Sellart, L., Materzynska, J., Vazquez, D., Lopez, A.M.: The SYNTHIA dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In: Proceedings of the Conference on Computer Vision and Pattern Recognition (2016)
47. Sevilla-Lara, L., Liao, Y., Güney, F., Jampani, V., Geiger, A., Black, M.J.: On the integration of optical flow and action recognition. In: Proceedings of the German Conference on Pattern Recognition. Springer (2018)
48. Sevilla-Lara, L., Zha, S., Yan, Z., Goswami, V., Feiszli, M., Torresani, L.: Only time can tell: Discovering temporal data for temporal modeling. In: Proceedings of the Winter Conference on Applications of Computer Vision (2021)
49. Simonyan, K., Zisserman, A.: Two-stream convolutional networks for action recognition in videos. In: Proceedings of the Conference on Advances in Neural Information Processing Systems (2014)
50. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
51. Soomro, K., Zamir, A.R., Shah, M.: UCF101: A dataset of 101 human actions classes from videos in the wild. arXiv preprint arXiv:1212.0402 (2012)
52. Sriastava, N., Manisomov, E., Salakhutdinov, R.: Unsupervised learning of video representations using LSTMs. In: Proceedings of the International Conference on Machine Learning (2015)
53. Teed, Z., Deng, J.: RAFT: Recurrent all-pairs field transforms for optical flow. In: Proceedings of the European Conference on Computer Vision (2020)
54. Thoma, M.: Analysis and Optimization of Convolutional Neural Network Architectures. Master’s thesis, University of the State of Baden-Wuerttemberg (2017)
55. Tran, D., Bourdev, L., Fergus, R., Torresani, L., Paluri, M.: Learning spatiotemporal features with 3D convolutional networks. In: Proceedings of the International Conference on Computer Vision (2015)

56. Tran, D., Wang, H., Torresani, L., Ray, J., LeCun, Y., Paluri, M.: A closer look at spatiotemporal convolutions for action recognition. In: Proceedings of the Conference on Computer Vision and Pattern Recognition (2018)
57. Troje, N.F.: Decomposing biological motion: A framework for analysis and synthesis of human gait patterns. *Journal of Vision* **2**(5), 2–2 (2002)
58. Ullman, S.: *The Interpretation of Visual Motion*. MIT Press (1979)
59. Vu, T.H., Olsson, C., Laptev, I., Oliva, A., Sivic, J.: Predicting Actions from Static scenes. In: Proceedings of the European Conference on Computer Vision (2014)
60. Wang, X., Girshick, R., Gupta, A., He, K.: Non-local Neural Networks. In: Proceedings of the Conference on Computer Vision and Pattern Recognition (2018)
61. Zhao, H., Wildes, R.P.: Interpretable deep feature propagation for early action recognition. arXiv preprint arXiv:2107.05122 (2021)
62. Zhou, B., Tang, X., Wang, X.: Coherent filtering: Detecting coherent motions from crowd clutters. In: Proceedings of the European Conference on Computer Vision (2012)
63. Zhu, Y., Li, X., Liu, C., Zolfaghari, M., Xiong, Y., Wu, C., Zhang, Z., Tighe, J., Manmatha, R., Li, M.: A comprehensive study of deep video action recognition. arXiv preprint arXiv:2012.06567 (2020)