

This ACCV 2020 paper, provided here by the Computer Vision Foundation, is the author-created version. The content of this paper is identical to the content of the officially published ACCV 2020 LNCS version of the paper as available on SpringerLink: https://link.springer.com/conference/accv

# Play Fair: Frame Attributions in Video Models

Will Price<sup>[0000-0003-2884-0290]</sup> and Dima Damen<sup>[0000-0001-8804-6238]</sup>

<first>.<last>@bristol.ac.uk, University of Bristol, UK

Abstract. In this paper, we introduce an attribution method for explaining action recognition models. Such models fuse information from multiple frames within a video, through score aggregation or relational reasoning. We break down a model's class score into the sum of contributions from each frame, *fairly*. Our method adapts an axiomatic solution to fair reward distribution in cooperative games, known as the Shapley value, for elements in a variable-length sequence, which we call the Element Shapley Value (ESV). Critically, we propose a tractable approximation of ESV that scales linearly with the number of frames in the sequence. We employ ESV to explain two action recognition models (TRN and TSN) on the fine-grained dataset Something-Something. We offer detailed analysis of supporting/distracting frames, and the relationships of ESVs to the frame's position, class prediction, and sequence length. We compare ESV to naive baselines and two commonly used feature attribution methods: Grad-CAM and Integrated-Gradients.

# 1 Introduction

Progress in Action Recognition has seen remarkable gains in recent years thanks to architectural advances [1–7] and large-scale datasets [8–13]. The task of action recognition is to classify the action depicted in a video, from a sequence of frames. We address a question previously unexplored in action recognition: given a video and a trained model, how much did each frame contribute to the model's output?

Determining the contribution of each frame is similar to feature attribution, which assigns a value to each input feature, representing its weight to a scalar output in the model's response. Note that this is distinct from feature selection which computes the value of the feature globally across a dataset, and not for a specific example. Previous works in feature attribution [14–18] have taken inspiration from cooperative game theory The Shapley value [19] was proposed as an axiomatic solution to *fairly* distributing a reward amongst the players in a cooperative game based on each player's contribution. We offer the first attempt to integrate the Shapley axioms of fairness for video understanding—hence our paper's title 'Play Fair'.

In contrast to general feature attribution, and similar approaches applied to images [20–25], our focus is on attributing elements (frames) in a time-varying sequence (video). We depict our approach, *Element Shapley Value* (ESV), in Fig. 1. ESV assigns an attribution to each frame representing its contribution to the model output. This attribution is computed from the change in the model



**Fig. 1.** We assign  $\phi_i^c$ , the Element Shapley Value, to each frame *i* in a video *fairly*, given a model *f* and a specific class *c*. We calculate differences in the class score for subsequences with and without the specified frame—e.g. (1,2,4) vs (1,4), and combine these for all subsequences (tiled boxes) to produce the element's attribution (e.g. for frame 2). Class-specific  $\phi_i^c$  are shown for two classes (right), the ground-truth class (in blue) and another class (in orange), highlighting positive/negative frame attributions.

output as the frame is added to different subsequences of the input video. We compute *class-specific* attributions (Fig. 1 right). Additionally, as attributions produced by ESV can be combined, we introduce *class-contrastive* attributions to determine which frames contribute more to one class than another.

Our contributions can be summarised as follows:

- 1. We introduce Element Shapley Value (ESV) to determine the contribution of each element in a variable-length sequence to a model's output<sup>1</sup>.
- 2. We propose and evaluate a tractable approximation to computing ESVs that linearly scales with the number of elements in the sequence.
- 3. We calculate class-specific and class-contrastive ESVs for two action recognition models on the fine-grained dataset Something-Something.
- 4. We demonstrate that some frames have a negative impact on the model's output. This deviates from the common misconception that utilising all frames in a sequence always improves performance.

We next explain our approach to determining the contribution of each element in a sequence to a model's response, which we term *element attribution*. Related work is deferred until the end of the paper, to offer the best insight into how our work relates to feature attribution and approaches in video explanation.

# 2 Element Attribution in Sequences

In this section, we introduce the Shapley value as an axiomatic solution to attribution. We discuss the limitations of feature attribution methods, showcasing how these can be overcome when attributing elements in sequences. We then introduce our approach to tractably computing ESVs for variable-length sequences.

### 2.1 Element Attribution and the Shapley Value

To answer our motivational question, how much did each frame contribute to the model's output?, we introduce element attribution as the task of determining

<sup>&</sup>lt;sup>1</sup> We focus on individual frames, but our approach naturally extends to determining attributions of video clips fed to a 3D CNN (e.g. [26]). As these models average clip scores, ESV can explain how much did each *clip* contribute to the model's output.

the contribution  $\phi_i$  of an element  $x_i$  in a sequence  $X = (x_i)_{i=1}^n$  to the output of the model f(X). For a classification model, where the output is a vector of class scores  $f(X) \in \mathbb{R}^C$ , we obtain the contribution  $\phi_i^c$  for each element  $x_i$  to the score  $f_c(X)$  of any class c.

Element attribution can be viewed as special case of grouped-feature attribution, where an element is represented by a group of features and elements are ordered by their observation time. Of the previously-proposed feature attribution methods, *additive* ones [16] have a desirable trait where the sum of attributions equals the model output, up to a constant additive bias  $b_c$  specific to f:

$$f_c(X) = b_c + \sum_{x_i \in X} \phi_i^c.$$
(1)

Dividing the model's output should be done *fairly*, reflecting the element's actual contribution to the model's output. A natural way of measuring the contribution of an element  $x_i$ , is to consider the change in the model's response when it is added<sup>2</sup> to a subsequence  $X' \subset X$ , where X' does not contain  $x_i$ :

$$\Delta_i^c(X') = f_c(X' \cup \{x_i\}) - f_c(X'), \tag{2}$$

this is known as the marginal contribution of  $x_i$  on X'. Note when  $\Delta_i^c(X')$  is positive, the addition of  $x_i$  increases the response of the model for class c and when negative, decreases it. However,  $\Delta_i^c(X')$  will differ depending on the subsequence X'. As such, it is necessary to consider the marginal contribution of  $x_i$  on all subsequences it can join to calculate element attributions. This has the benefit of capturing the effects of resampling the sequence at all possible rates.

Lloyd Shapley (1953) defined a set of axioms that  $\phi_i^c$  should satisfy if it is to be considered *fair* [19] (later refined by Young [27] in 1985), we explain these in the context of element attribution:

A1: Efficiency. The sum of element attributions should equal the output of the model on the input X, minus the output of the model on the empty sequence  $\emptyset$ :

$$\sum_{x_i \in X} \phi_i^c = f_c(X) - f_c(\emptyset).$$
(3)

We call this difference,  $f'_c(X) = f_c(X) - f_c(\emptyset)$ , the evidential score of X w.r.t to class c, as it is the difference in class score when X is observed over observing nothing. We use the empirical class distribution of the training set as  $f_c(\emptyset)$ .

**A2: Symmetry.** Any pair of elements  $x_i$  and  $x_j$  should be assigned equal contributions,  $\phi_i^c = \phi_j^c$ , if  $f_c(X' \cup \{x_i\}) = f_c(X' \cup \{x_j\})$  for all  $X' \subseteq X \setminus \{x_i, x_j\}$ . **A3: Monotonicty.** For any pair of classes  $c_1, c_2$ , if  $\Delta_i^{c_1}(X') \ge \Delta_i^{c_2}(X')$  for all subsequences  $X' \subseteq X \setminus \{x_i\}$  then  $\phi_i^{c_1} \ge \phi_i^{c_2}$ .

There is a unique formulation of  $\phi_i^c$  satisfying the above axioms, known as the Shapley value, which consequently assigns attributions *fairly*. The Shapley

<sup>&</sup>lt;sup>2</sup> We adapt set notation to sequences using subset ( $\subset$ ) and union ( $\cup$ ) operations to form new subsequences. Elements preserve the same ordering as in the full sequence.

#### 4 W. Price and D. Damen

value [19] was originally proposed as a solution to the problem of fairly distributing a reward amongst players in a cooperative game. A cooperative game is one where a group of players known as a *coalition* collaborate to earn some reward. Players are treated as sets, as opposed to sequences. By analogy, each element in the sequence acts as a player and subsequences act as ordered coalitions. Accordingly, we refer to our element attribution method, that satisfies the Shapley axioms, as the Element Shapley Value (ESV). This is the weighted average of the marginal contributions of  $x_i$  to all subsequences  $X' \subseteq X \setminus \{x_i\}$  it can join:

$$\phi_i^c = \sum_{X' \subseteq X \setminus \{x_i\}} w(X') \Delta_i^c(X') \qquad w(X') = \frac{(|X| - |X'| - 1)! |X'|!}{|X|!} \,. \tag{4}$$

where w(X') can be interpreted as the probability of drawing a subsequence X'from  $X \setminus \{x_i\}$ , considering each way of forming X by adding an element at a time starting from  $\emptyset$ , going via X' and  $X' \cup \{x_i\}$ , as equally likely. Consequently,  $\phi_i^c$  can be rewritten as the expectation over a random variable X' with sample space  $2^{X \setminus \{x_i\}}$  and probability mass function w(X') (proof in Appendix B.1):

$$\phi_i^c = \mathop{\mathbb{E}}_{X'} \left[ \Delta_i^c(X') \right],\tag{5}$$

which can be reformulated again in terms of the expected model response on the sequences before and after  $x_i$  is added to X':

$$\phi_i^c = \mathop{\mathbb{E}}_{X'} \left[ f_c(X' \cup \{x_i\}) \right] - \mathop{\mathbb{E}}_{X'} \left[ f_c(X') \right] \,. \tag{6}$$

This is an easier form to compute, as we don't have to measure the marginal contributions of each element on each subsequence, but instead measure the expected model score on subsequences containing  $x_i$  and those not containing  $x_i$ . We take one final step towards the form that we will actually use to compute element attributions, expanding Eq. (6) via the law of total expectation to

$$\phi_i^c = \mathbb{E}_s \left[ \mathbb{E}_{X'|s} \left[ f_c(X' \cup \{x_i\}) \right] - \mathbb{E}_{X'|s} \left[ f_c(X') \right] \right], \tag{7}$$

where the subsequence length s ranges from 0 to |X| - 1 with equal probability, and the random variable X' conditioned on s takes on values from  $\{X' \subseteq X \setminus \{x_i\} : |X'| = s\}$  with equal probability (proof in Appendix B.1). This allows us to compute  $\phi_i^c$  by first considering the expected marginal contributions at each scale s and then combining these across scales.

#### 2.2 Element Attribution in Variable-Length Sequences

As shown in Eq. (4), computing element attributions requires evaluating the model on different subsequences. We first consider approaches used to evaluate a model on feature subsets, then the options we have in element attribution.

**Feature Ablation.** Evaluating a model on feature subsets is challenging as models rarely support the notion of a 'missing feature'. Two approaches exist: re-training the model on all combinations of features [28] or substituting missing features with those from a reference [14,16,22,24,25,29,30], but both approaches have limitations. Retraining is computationally infeasible for more than a handful of features, and the choice of reference in feature substitution has a significant impact on the resulting attribution values [18,31].

**Element Ablation.** Training an exponential number of models for each subsequence is infeasible for large models and thus ruled out. We could substitute frames by a *reference*, e.g. the mean element from the training set, or the nearest element still present, amongst other reasonable alternatives. However, similar to feature substitution [31], the choice of reference will have a large impact on the resulting attribution values. Feeding sequences that are out-of-distribution (such as one created by duplicating frames) to an action recognition model can result in uncharacteristic model responses [32]. We instead propose a more attractive option: utilising models that support variable-length input sequences.

**Supporting Variable Length Inputs.** When the model to be explained does not support variable-length inputs, we take inspiration from multi-scale models (e.g. TRN [1]) to build a model  $f^{\text{ms}}$  capable of operating over variable-length sequences. Let  $f^s(X)$  be a fixed-scale model which takes a sequence of length s as input. We construct a set of models operating at different scales  $\{f^s\}_{s=1}^{n_{\max}}$ , up to a maximum subsequence length  $n_{\max}$ . We then combine these, such that each scale contributes equally to the overall output:

$$f^{\rm ms}(X) = \mathop{\mathbb{E}}\limits_{s} \left[ \mathop{\mathbb{E}}\limits_{X'|s} \left[ f^s(X') \right] \right], \tag{8}$$

where s is a random variable over subsequence lengths  $\{1 \dots \min(|X|, n_{\max})\}$ , all equally likely, and X' is a random variable over subsequences of size s. This has similarities to the retraining approach [28], however we can leverage the homogeneity of our input to reduce the number of models from  $\mathcal{O}(2^{|X|})$  (one for each possible subsequence) to  $\mathcal{O}(n_{\max})$ . Another contrast is that the same models are used in inference and element attribution, unlike the work in [28].

### 2.3 Tractable Approach to Element Shapley Values

We have now defined the Element Shapley Value (ESV), as well as how these can be calculated using models supporting variable-length inputs. We next show how we can tractability compute the ESVs for all elements in a variable-length sequence in parallel culminating in Algorithm 1.

**Bottom-up Computation.** We compute the model's output for each subsequence once by deriving a recurrence relation between  $f^{\text{ms}}(X')$  and  $f^{\text{ms}}(X)$ , where  $X' \subset X$ , |X| = n, and |X'| = n - 1 (proof in Appendix B.2),

$$f^{\rm ms}(X) = \begin{cases} f^1(X) & n = 1\\ \frac{1}{n} \left[ f^n(X) + (n-1) \mathbb{E}_{X'|s=n-1} \left[ f^{\rm ms}(X') \right] \right] & n \le n_{\rm max} \\ \mathbb{E}_{X'|s=n-1} \left[ f^{\rm ms}(X') \right], & n > n_{\rm max} \end{cases}$$
(9)

### 6 W. Price and D. Damen

To compute  $f^{\mathrm{ms}}(X)$ , we start with subsequences of one element, moving up one scale at a time by combining results according to Eq. (9), obtaining  $f^{\mathrm{ms}}(X')$  for all  $X' \subseteq X$  in the process. To compute ESVs simultaneously, we compute the expected model scores on subsequences with/without each element at each scale. We then combine these to obtain all elements' ESVs as in Eq. (7).

**Sampling Subsequences.** The bottom-up computation improves the efficiency, but doesn't deal with the exponential number of subsequences. We apply a sampling approximation to the expectations over subsequences in the definition of  $f^{\rm ms}$  (Eq. (9)) and the ESV (Eq. (7)). This is similar to the Monte-Carlo approach used in prior applications of the Shapley value [14,15] but is more sample efficient. Our sampling approach aims to maximise the number of subsequence relationships across scales to best approximate  $f^{\rm ms}$ . Given a sample of subsequences  $\mathcal{X}^s$  at scale s, we grow subsequences by one element to obtain subsequences at the next scale, We then form a pool of all possible candidate subsequences of length s + 1:

$$\mathcal{C}^{s+1} = \bigcup_{\mathcal{X}_j^s \in \mathcal{X}^s} \{\mathcal{X}_j^s \cup \{x\} : x \in X \setminus \mathcal{X}_j^s\},$$
(10)

from which we sample at most m subsequences to construct  $\mathcal{X}^{s+1}$ . We start our sampling approach with all single element subsequences:  $\mathcal{X}^1 = \{(x_i)\}_{i=1}^n$ .

**Tying it All Together.** Combining the above techniques, we present our approach in Algorithm 1 for computing ESVs. When the number of sampled subsequences, m, is chosen to be  $\max_k \binom{|X|}{k}$ , it computes the *exact* ESVs  $\phi_i^c$  for  $f^{\text{ms}}(X)$ . When m is less than this, the algorithm computes the approximate ESVs  $\phi_i^c$ . We repeat the inner loop that computes the marginal contributions across all scales a number of iterations to improve the accuracy of the approximation.

Additional Definitions. Elements whose  $\phi_i^c > 0$  are termed supporting elements for class c, otherwise, if  $\phi_i^c \leq 0$ , they are termed distracting elements.

A common question when diagnosing a model is to understand why it classified a sequence incorrectly. We utilise the linear property of ESVs to compute a class-contrastive ESV  $\delta_i$ . This is the ESV for the class-contrastive (cc) model  $f_{cc}(X) = f_{gt}(X) - f_{pt}(X)$ , where gt is the ground-truth class and pt is the incorrectly predicted class. The class-constrastive ESV can thus be computed as:  $\delta_i = \phi_i^{gt-pt} = \phi_i^{gt} - \phi_i^{pt}$ . When  $\delta_i > 0$ , the element contributes more to the ground-truth class than the predicted class.

**Limitations.** We foresee two limitations to our current implementation of ESVs: (i) when the model's output for a given class is too small  $\epsilon$ , frame-level ESVs are uninformative; (ii) when frame-level features are similar for multiple frames in the input sequence, our implementation calculates ESV for each independently, and does not benefit from feature similarities.

### 3 Experiments

We have now explained how ESVs can be computed in a tractable manner for all elements in a variable-length sequence. While our approach is applicable Algorithm 1 Element Shapley Value (ESV) computation using  $f^{\rm ms}$ 

**Input:** A sequence  $X = (x_i)_{i=1}^n : \mathbb{R}^{n \times D}$ , Single-scale models  $\{f^s : \mathbb{R}^{s \times D} \to \mathbb{R}^C\}_{s=1}^{n_{\max}}$ comprising  $f^{ms}$ , and a class c to explain. **Output**: Element Shapley Values  $\phi_i^c$  for all  $x_i \in X$ Intermediates:  $\mathcal{X}_{i}^{s}$ : A subsequence of X of length s.  $\mathcal{S}_i^s, \bar{\mathcal{S}}_i^s$ : Sum of scores over sequences that contain/don't contain element  $x_i$  at scale s.  $\mathcal{N}_i^s, \overline{\mathcal{N}}_i^s$ : Number of sequences that contain/don't contain element  $x_i$  at scale s.  $\mathcal{F}_{i}^{s}$ : Results of  $f_{c}^{\mathrm{ms}}$  on the subsequence  $\mathcal{X}_{i}^{s}$ . 1:  $\bar{\mathcal{S}}_i^0 \leftarrow f_c(\emptyset), \ \bar{\mathcal{N}}_i^0 \leftarrow 1, \ \mathcal{X}^0 \leftarrow \{\{\}\}\}$ 2: for iteration from 0 to max iterations do for scale s from 1 to n do 3:  $\mathcal{C}^s$ : Form extended subsequence candidate pool according to Eq. (10). 4:  $\mathcal{X}^s \leftarrow \text{Randomly select } \min(m, |\mathcal{C}^s|) \text{ subsequences from } \mathcal{C}^s.$ 5: $Z_j \leftarrow \sum_k \mathbb{1}[\mathcal{X}_k^{s-1} \subset \mathcal{X}_j^s]$ 6: if s = 1 then 7:  $\mathcal{F}_j^1 \leftarrow f_c^1(\mathcal{X}_j^s)$ 8: else if  $s \leq n_{\max}$  then  $| \mathcal{F}_j^s \leftarrow \frac{1}{s} \left( f_c^s(\mathcal{X}_j^s) + (s-1) \frac{1}{Z_j} \sum_k \mathbb{1}[\mathcal{X}_k^{s-1} \subset \mathcal{X}_j^s] \mathcal{F}_k^{s-1} \right)$ 9: 10: $\begin{array}{c} \textbf{else} \\ \big| \quad \mathcal{F}_{j}^{s} \leftarrow \frac{s-1}{s} \frac{1}{Z_{j}} \sum_{k} \mathbbm{1}[\mathcal{X}_{k}^{s-1} \subset \mathcal{X}_{j}^{s}] \mathcal{F}_{k}^{s-1} \end{array}$ 11: 12:
$$\begin{split} \mathcal{S}_i^s &\leftarrow \mathcal{S}_i^s + \sum_j \mathbb{1}[x_i \in \mathcal{X}_j^s] \mathcal{F}_j^s, \quad \bar{\mathcal{S}}_i^s \leftarrow \bar{\mathcal{S}}_i^s + \sum_j \mathbb{1}[x_i \notin \mathcal{X}_j^s] \mathcal{F}_j^s \\ \mathcal{N}_i^s &\leftarrow \mathcal{N}_i^s + \sum_j \mathbb{1}[x_i \in \mathcal{X}_j^s], \quad \bar{\mathcal{N}}_i^s \leftarrow \bar{\mathcal{N}}_i^s + \sum_j \mathbb{1}[x_i \notin \mathcal{X}_j^s] \end{split}$$
13:14: 15:  $\phi_i^c \leftarrow \frac{1}{|X|} \sum_{s=1}^{|X|} \mathcal{S}_i^s / \mathcal{N}_i^s - \bar{\mathcal{S}}_i^{s-1} / \bar{\mathcal{N}}_i^{s-1}$ 

to sequences of any type of data, in this paper we focus on frames in video sequences. We compare ESV against baselines and two other commonly used feature attribution methods in a frame ablation experiment. We show that the choice of attribution method can have a big impact on the frames' attribution values. We then analyse exact ESVs across a range of facets demonstrating how they can be used to understand models' behaviour. Finally, we evaluate our proposed approximation showing that we can scale up our method to compute ESVs for all frames in variable-length videos.

**Experimental Setup.** Our experiments are conducted on the validation set of the large-scale Something-Something v2 dataset [10], frequently used to probe video models [1,29,33] due to its fine-grained nature, large number of classes (174), and the temporal characteristics of most of its classes e.g. "Pulling [...] from behind of [...]". Unless otherwise stated,  $\phi_i^c$  is calculated for the ground-truth class of the example video. We implement  $f^{\rm ms}$  using single hidden-layer MLPs atop of extracted features with  $n_{\rm max} = 8$ .

ESV is best applied to models where frame level features are extracted and combined through a temporal module. In this case, only the temporal module is evaluated to compute the ESVs. We primarily explain one commonly used action recognition model: Temporal Relational Network (TRN) [1]. However, we also use



**Fig. 2.** Element attributions for two example sequences comparing our method (ESV) to Integrated Gradients (IG), and Grad-CAM (GC).

ESVs to explain Temporal Segment Network (TSN) [34], and provide open-source code for explaining other models<sup>3</sup>. Further implementation and training details can be found in Appendix C in addition to how we compute ESVs for TSN.

### 3.1 Analysing Element Shapley Values

Our analysis is provided by answering questions that give new insights into how individual frames contribute to the model's output.

How do ESVs differ from other attribution methods? We compare ESV against Integrated Gradients (IG) [17] and Grad-CAM<sup>4</sup> (GC) [20]. These methods produce attributions for each feature, so to obtain element attributions, we marginalise the spatial and channel dimensions by averaging. We chose GC as the most-commonly used approach in explaining networks for video-understanding (e.g. [35–38]) and IG as an axiomatic approach based on Aumann-Shapley values [39], an extension of the Shapley value to infinite player games. IG computes the integral of gradients along a path from a reference to the input being explained (we use the mean frame feature as our reference). We use the public implementation of these approaches from Captum [40].

We first present a qualitative comparison on two sequences in Fig. 2 demonstrating disagreement amongst the methods in their element attributions. The top example ("Closing [...]") shows that our method (ESV) highlights the first frame showing the open cupboard as well as the frame where the cupboard door has just been closed as the most important. Both IG and GC only highlight the first frame, missing the action completion, with GC considering the completion as marginally distracting. In the bottom example, ESV assigns the first three frames negative attributions, with positive attributions restricted to frames after the hand appears. ESV considers the frame with the visibly spinning key as most important while IG and SG highlight the last frame instead.

We compute the Pearson correlation between the attribution values produced by each pair of methods on the validation set. For  $ESV \times IG$ ,  $ESV \times GC$  and

<sup>&</sup>lt;sup>3</sup> https://github.com/willprice/play-fair/

<sup>&</sup>lt;sup>4</sup> We remove the ReLU from the formulation to produce signed attributions.



**Fig. 3.** TRN class score and accuracy after iteratively discarding frames in order of their attribution rank (ascending  $\blacktriangle$  vs descending  $\blacktriangledown$ ). We compare our method (ESV) to baselines (a,b) and two alternate attribution methods (c): GradCam (GC) and Integrated Gradients (IG). We keep figures (b) and (c) separate for readability. Removing frames with the highest ESV first causes the quickest degradation, whilst removing frames with the lowest ESV improves performance by avoiding distractor frames.



**Fig. 4.** (a) The percentage of supporting frames for correctly/incorrectly classified examples, comparing TSN to TRN. (b) % of examples where the frame is the most supporting/distracting show that TSN/TRN make use of frames in different positions.

 $IG \times GC$  we get 0.60, 0.56 and 0.81 respectively for TRN compared to 0.85, 0.89 and 0.82 for TSN. We note that agreement differs per model. IG is more similar to GC when analysing TRN, but more similar to ESV when analysing TSN model attributions. Critically, we believe our fairness axioms and avoidance of out-of-distribution inputs make ESV a more founded technique for element attribution as we will demonstrate next.

How does performance change when we discard frames according to their attributions? For TRN, we iteratively discard frames from the input sequence in descending/ascending order of their attributions. We compare this approach to four baselines: discarding frames from the center frame outwards, from the edges of the sequence inwards, uniformly, and randomly. Figures 3a and 3b report results of this investigation on the full validation set uniformly sampling 8 frames from each video. Whilst a boost in performance is expected when removing frames with the lowest attribution value, since this approach has privileged knowledge of the ground-truth class, these figures show that (i) on average, 4 frames in an 8-frame sequence negatively influence the class score, and (ii) it is possible to maintain the model's accuracy whilst discarding the majority (6/8) of the frames.



**Fig. 5.** Comparing the percentage of videos where  $\phi_i^c > \phi_j^c$  for TRN. This shows that, across the dataset, some frames are on average more informative than others, and for certain classes this can be quite different to the overall trend.

We perform the same test, discarding frames by attribution rank for ESV, IG and GC in Fig. 3c. When discarding 4 frames by their decreasing rank, the model's accuracy increases by 20% for ESV compared to 8% for GC and 7% for IG. These results demonstrate that ESV's attribution values are more representative of how the model values frames compared to the other attribution methods.

What can we learn from *supporting* and *distracting* frames? We analyse the proportion of supporting frames across correctly and incorrectly classified examples in Fig. 4a, comparing TRN to TSN. Correctly classified examples are more likely to have a larger proportion of supporting frames, however a number of correctly classified videos contain *distracting* frames. There are more supporting frames for TRN compared to TSN for correctly classified examples.

Is there a correlation between a frame's position in the sequence and its ESV? We consider the proportion of videos where each frame, by its position in the sequence, is the most helpful or harmful in Fig. 4b. The first/last frames are often the most impactful for TRN, as the model learns temporal relationships.

We further analyse TRN, reporting the percentage of videos for which  $\phi_i^c > \phi_j^c$ for all combinations of frame positions *i* and *j* across the validation set as well as per-class. Fig. 5 shows that (i) frames residing in the middle of the sequence have higher attributions on average than those at the edges; (ii) class trends can deviate from this distinctly. The later frames for the class "Plugging [...] into [...]" are most important, as these are the ones that discriminate it from the similar action "Plugging [...] into [...] but pulling it right out"; (iii) Early frames from "Moving [...] and [...] away from each other" contribute most, as the objects move further from one another as the video progresses.

How do subsequences of various lengths contribute to the ESV? Since ESV is the average of the marginal contributions of subsequences at each scale, we analyse the per-scale average marginal contributions to probe the contribution of subsequences of a certain length s (Fig. 6a). The average marginal contributions steadily decreases as longer subsequences are considered, indicating that the majority of the frame's attribution is already extracted from shorter subsequences. We then consider whether there is a relationship between the number of frames fed to the model and the proportion of supporting frames in Fig. 6b. The longer the subsequence, the more likely it is some frames become redundant and potentially



**Fig. 6.** TRN: (a) Box-plot of marginal contributions at each scale. (b) Increasing the number of frames fed to the model decreases the percentage of supporting frames. (c) Change in ESV when compared to neighbouring frames.

distracting. This is an interesting finding—it showcases that utilising all frames in the sequence could harm model performance in the presence of distracting frames.

Are ESVs temporally smooth? We modify 16-frame sequences by replacing one frame in the uniformly sampled sequence with the frame 1, 2, or 3 places to the left/right. We compute the normalised difference in  $\text{ESV}^5$  for the original and replaced frame. We plot the distribution of changes in Fig. 6c. The figure shows a general trend of temporal smoothness, symmetric for both left/right shifts; the further the shift in either direction, the greater the change in ESV.

What about class-contrastive ESVs? Up to this point we have focused on ESVs for one class, the ground truth class of a video. We present qualitative examples of class-contrastive ESVs in Fig. 7. These show which frames contribute more to the ground-truth class than the predicted class. In the top sequence, the first frames confuse the model contributing more to "Pulling [...] from left to right". Frames when the battery is rolling contribute highly to the ground truth class. The second sequence shows that plugging the cable into the laptop contributes to both the predicted and ground truth class, but there is insufficient support from the frame where the cable is unplugged to make the correct classification.

### 3.2 ESV Approximation Evaluation

We now evaluate our ability to approximate Shapley values through our tractable approach proposed in Section 2.3. We first compare approximate to exact ESVs on 16 frame sequences (longer sequences limited by GPU memory). We then scale up our analysis to compute ESVs for all frames in variable-length videos and show approximate ESVs are consistent with those computed exactly for shorter sequences.

In Algorithm 1, we sample m subsequences per scale every iteration. We compare the approximate ESV  $\hat{\phi}_i^c$  produced by subsequence sampling against exact ESV  $\phi_i^c$ , to assess the error introduced by approximation. For this assessment, we

<sup>&</sup>lt;sup>5</sup> As the contributions vary between sequences, we first normalise the ESVs so that the mass of contributions sum to one before comparing them, i.e.  $\sum |\phi_i^c| = 1$ .



Fig. 7. Class-contrastive ESVs comparing the ground truth class (gt) to the predicted class (pt).  $\bullet$  0  $\bullet$ 

				Re	ativ	LAD regression slope										Pearson correlation															
d)	32	93.3	64.7	53.0	46.5	41.9	38.3	35.6	33.4	31.6	30.0	0.93	0.96	0.96	0.96	0.97	0.97	0.98	0.98	0.98	0.98	0.	53 0.7	8 0.8	0.87	0.89	0.91	0.92	0.92	0.93	0.94
i.	64	62.2	44.6	36.6	32.0	28.9	26.6	24.8	23.4	22.2	21.2	0.97	0.98	0.98	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.	30 O.8	8 0.93	0.93	0.94	0.95	0.96	0.96	0.97	0.97
e,	128	41.2	29.6	24.4	21.4	19.4	17.8	16.6	15.7	15.0	14.3	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.	39 O.9	4 0.90	6 0.97	0.97	0.98	0.98	0.98	0.98	0.99
d	256	27.2	19.6	16.4	14.4	13.0	12.1	11.4	10.8	10.3	09.9	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.99	0.	95 0.9	7 0.98	8 0.99	0.99	0.99	0.99	0.99	0.99	0.99
an	512	16.6	12.1	10.2	09.0	08.2	07.6	07.2	06.8	06.6	06.3	1.00	1.00	1.00	1.00	1.00	0.99	1.00	1.00	1.00	1.00	0.	8 0.9	9 0.99	0.99	1.00	1.00	1.00	1.00	1.00	1.00
S	1024	09.5	06.9	05.8	05.2	04.8	04.5	04.3	04.1	03.9	03.8	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.	9 1.0	0 1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
		1	2	3	4	5	6	7	8	9	10	1	2	3	4	5 of It	6 orat	7	8	9	10	1	2	3	4	5	6	7	8	9	10
															١ <b>0</b> .		erat	.1011:	>										_		
		Scale sample size $m$												35	2	6	4	12	8	25	6	51	2	10	24	_					
		$\%$ of total subsequences/iteration $0.68\ 1.32\ 2.56\ 4.71\ 9.01\ 16$														16.	.19														

Fig. 8. Comparing approximate vs. exact ESVs for 16-frame videos, as we vary m, the maximum number of subsequences at each scale, (y-axis) and the number of iterations (x-axis). For relative error, lower is better, and for LAD regression slope and Pearson correlation, closer to one is better. The table shows % of subsequences sampled per iteration as m varies.

randomly sample 1,000 videos where  $f'_c(X) \ge 0.05$  so the ESVs aren't so small as to compromise the approximation evaluation. We then uniformly sample 16 frames from each video. We consider three evaluation metrics to assess the effect of sampling:

- 1. Relative error, we compute the normalised error between approximate and exact ESVs per element  $|(\hat{\phi}_i^c \phi_i^c)|/A$  where  $A = \frac{1}{|X|} \sum_j |\phi_j^c|$  is the video-level mean of the unsigned ESVs.
- 2. Bias, we fit a Least Absolute Deviance (LAD) regression between  $\phi_i^c$  and  $\phi_i^c$ . A regression slope < 1 shows an over-estimate.
- 3. Correlation, we compute Pearson's r between the approximate and exact ESVs computed for each video.

These metrics are computed per video and averaged across all videos. We present the results in Fig. 8 and demonstrate the efficiency of the sampling by reporting the percentage of subsequences sampled per iteration, for various m values in the table underneath. For instance, at m = 256 over 4 iterations, we would have considered less than 19% of all subsequences, but would achieve r = 0.99and no bias in the approximation. Increasing the sample size m and/or number



Fig. 9. ESVs for all frames computed using our approximation method (we plot the mean ESVs over 10 runs of our algorithm and shade  $\mu \pm \sigma$ , m = 1024 with 4 iterations). We also plot the approximate and exact ESVs for 16 uniformly sampled frames demonstrating these are representative of the ESVs for longer sequences.

of iterations improves all metrics. Similar to Fig. 1, in Fig. 9 we plot ESVs computed for all 49 frames in a video. Without approximation we would require  $10^{14}$  evaluations of  $f^{\rm ms}$ , but with our approximation requires only 49k evaluations.

### 4 Related Work

Feature Attribution. These techniques can be categorised into backprop [17, 20, 23, 23, 35, 41-44] or perturbation based [16, 22, 24, 43, 45] methods. Backprop methods either use standard or modified backpropagation rules to work out a feature's contribution to an output neuron. These are more attractive computationally, as they don't require multiple model evaluations, but often don't satisfy expected principles of an attribution method such as implementation invariance, sensitivity to the input or the model parameters [17, 46]. Grad-CAM [20] is most popular, but a variety of others exist [21, 23, 35, 42-44].

In contrast, perturbation-based methods evaluate the model on perturbed inputs. Occlusion [43] was first utilised, sliding an occluding patch over the image and evaluating the change in the model's output. Štrumbelj and Kononenko [14] proposed the use of Shapley values as a model explanation, initially retraining a model on every possible subset of features, but later [15] using feature substitution. SHAP [16] unified various attribution methods [23, 25, 44] by showing how each could compute or approximate Shapley values under certain assumptions (e.g. feature independence, model linearity).

Recent hybrid approaches combine perturbation with backprop [17, 22, 24]. Integrated gradients (IG) [17] uses Aumann-Shapley values [39] to compute feature attributions for differentiable models, however the method requires a reference and evaluates the model on out-of-distribution examples generated by linearly interpolating between the input and the reference. Meaningful perturbations [24] learn a mask to retain the most important pixels in an input image to a class

#### 14 W. Price and D. Damen

neuron. They add regularisers to encourage sparsity and contiguous region formation. Extremal perturbations [22] address some of the shortcomings of [24] such as the non-uniqueness of masks. Both works [22,24] produce binary masks, rather than pixel attributions.

**Explaining Video Models.** Most works [29, 35-38, 41] in attribution for video understanding use backprop methods originally designed for images, such as Grad-CAM [20] or EBP [21]. A few recent works propose video-specific attribution methods. EBP-RNN [41] extends EBP [21] to CNN+RNN to explain video models. Saliency tubes [47] use a Grad-CAM like formulation to provide spatio-temporal attributions of the last convolutional layer. They extend the approach in [48] by backpropagating gradients to different depths in the network. Mänttäri et al. [29] apply meaningful perturbations [24] to learn a temporal mask over the input. To keep the number of input frames fixed, they replace missing frames by duplication. Li et al. [30] learn a spatio-temporal mask via extremal perturbations [22]. They replace missing voxels with a reference of a blurred voxel, analysing R(2+1)D [49] and a CNN+LSTM model.

In contrast to these works, we offer the first perturbation method for video explanation that is based on the Shapley axioms of fairness. Our method obeys principled criteria (Section 2.1) and does not feed models with out-of-distribution examples as in [17, 22, 24].

Alternate efforts in explaining video models take a network-centric approach. Feichtenhofer et al. [50] use activation maximisation [42] to explain a variety of two-stream architectures [34, 51, 52], synthesizing inputs that maximally activate a chosen class neuron. Huang et al. [32] explain the temporal signal in video by training a temporal generative model to produce within-distribution videos of shorter or re-ordered frames. They explain C3D [53] on a fixed-length sequence, analysing the model's performance drop. These approaches offer a complimentary viewpoint to model explanations, and do not attempt frame attributions.

# 5 Conclusion

In this paper, we introduced the notion of element attribution for determining the contribution of each frame to the output of an action recognition model. We adopted the Shapley axioms as a way of determining these contributions in a fair and principled manner in our proposed Element Shapley Value (ESV). We side-step the issues present in feature-attribution by reformulating the Shapley value and utilise multi-scale models to determine the marginal contributions of frames without having to substitute them. We used ESV to analyse frame-based action recognition models on the Something-something dataset, showing insights into how trained models make classification decisions.

Acknowledgement. Research supported by EPSRC UMPIRE (EP/T004991/1) and EPSRC Doctoral Training Partnershipts (DTP).

## References

- Zhou, B., Andonian, A., Oliva, A., Torralba, A.: Temporal Relational Reasoning in Videos. In: Proceedings of the European Conference on Computer Vision (ECCV). (2018) 803–818
- 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 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2018) 6450–6459
- Hussein, N., Gavves, E., Smeulders, A.W.M.: Timeception for Complex Action Recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2019) 254–263
- Girdhar, R., Carreira, J., Doersch, C., Zisserman, A.: Video Action Transformer Network. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2019) 244–253
- Chen, Y., Kalantidis, Y., Li, J., Yan, S., Feng, J.: Multi-fiber Networks for Video Recognition. In: Proceedings of the European Conference on Computer Vision (ECCV). (2018)
- Lin, J., Gan, C., Han, S.: TSM: Temporal Shift Module for Efficient Video Understanding. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2019) 7083–7093
- Feichtenhofer, C., Fan, H., Malik, J., He, K.: SlowFast Networks for Video Recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2019) 6202–6211
- Abu-El-Haija, S., Kothari, N., Lee, J., Natsev, P., Toderici, G., Varadarajan, B., Vijayanarasimhan, S.: YouTube-8M: A Large-Scale Video Classification Benchmark. ArXiv160908675 Cs (2016)
- Kay, W., Carreira, J., Simonyan, K., Zhang, B., Hillier, C., Vijayanarasimhan, S., Viola, F., Green, T., Back, T., Natsev, P., Suleyman, M., Zisserman, A.: The Kinetics Human Action Video Dataset. ArXiv170506950 Cs (2017)
- Goyal, R., Ebrahimi Kahou, S., Michalski, V., Materzynska, J., Westphal, S., Kim, H., Haenel, V., Fruend, I., Yianilos, P., Mueller-Freitag, M., Hoppe, F., Thurau, C., Bax, I., Memisevic, R.: The "Something Something" Video Database for Learning and Evaluating Visual Common Sense. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). (2017)
- Diba, A., Fayyaz, M., Sharma, V., Paluri, M., Gall, J., Stiefelhagen, R., Van Gool, L.: Large Scale Holistic Video Understanding. In: Proceedings of the European Conference on Computer Vision (ECCV). (2020)
- 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 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2018) 6047–6056
- Damen, D., Doughty, H., Maria Farinella, G., Fidler, S., Furnari, A., Kazakos, E., Moltisanti, D., Munro, J., Perrett, T., Price, W., Wray, M.: Scaling Egocentric Vision: The EPIC-KITCHENS Dataset. In: Proceedings of the European Conference on Computer Vision (ECCV). (2018) 720–736
- Štrumbelj, E., Kononenko, I.: An Efficient Explanation of Individual Classifications using Game Theory. J. Mach. Learn. Res. 11 (2010) 1–18

- 16 W. Price and D. Damen
- Strumbelj, E., Kononenko, I.: Explaining prediction models and individual predictions with feature contributions. Knowl Inf Syst 41 (2014) 647–665
- Lundberg, S.M., Lee, S.I.: A Unified Approach to Interpreting Model Predictions. In Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R., eds.: Advances in Neural Information Processing Systems 30 (NeurIPS). Curran Associates, Inc. (2017) 4765–4774
- Sundararajan, M., Taly, A., Yan, Q.: Axiomatic Attribution for Deep Networks. In: Proceedings of the 34th International Conference on Machine Learning (ICML). ICML'17, Sydney, NSW, Australia, JMLR.org (2017) 3319–3328
- Sundararajan, M., Najmi, A.: The many Shapley values for model explanation. In: Proceedings of the 37th International Conference on Machine Learning (ICML). ICML'20 (2020)
- Shapley, L.S.: A Value for n-Person Games. In: Contributions to the Theory of Games (AM-28), Volume II. Volume 2. Princeton University Press, Princeton (1953)
- Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-CAM: Visual Explanations From Deep Networks via Gradient-Based Localization. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). (2017) 618–626
- 21. Zhang, J., Bargal, S.A., Lin, Z., Brandt, J., Shen, X., Sclaroff, S.: Top-Down Neural Attention by Excitation Backprop. Int J Comput Vis **126** (2018) 1084–1102
- Fong, R., Patrick, M., Vedaldi, A.: Understanding Deep Networks via Extremal Perturbations and Smooth Masks. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). (2019) 2950–2958
- Bach, S., Binder, A., Montavon, G., Klauschen, F., Müller, K.R., Samek, W.: On Pixel-Wise Explanations for Non-Linear Classifier Decisions by Layer-Wise Relevance Propagation. PLOS ONE 10 (2015) e0130140
- Fong, R.C., Vedaldi, A.: Interpretable Explanations of Black Boxes by Meaningful Perturbation. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV). (2017) 3429–3437
- Ribeiro, M.T., Singh, S., Guestrin, C.: "Why Should I Trust You?": Explaining the Predictions of Any Classifier. In: Proceedings of the 22nd ACM International Conference on Knowledge Discovery and Data Mining (SIGKDD). KDD '16, San Francisco, California, USA, Association for Computing Machinery (2016) 1135–1144
- Carreira, J., Zisserman, A.: Quo Vadis, Action Recognition? A New Model and the Kinetics Dataset. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2017) 6299–6308
- Young, H.P.: Monotonic solutions of cooperative games. Int J Game Theory 14 (1985) 65–72
- Štrumbelj, E., Kononenko, I., Robnik Šikonja, M.: Explaining instance classifications with interactions of subsets of feature values. Data Knowl. Eng. 68 (2009) 886–904
- Mänttäri, J., Broomé, S., Folkesson, J., Kjellström, H.: Interpreting video features: A comparison of 3D convolutional networks and convolutional LSTM networks. ArXiv200200367 Cs (2020)
- Li, Z., Wang, W., Li, Z., Huang, Y., Sato, Y.: A Comprehensive Study on Visual Explanations for Spatio-temporal Networks. ArXiv200500375 Cs (2020)
- Sturmfels, P., Lundberg, S., Lee, S.I.: Visualizing the Impact of Feature Attribution Baselines. Distill 5 (2020) e22
- 32. Huang, D.A., Ramanathan, V., Mahajan, D., Torresani, L., Paluri, M., Fei-Fei, L., Carlos Niebles, J.: What Makes a Video a Video: Analyzing Temporal Information in

Video Understanding Models and Datasets. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2018) 7366–7375

- 33. Dwibedi, D., Sermanet, P., Tompson, J.: Temporal Reasoning in Videos Using Convolutional Gated Recurrent Units. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). (2018) 1111–1116
- 34. Wang, L., Xiong, Y., Wang, Z., Qiao, Y., Lin, D., Tang, X., Van Gool, L.: Temporal segment networks: Towards good practices for deep action recognition. In: Proceedings of the European Conference on Computer Vision (ECCV), Springer (2016) 20–36
- Chattopadhay, A., Sarkar, A., Howlader, P., Balasubramanian, V.N.: Grad-CAM++: Generalized Gradient-Based Visual Explanations for Deep Convolutional Networks. In: 2018 IEEE Winter Conference on Applications of Computer Vision (WACV). (2018) 839–847
- Doughty, H., Damen, D., Mayol-Cuevas, W.: Who's Better? Who's Best? Pairwise Deep Ranking for Skill Determination. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2018) 6057–6066
- 37. Goyal, R., Mahdisoltani, F., Berger, G., Gharbieh, W., Bax, I., Memisevic, R.: Evaluating visual "common sense" using fine-grained classification and captioning tasks. In: 6th International Conference on Learning Representations, Workshop Track Proceedings (ICLRW). (2018)
- Srinivasan, V., Lapuschkin, S., Hellge, C., Müller, K.R., Samek, W.: Interpretable human action recognition in compressed domain. In: 2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). (2017) 1692– 1696
- Aumann, R.J., Shapley, L.S.: Values of Non-Atomic Games. Princeton University Press (1974)
- Kokhlikyan, N., Miglani, V., Martin, M., Wang, E., Reynolds, J., Melnikov, A., Lunova, N., Reblitz-Richardson, O.: PyTorch captum. GitHub Repos. (2019)
- 41. Adel Bargal, S., Zunino, A., Kim, D., Zhang, J., Murino, V., Sclaroff, S.: Excitation Backprop for RNNs. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2018) 1440–1449
- 42. Simonyan, K., Vedaldi, A., Zisserman, A.: Deep Inside Convolutional Networks: Visualising Image Classification Models and Saliency Maps. In Bengio, Y., LeCun, Y., eds.: 2nd International Conference on Learning Representations (ICLR). (2014)
- 43. Zeiler, M.D., Fergus, R.: Visualizing and Understanding Convolutional Networks. In Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T., eds.: Proceedings of the European Conference on Computer Vision (ECCV). Lecture Notes in Computer Science, Cham, Springer International Publishing (2014) 818–833
- 44. Shrikumar, A., Greenside, P., Kundaje, A.: Learning Important Features Through Propagating Activation Differences. ArXiv170402685 Cs (2019)
- Petsiuk, V., Das, A., Saenko, K.: RISE: Randomized Input Sampling for Explanation of Black-box Models. In: British Machine Vision Conference 2018 (BMVC), BMVA Press (2018) 151
- Adebayo, J., Gilmer, J., Muelly, M., Goodfellow, I., Hardt, M., Kim, B.: Sanity Checks for Saliency Maps. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R., eds.: Advances in Neural Information Processing Systems 31 (NeurIPS). Curran Associates, Inc. (2018) 9505–9515
- 47. Stergiou, A., Kapidis, G., Kalliatakis, G., Chrysoulas, C., Veltkamp, R., Poppe, R.: Saliency Tubes: Visual Explanations for Spatio-Temporal Convolutions. In: 2019 IEEE International Conference on Image Processing (ICIP). (2019) 1830–1834

- 18 W. Price and D. Damen
- Stergiou, A., Kapidis, G., Kalliatakis, G., Chrysoulas, C., Poppe, R., Veltkamp, R.: Class Feature Pyramids for Video Explanation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops (ICCVW). (2019) 4255– 4264
- Xie, S., Sun, C., Huang, J., Tu, Z., Murphy, K.: Rethinking Spatiotemporal Feature Learning: Speed-Accuracy Trade-offs in Video Classification. In: Proceedings of the European Conference on Computer Vision (ECCV). (2018) 305–321
- Feichtenhofer, C., Pinz, A., Wildes, R.P., Zisserman, A.: Deep Insights into Convolutional Networks for Video Recognition. Int J Comput Vis 128 (2020) 420–437
- Feichtenhofer, C., Pinz, A., Wildes, R.: Spatiotemporal Residual Networks for Video Action Recognition. In Lee, D.D., Sugiyama, M., Luxburg, U.V., Guyon, I., Garnett, R., eds.: Advances in Neural Information Processing Systems 29 (NeurIPS). Curran Associates, Inc. (2016) 3468–3476
- Feichtenhofer, C., Pinz, A., Zisserman, A.: Convolutional Two-Stream Network Fusion for Video Action Recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2016) 1933–1941
- 53. Tran, D., Bourdev, L., Fergus, R., Torresani, L., Paluri, M.: Learning Spatiotemporal Features with 3D Convolutional Networks. In: Proceedings of the 2015 IEEE International Conference on Computer Vision (ICCV). ICCV '15, Washington, DC, USA, IEEE Computer Society (2015) 4489–4497