

Class Feature Pyramids for Video Explanation

Alexandros Stergiou¹Georgios Kapidis^{1,4}
Ronald Poppe¹Grigorios Kalliatakis²
Remco Veltkamp¹Christos Chrysoulas³¹Utrecht University
Utrecht
The Netherlands{a.g.stergiou, g.kapidis,
r.w.poppe, r.c.veltkamp}@uu.nl²University of Essex
Colchester
United Kingdom

gkallia@essex.ac.uk

³London South Bank University
London
United Kingdom

chrysouc@lsbu.ac.uk

⁴Noldus IT
Wageningen
The Netherlandsgeorgios.kapidis
@noldus.nl

Abstract

Deep convolutional networks are widely used in video action recognition. 3D convolutions are one prominent approach to deal with the additional time dimension. While 3D convolutions typically lead to higher accuracies, the inner workings of the trained models are more difficult to interpret. We focus on creating human-understandable visual explanations that represent the hierarchical parts of spatio-temporal networks. We introduce Class Feature Pyramids, a method that traverses the entire network structure and incrementally discovers kernels at different network depths that are informative for a specific class. Our method does not depend on the network's architecture or the type of 3D convolutions, supporting grouped and depth-wise convolutions, convolutions in fibers, and convolutions in branches. We demonstrate the method on six state-of-the-art 3D convolution neural networks (CNNs) on three action recognition (Kinetics-400, UCF-101, and HMDB-51) and two egocentric action recognition datasets (EPIC-Kitchens and EGTEA Gaze+).¹

1. Introduction

Deep neural networks have revolutionized many domains, *e.g.*, image recognition, speech recognition and knowledge discovery [22, 25, 26]. With the success of those modern neural networks, comes the need to explain their decisions – including understanding how they will behave in the real world, detecting model bias, and for scientific curiosity. However, their nested non-linear structure makes them highly non-transparent, *i.e.*, neurons and layers in the network structure that significantly contribute to a class and instance-specific features cannot be accurately determined

¹Our code is available at <https://git.io/fjDCW>

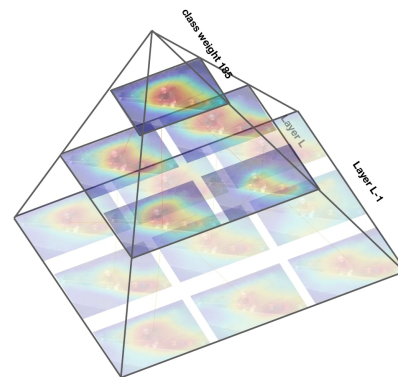


Figure 1: **Class Feature Pyramid.** A top-down view of class-specific feature information at each layer, represented as a hierarchical pyramid. The figure is created for an *ab-seiling* video from Kinetics-400 on a 3D-ResNet152. Figure is animated.

solely by the model's decisions. Therefore these models are typically regarded as *black boxes*.

We focus on convolutional neural networks (CNNs), deep neural networks that use kernels to process the input data. Initial efforts to provide visual explanation of 2D CNNs followed the approach of learning a separate linear layer with global feature representations [55]. This approach has been widely used in order to discover the image regions that deep and complex convolutional architectures find as most informative [31, 56]. The functionality of such methods stretches beyond simple information representation, providing additional benefits for various tasks. Recent works have studied class activations for further empowering models and as the basis for teacher-student networks [20, 29, 52, 54]. Others have focused on utilizing

feature information for restricting model capacity through ranking of how globally informative features are [9, 30].

Although impressive progress has been made in the domain of visual explanations in the 2D image domain [2, 18, 35, 38], visualization techniques for 3D convolution operations are scarce. This is mainly due to the additional time dimension of spatio-temporal frame sequences.

In this work, we bridge the gap between video classification and the interpretability of 3D-CNNs with a novel method *Class Feature Pyramids*: a plug-and-play visualization method specifically designed for 3D-CNNs. Class Feature Pyramids enable *back-stepping*² across multiple network layers and generating a class dependency graph that provides a tree-like hierarchy of the most informative features and their neural connections to feature activations in the preceding and succeeding layers. This enables the creation of a pyramid-like representation of class-informative layer features as seen in Figure 1.

Our contributions are summarized as follows:

- We introduce Class Feature Pyramids to improve the interpretability of 3D-CNNs by finding class-specific features across all network’s layers and neglecting 3D convolutional kernel’s spatio-temporal locality. Class Feature Pyramids are developed as a generic, plug-and-play visualization technique that can be used for any architecture that includes different types of convolutions or connections without any adjustments in the network architecture.
- We extensively study the properties of Class Feature Pyramids on visualizing layer-wise features as a concatenation of multiple extracted features and individual activation maps, based on their connection with other features with high activations from a previous layer.
- We apply Class Feature Pyramids to six top-performing spatio-temporal models in five benchmark action recognition datasets Kinetics-400 [6], UCF-101 [41], and HMDB-51 [23] and egocentric action recognition datasets EPIC-kitchens [10] and EGTEA Gaze+ [27]. We show that for action recognition, our method helps to uncover the time-space regions that networks focus on at different depths, while for egocentric action recognition, our method exposes the salient areas of the hand and object movements that the network associates with different action classes.

The remainder of the paper is structured as follows. We proceed with a discussion of related work on visualization techniques for CNNs. We introduce Class Feature Pyramids in Section 3. A quantitative and qualitative analysis of their

²We define the exploration of cross-layer features as *back-stepping* through the network in order to construct an association between high-level and low-level features.

performance appears in Sections 4 and 5, respectively. We conclude in Section 6.

2. Related work

While modern CNNs involve high complexity[5, 19, 46, 58], the connection between model architectural evaluation and interpretability has become closely associated with the overall model comprehensibility [16]. In most works that focus on network explainability, general feature information is extracted and presented by methods that only target the most prevalent high-level features learned by a network [36, 39, 49, 55]. They do not provide an extensive approach that can accommodate features used in different parts of the architecture. The rise of model complexity further indicates that explanation methods that solely consider the output of the entire model do not provide sufficient information when comparing different large architectures.

One of the first attempts to directly evaluate the importance and necessity of different parts of the network was that of Springenberg *et al.* [42]. Their *guided back-propagation* method made use of deconvolutions without the need of sub-sampling information in order to simulate a backward pass through the network. This also excluded negative gradients during the backward pass which relates to information from neural weights that decrease the high-layer activations. Through this, models can be decomposed to provide an understanding of the pixel-wise information extracted by multi-layer architectures [4]. The proposed *Layer-wise Relevant Propagation* (LRP) [24] further enhances the cross-layer activation search by relating image parts to the corresponding class. Supplementary to LRP, *DeepLIFT* [37] used a score-keeping technique to compare neural activations through back-propagating over multiple network layers. The *DeepLIFT* approach primarily focuses on understanding the importance of cross-layer connections through positive or negative score assignment. Other works on back-propagating information include optimization in order to find the activations that excite specific neurons through iteratively calculating the derivatives for permuting the neuron input [32].

For video recognition, advances in explaining the features learned by spatio-temporal models have been scarce. Apart from the increased complexity in terms of the represented information in clips, visual descriptions in video models have been a challenging task also because of the lack of a fixed strategy based on which spatio-temporal models are built. Ways of representing the time dimension include Two-Stream networks [40] that additionally use optical flow for capturing movement [12, 33, 50], while others make use of recurrent cells [14, 28, 57]. One of the most dominant approaches has been the extension of 2D convolutions to 3D convolutions by including time as an additional dimension [3, 21]. 3D convolutions have yielded significant

interest in the video recognition field [6, 15, 48, 47, 13]. Since the popularization of 3D-CNNs for video recognition, some early attempts have been made for visualizing the activations of the learned features. An approach extending LRP to space-time volumes was introduced by Anders *et al.* [1] in which a border effect was used for determining the objects and parts in frames that the network focuses on.

In this paper, in order to enable a hierarchical representation of kernel layer activations in 3D-CNNs, we propose *Class Feature Pyramids* for discovering kernel correspondence on the basis of how informative kernel activations are considered for classes. The visualization of kernel activations are based on the *Saliency Tubes* approach [43] which uses the spine interpolate of spatio-temporal kernel activations in order to create a representation conjoined with the used clip.

3. Discovery of prominent spatio-temporal feature combinations

A number of previous works [7, 43] have focused on creating representations of the regions in space and time that 3D-CNNs focus on when considering a particular class instance. These regions are class-specific and correspond to the activation maps produced in the last convolutional operation of the network. A shortfall of these methods is the fact that visually similar classes are salient in the same spatio-temporal region because the majority of their activations are the same, while only a small number of activations are distinct to a specific class. *Class Feature Pyramids* explicitly target such features and their corresponding hierarchy, based on activations from previous layers of the network.

In this section, we will describe how cross-layer feature associations are found and how different networks can be back-propagated in a hierarchical way given their architectural building blocks. In Section 3.1 and through Figure 2 we demonstrate how class information can be back-propagated through earlier layers. In Sections 3.2 and 3.3, we provide an overview of how cross-layer feature dependencies are discovered. A description of how Class Feature Pyramids are used over different networks is found at Section 3.4.

3.1. Information propagation from class prediction

We denote the class weight tensor of the final prediction layer ($w^{[p]}$) with ($[p]$) the layer index. We use Softmax to identify the class c with the maximum probability and find corresponding class weight vector ($w_c^{[p]}$), denoted with red in Figure 2. To discover how each feature value of the selected weight vector can affect the final predictions, we perform a channel-wise multiplication between the prediction layer activation map ($a^{[p]}$) and the discovered class weight tensor. The produced class-based activation map ($a_c^{*[p]}$) has

the same dimensionality as the initial input activation map. However, each of the new features is solely based on the selected class ($[c]$). A min-max feature scaling is applied to the activations in order to have a probabilistic distribution over the channel and feature activations, for each feature index ($i \in \{0, \dots, d\}$):

$$\bar{a}_c^{*[p]} = \frac{a_{c,i}^{*[p]} - \min(a_c^{*[p]})}{\max(a_c^{*[p]}) - \min(a_c^{*[p]})} \forall i \in \{0, \dots, d\} \quad (1)$$

In order to use globally aggregated class-based information, normalized over a probabilistic distribution ($\bar{a}_c^{*[p]}$), we proceed with a third operation that explores the channel-wise dependencies of the class and the features extracted in a particular layer. This operation needs to be able to deal with both small probabilistic feature distributions, in the case of visually similar classes, and broader probability distributions, for easily distinguishable cases. Additionally, we need to be able to highlight multiple features in contrast to focusing on a single value in a one-hot fashion. These requirements led to the use of a simple monotonic shifted logistic sigmoid function, given a user-defined threshold value (θ), such as in Equation 2. This also helps to mitigate the overall complexity of the method and increases the computational efficiency. The parameterization of the gate mechanism based on this threshold creates a bottleneck which can effectively reduce the pooled activation map's channel dimensionality to solely include high activations. We discuss the threshold parameter in Section 4.1.

$$feats_i = \{i : F_i^{[p]} > 0\} \text{ where } \mathbf{F}^{[p]} = \frac{1}{1 + e^{-x+\theta}} \quad (2)$$

Through Equation 2, the indexes ($feats_i$) of the most dominant class features can be identified. The features detected by the sigmoid function have a direct correspondence to those in the previous layer's output, which in turn are concatenated given a set of kernels ($k^{[l]}$). Therefore, each dimension in the activations volume is directly related to a specific kernel. The influence of this specific kernel on the final class prediction is based on the aforementioned logistic function, shown in Figure 2 with the orange cross-layer connections.

3.2. Cross-layer feature dependencies

As seen in Figure 2, the challenge when back-stepping through earlier time-space features extracted by the network, is the complexity of these features and their correspondence to class features used in predictions. This relates to the curse of dimensionality as there is no straightforward approach to represent higher-dimensional signals (such as those found in deeper layer activations) to a lower-dimensional space (e.g. early network layers) and vice

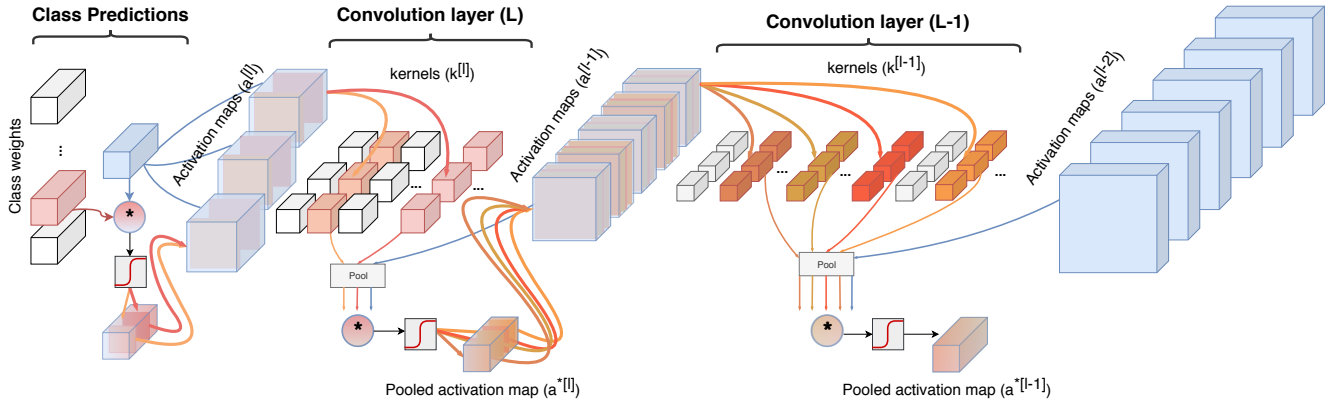


Figure 2: **Back-step process for discovering significant class-specific features.** Layer weights and input activation maps are pooled to create global representations of the information in each space-time volume, alleviating the locality of convolutions. The selected kernels are multiplied element-wise to create new pooled class-based activation maps specific to only the selected features in the layer that exhibit large contribution for the specific class. The location of the highest activation in the class-based activations is found through a sigmoid applied over the volume. Features in a previous layer are iteratively discovered for each high activation in the current layer, thus creating a class feature hierarchy.

versa. Even in the context of maintaining the same dimensional space in terms of size, as the operations performed are followed by non-linearities, the problem of cross-layer feature correspondence persists as the feature space representing information is different across pairs of non-linear layers. An additional problem specific to CNNs is the strict locality of their operations. Since a kernel's receptive field is determined by the layer number, with the input volume gradually decreasing in size through a forward pass, later layers tend to probe larger spatio-temporal patches with larger number of features. Therefore the feature region that each kernel is applied to, also varies in size.

In *Class Feature Pyramids*, information is considered in a global manner, which implies that the localities of kernels are transformed in order to hold the accumulated information. This is done during the feature correspondence step which does not handle the detection of specific regions. Specifically, to back-step features from a convolution layer (L) to $(L - 1)$, based on the indexes ($feats_i$) in the activations, the corresponding kernels ($kernels_j^{[L]}$) are selected. Traversal through features in kernels and in the activation maps ($a^{[L]}$) is accomplished by pooling both vectors to sizes equal to that of their features:

$$a^{*[L]} = \prod_{j=0}^{a^{[L]}} a^{[L-1]} * w_j^{[L]}, \forall j \in feats_i$$

where, $a^{[L-1]} = Pool(a^{[L-1]})$ and $w^{[L]} = Pool(w^{[L]})$

(3)

As the created activation map ($a^{*[L]}$) now correlates with the layer's class information in the absence of local fea-

tures, the rest of the method is similar to how feature correspondence between layers is discovered for the class prediction layer in the beginning. The class-feature activation map ($a^{*[L]}$) is normalized to a probabilistic distribution over the features (Equation 1). Based on the probabilistic class-feature vector, the features with the highest activations in layer $(L - 2)$ can be identified through the sigmoid activation function of Equation 2, in accordance with threshold value θ .

3.3. Layer-wise and feature-wise hierarchies

An intrinsic part of the network's overall explainability is to examine the value of kernels of a specific layer (L) individually or as group of feature activations of a previous layer $(L - 1)$. The necessity for understanding the cross-layer activation dependencies becomes clear when visualizing shallow layers. Such layers have limited feature complexity and kernels is activated by many kernels in deeper layers. This one-to-many association impedes the creation of a coherent kernel dependency graph. To ensure a balanced view for connection across different layers, the activations found from the back-step process can be viewed both in individual kernels as separate concatenation of activations from the previous layer, as well as for the entire layer providing a concatenation of all informative kernels in the layer.

Feature-wise kernel association: Individual kernels in layer (L) can be represented as the average sum of activations from kernels in the previous layer $(L - 1)$. This is based on discovering the indexes of the most descriptive features in the scaled and pooled activation ($a_k^{*[L]}$, where $k \in kernels_j^{[L]}$) and performing a back-step to

concatenate all activation maps in the previous layer that are part of these indexes. Then, the identified activation maps are averaged in order to be represented on top of the original input clip.

Layer-wise feature relationships: Supplementary to the singular kernel-oriented approach, layer-wise class activations aim towards a compact representation of the features associated with a specific class in an entire layer. In comparison to the kernel-based approach, and on top of their compact feature representation, layer-wise visualization also has the advantage that the number of occurrences of activations of specific kernels in the layer-wise volume is based on the number of connections that the filters have to features of previous layers. Therefore, the layer representation directly relates to the number of cross-layer connections, as informative features for multiple kernels in the following layer will have higher values than those with lower numbers of connections.

3.4. Convolution block type invariance

Visualizations produced by *Class Feature Pyramids* are invariant to the convolution types and connections used. In comparison to other feature visualization methods [44, 45, 53], the nature of the convolution operation has no impact on our method and can be used even in architectures that include convolutions performed over the channel volumes, in parallel, and with varying kernel and channels dimensions performed at the same layer. In particular, our work addresses three different convolution operation types that exhibit a high degree of complexity in the way information is connected, and how they are performed. We illustrate our approach for these specific cases in Figure 3.

Residual connections: Spatio-temporal networks that employ residual connections have been widely used in video recognition [11, 17, 34, 48]. In cases where bottleneck blocks are used, the back-step process becomes more complex as information is divided between two paths. This does not allow a straightforward hierarchical description of a backward pass through the network. We therefore include such cases in our method by creating tensors with values of one. The purpose of the created tensors is to act as a direct link between the proceeding and succeeding layers since the single value translates to an always-active state for the gate-keeper activation function. This translates to all the discovered feature indexes in the previous layer with high activations being passed through the next layer directly. Through this technique, activation indicxs are shared between previous network layers in the residual branch in the same manner as those in the main pathway of the block.

Grouped convolutions: Additional operations include the use of cross-channel convolutions performed in groups [17, 47]. These types of convolutions are especially challenging for back-stepping and visualization because their

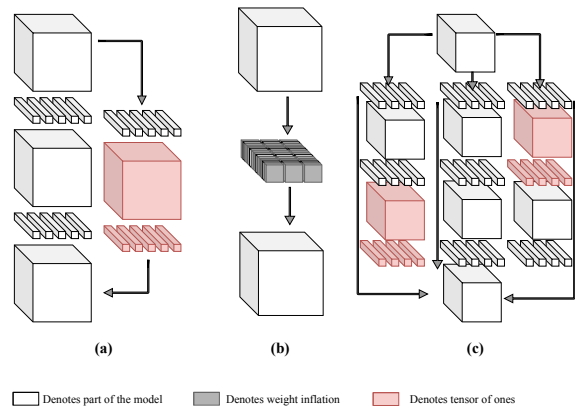


Figure 3: **Back-step for different convolution types.** (a) **Residual connections**, kernels and activations of ones (red) are used to allow the discovery of channel associations with layers that are connected with residual connections. (b) **Grouped convolutions** in which the kernels are inflated (denoted with gray) to hold the same dimensions as the input. (c) **Convolutions in branches** are performed in parallel with the possibility of uneven depth. In these cases, the branch with the maximum number of convolutions is selected as the base with tensors of ones added to the other branches.

filters are of various channel depth. Consequently, there is no immediate association between the features that the filter uses and the total number of activations in the layer. This is dealt with by explicitly inflating each of the grouped kernels to correspond to the same dimensional space as the activation maps. This simulates how channel-wise operations are executed.

Convolutions in branches: The approach of decoupling information to multiple branches and streams has also been widely used for video classification [6, 8, 13, 51]. Multiple pathways are created from the same activation maps, in which different operations occur. The variation in the type of operations and the number of operations adds a degree of ambiguity for constructing such blocks in a hierarchical manner. In these cases, back-steps through branches and pathways are accomplished with one-valued kernels and activation maps that act as small sub-structures that allow information to be passed directly across various paths.

4. Computational complexity and latency

Class Feature Pyramids is a generic visualization method that can be applied regardless of the overall complexity of the convolution operations or the between-layer connections. In this section, we describe the different spatio-temporal networks we employed to demonstrate our method and discuss the running times required for identify-

ing the kernels with the highest activations per layer.

4.1. Inference and running times

We evaluate our approach on six different spatio-temporal networks to compare the times required during the back-step kernel search with different thresholds. The results are summarized in Table 1. We progressively increase the threshold value only to be proportional to the complexity of the model to ensure that a sufficient number of features is found for each architecture. Apart from the network architecture and the threshold value, latency also depends on the number of layers that class features are back-stepped to. This relates to features in earlier layers being reached slower than features in deeper layers. This is primarily based on the complexity of the considered features as more general features in early layers will have significantly more connections to high-layer features. Higher-layer features are more specific for certain classes.

Network	GFLOPS	Back-step time (msec)	# layers	θ
Multi-FiberNet [8]	22.70	24.43	3	0.6
I3D [6]	55.79	23.21	1 + mixed5c	0.65
ResNet50-3D [17]	80.32	21.39	3	0.55
ResNet101-3D [17]	110.98	39.48	3	0.6
ResNet152-3D [17]	148.91	31.06	3	0.6
ResNeXt101-3D [17]	76.96	70.49	3	0.6

Table 1: **Evaluation of running times.** The threshold value (θ) is based on the model complexity. All architectures are back-stepped for three layers. For I3D, this corresponds to the class filters and the last *mixed* block. All times are obtained on a machine with $2 \times$ Nvidia GTX 1080 Ti GPUs.

When creating class-based activation maps ($a_c^{*[p]}$), the choice of global vectorized activations and features instead of performing in-layer convolutions is primarily attributed towards a significant reduction in operation time. The speed-up gained by using a global representation instead of the product of iterative local operations is proportional to the size of the spatio-temporal activation maps divided by the kernel’s dimensions. Therefore, if we consider the input activations of size $D \times H \times W$ with D , H , and W the number of frames, height and width, respectively, and the kernel sizes of $F_d \times F_h \times F_w$, the time required for iterating over the input activation map’s channels changes from $O^{n \times (F_d/D \times F_h/H \times F_w/W)}$ to O^n with vectorized volumes.

4.2. Hierarchical visual interpretability outline

We propose two ways to visualize features when generating network explanations. The first focuses on in-layer kernels being specific to the features extracted in individual layers. For the representation of each of the activations found, we create a spatio-temporal mask that contains the produced activation map of dimension (i). In the second

approach, multiple class activations are stacked together to form a multi-feature visualization for the class layer features. To deal with the curse of dimensionality, in both cases, we reshape the volume through a polynomial spline interpolation and thus create a new activation map. More specifically, we define each representation of the activation volume as a multi-dimensional point. For each pair of produced points, we define a function $f_{(x,y)}$ where (x, y) denotes the spatial extend of each frame in the clip and we find in-between points through a piece-wise polynomial function $S_{(x,y)}$ composed of an n -degree polynomial. For each activation, n is equal to T/t , where t is the temporal extend of the activation map and T is the temporal extend of the original clip. The function $S_{(x,y)}$ is defined as:

$$S(x, y) = P_i(x, y) (x_{i-1}, y_{i-1}) < (x, y) < (x_i, y_i) \forall 0 \leq i \leq n \quad (4)$$

5. Visualization results

In this section, we visualize the outputs of the back-step process as they are created from the representative layers of the network given the threshold from Equation 2, down to the selected layer or selected kernels of a layer. We explain how the proposed method can be used to visualize class features of different layers from various model architectures in two different video recognition tasks. Results on different types of class features extracted by six network architectures and 3D convolution operation variations are presented in Section 5.1. All models are evaluated on Kinetics-400, UCF-101 and HMDB-51. In Section 5.2 we present instances of class features for egocentric actions for EPIC-Kitchens and EGTEA Gaze+ that demonstrate how kernel groups of the network distinctively follow parts of the scene.

5.1. Class feature exploration in action recognition

In Figure 4, we show how features in different depths are visualized for different network architectures on the three action recognition datasets. We demonstrate the visual variations in class features observed considering how each network processes each clip. We examine the activations of spatio-temporal features across different depths of the network, rather than the regions with the highest activations for class weights. This means that the activations presented should include much more targeted regions and have significantly smaller duration. A full list of the kernels used can be found in the supplementary material.

In the first column of Figure 4 the selected example is of class *building lego* from the Kinetics-400 dataset. Feature activation visualizations between different networks vary significantly. In the case of the convolution operations that are grouped into fibers such as in MultiFiberNet, multiple

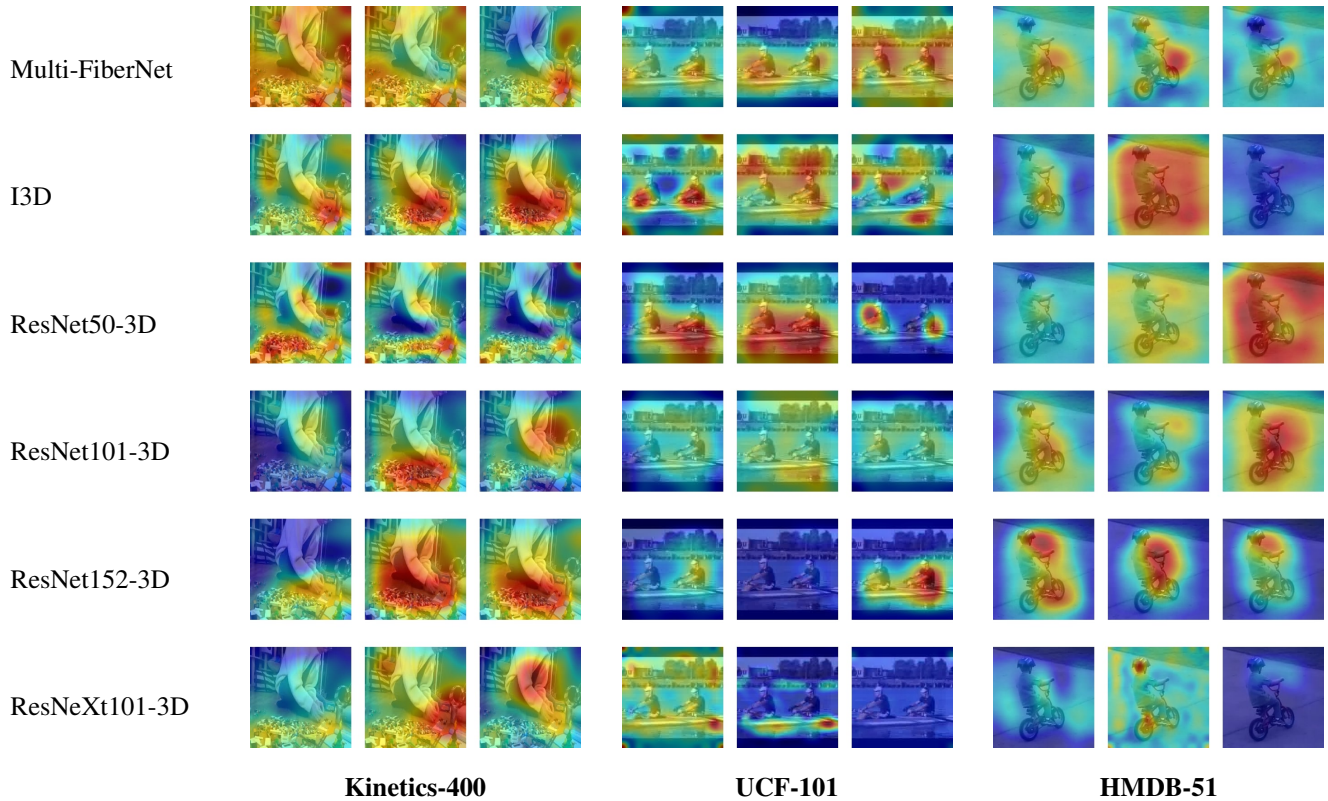


Figure 4: **Kernel visualizations for action recognition datasets.** Each row of three frames corresponds to activations of different kernels at different depths. We define the layer depth of each kernel, its indexes and the number of connections that the kernels have to the preceding layer in the supplementary material. We use different threshold values for each of the networks. The first column corresponds to an example of *building lego* in Kinetics-400. The threshold values used for the networks (row-wise) are [0.7, 0.7, 0.65, 0.65, 0.7, 0.7]. The second column is a *rowing* clip from UCF-101 with thresholds [0.8, 0.75, 0.7, 0.75, 0.8, 0.8]. Last example is from HMDB-51 of class *ride bike* with the same thresholds as UCF-101.

small spatio-temporal regions are extracted at each layer because each of the kernels is performed in a sub-volume of the layer’s activations. The same effect in a lesser extend is also present in both group-based and branch-based operations of the ResNeXt and I3D architectures, respectively. In contrast, when using models with convolutions performed over the entire activation map per single operation (*i.e.* kernels and activation maps have equal dimensions) the regions of single-class informative features are more evident. In the *building lego* class, ResNet architectures focus much more on either the hand region when picking or combining lego pieces (strong temporal cues) or the pile of pieces (strong appearance cue).

The second column presents a clip of *rowing* from UCF-101. In this example, we visualize kernels from early layers where the effects of convolution type are more notable. A direct comparison between MultiFiberNet, with each operation corresponding to 16 fibers, and ResNeXt, with 32 kernel groups, shows that both networks focus on multiple regions.

However, based on their key architectural difference, fibers consider a part of the input activation maps though feature slices. This makes the product of each operation discrete, while in contrast grouped convolutions in ResNext iterate channel-wise in a single layer of the proposed ResNeXt block. For I3D, we visualize a layer inside a *mixed branch* that includes up to four different regions each of which corresponds to the different paths inside the branch. In the case of Residual networks, the feature activation regions remain distinct to specific movements and appearance features.

For the final clip, we select an instance of *ride bike* from HMDB-51. In terms of network depth, we demonstrate a combination between features across multiple depths. In MultiFiberNet, the first visualizations are from the last convolution layer of the network while the proceeding two feature activations are from earlier network layers. Similarly, for the grouped convolutions in ResNeXt, the middle representation is drawn from earlier networks layers and is therefore focusing on a large number of spatio-temporal regions.

In contrast, the feature activations in the two side visualizations are for the final convolution layers that specifically focus on certain parts of the movement. The feature activations chosen for the I3D, and ResNet variants all are based on the last three convolution layers.

5.2. Egocentric spatio-temporal feature regions

In Figure 5, we show the input frame sequences, overlaid with the activations of the class corresponding kernels that are above θ for a layer. Results are for the MultiFiberNet network. The first four rows contain samples for EPIC-Kitchens and the last row is a clip from EGTEA Gaze+.

In the first row, we show how the network follows a ‘cut’ action. In the last kernel visualization it focuses on the area where the knife meets the carrot. Eventually, it follows its movement along with the hands. Kernels in deeper layers seem to specialize on specific parts of the frames or the content, for example the cutting board, the movements of the knife and the shoulders. Similarly, for the ‘open’ action in the second row, the network has higher class feature activations when the lid opens and the two box pieces are distinguished. This is observable by the strong activations in that area. Evidently, the other kernels are following the same pattern of specialization in different areas spatially and temporally.

Rows 3 and 4 show the same segment from EPIC-Kitchens but for a different action class. In this case, the ground truth label is ‘insert’ however the network misclassifies it for ‘put’. In the third row we visualize activation maps for action ‘insert’, while in the fourth for ‘put’. We argue that not only there is a semantic relationship between the false predictions, but the almost overlapping activated regions (rows 3 and 4, first two columns) indicate the connection between the classes in deeper network layers as well. That leads to the conclusion that the network understands the two classes from similar features and only differentiates from kernels that have importance below our threshold.

Finally, we add a visualization from EGTEA Gaze+ to show how a modified perspective does not change the visualized features. We show the features for class ‘wash’, which is the ground truth for the segment. In the first and second columns, the network tries to capture the hands, the sink and the water. In the third column, it concentrates further on the water, seemingly associating it with washing. Finally, in deeper layers it primarily captures salient objects such as the faucet and the sink’s boundaries.

6. Conclusion

Research on the explainability of spatio-temporal 3D-CNNs is limited, partly because of a scarcity of proper visualization methods. To overcome this deficit, we propose a lightweight, hierarchical approach named *Class Feature Pyramids* that captures and presents informative features

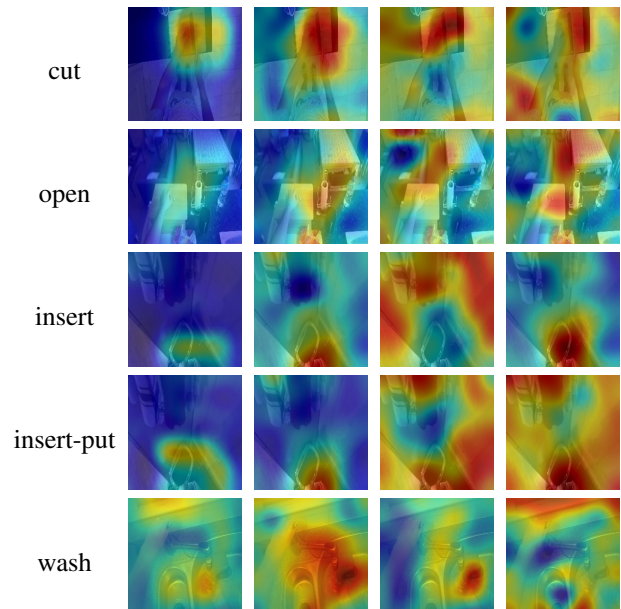


Figure 5: **Kernel visualizations for egocentric datasets.** Visual results from MultiFiberNet on datasets EPIC-Kitchens (first four rows) and EGTEA Gaze+ (last row). From left to right, class activation maps from the last layer to gradually deeper ones. The threshold values for the visualizations are [0.84, 0.64, 0.84, 0.84, 0.84].

over different layers that are specific for a class. Our method is independent of the network type and can be employed regardless of the type of 3D convolution operation. Additionally, it enables the visualization of activations in layer-wise, group-wise or kernel-wise format. Our method is therefore suitable to visualize and, subsequently, to better understand what kind of features are learned to identify a specific class. This will aid in explaining the inner workings of 3D-CNNs, identifying potential biases of trained models, and to interpret the success or failure of a classification.

We have demonstrated the merits of our approach on six 3D-CNN architectures that include different connections and convolution operations. Our experiments are performed on five common third-person and egocentric action recognition datasets. The results provide insight into the learned features. For egocentric videos, they reveal attention towards both salient objects and the movement of the hands.

7. Acknowledgments

This work is supported by the Netherlands Organization for Scientific Research (NWO) with a TOP-C2 grant for Automatic recognition of bodily interactions (ARBITER) and the EU H2020 research and innovation program under the Marie Skłodowska Curie grant agreement No 676157 (ACROSSING).

References

- [1] C. Anders, G. Montavon, W. Samek, and K.-R. Müller. Understanding patch-based learning by explaining predictions. *arXiv preprint arXiv:1806.06926*, 2018. 3
- [2] L. Anne Hendricks, R. Hu, T. Darrell, and Z. Akata. Grounding visual explanations. In *European Conference on Computer Vision (ECCV)*, pages 264–279. (Springer), 2018. 2
- [3] M. Baccouche, F. Mamalet, C. Wolf, C. Garcia, and A. Baskurt. Sequential deep learning for human action recognition. In *International Workshop on Human Behavior Understanding (HBU)*, pages 29–39. (Springer), 2011. 2
- [4] S. Bach, A. Binder, G. Montavon, F. Klauschen, K.-R. Müller, and W. Samek. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PLoS one*, 10(7):e0130140, 2015. 2
- [5] M. Berman, H. Jégou, A. Vedaldi, I. Kokkinos, and M. Douze. Multigrain: a unified image embedding for classes and instances. *arXiv preprint arXiv:1902.05509*, 2019. 2
- [6] J. Carreira and A. Zisserman. Quo vadis, action recognition? A new model and the Kinetics dataset. In *Computer Vision and Pattern Recognition (CVPR)*, pages 4724–4733. (IEEE), 2017. 2, 3, 5, 6
- [7] A. Chattopadhyay, A. Sarkar, P. Howlader, and V. N. Balasubramanian. Grad-cam++: Generalized gradient-based visual explanations for deep convolutional networks. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 839–847. (IEEE), 2018. 3
- [8] Y. Chen, Y. Kalantidis, J. Li, S. Yan, and J. Feng. Multi-fiber networks for video recognition. In *European Conference on Computer Vision (ECCV)*, 2018. 5, 6
- [9] J. Cheng, P.-s. Wang, G. Li, Q.-h. Hu, and H.-q. Lu. Recent advances in efficient computation of deep convolutional neural networks. *Frontiers of Information Technology & Electronic Engineering*, 19(1):64–77, 2018. 2
- [10] D. Damen, H. Doughty, G. Maria Farinella, S. Fidler, A. Furnari, E. Kazakos, D. Moltisanti, J. Munro, T. Perrett, W. Price, et al. Scaling egocentric vision: The epic-kitchens dataset. In *European Conference on Computer Vision (ECCV)*, pages 720–736. (Springer), 2018. 2
- [11] A. Diba, M. Fayyaz, V. Sharma, M. Mahdi Arzani, R. Yousefzadeh, J. Gall, and L. Van Gool. Spatio-temporal channel correlation networks for action classification. In *European Conference on Computer Vision (ECCV)*, pages 284–299. (Springer), 2018. 5
- [12] A. Diba, V. Sharma, and L. Van Gool. Deep temporal linear encoding networks. In *Computer Vision and Pattern Recognition (CVPR)*, pages 2329–2338. (IEEE), 2017. 2
- [13] C. Feichtenhofer, H. Fan, J. Malik, and K. He. Slow-fast networks for video recognition. *arXiv preprint arXiv:1812.03982*, 2018. 3, 5
- [14] H. Gammulle, S. Denman, S. Sridharan, and C. Fookes. Two stream LSTMs: A deep fusion framework for human action recognition. In *Applications of Computer Vision (WACV)*, pages 177–186. (IEEE), 2017. 2
- [15] D. Ghadiyaram, D. Tran, and D. Mahajan. Large-scale weakly-supervised pre-training for video action recognition. In *Computer Vision and Pattern Recognition (CVPR)*, pages 12046–12055. (IEEE), 2019. 3
- [16] R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti, and D. Pedreschi. A survey of methods for explaining black box models. *ACM computing surveys (CSUR)*, 51(5):93, 2018. 2
- [17] K. Hara, H. Kataoka, and Y. Satoh. Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet? In *Computer Vision and Pattern Recognition (CVPR)*, pages 18–22. (IEEE), 2018. 5, 6
- [18] L. A. Hendricks, Z. Akata, M. Rohrbach, J. Donahue, B. Schiele, and T. Darrell. Generating visual explanations. In *European Conference on Computer Vision (ECCV)*, pages 3–19. (Springer), 2016. 2
- [19] Y. Huang, Y. Cheng, D. Chen, H. Lee, J. Ngiam, Q. V. Le, and Z. Chen. Gpipe: Efficient training of giant neural networks using pipeline parallelism. *arXiv preprint arXiv:1811.06965*, 2018. 2
- [20] S. Jetley, N. A. Lord, N. Lee, and P. H. Torr. Learn to pay attention. In *International Conference on Learning Representations (ICLR)*, 2018. 1
- [21] S. Ji, W. Xu, M. Yang, and K. Yu. 3D convolutional neural networks for human action recognition. *Transactions on Pattern Analysis and Machine Intelligence*, 35(1):221–231, 2013. 2
- [22] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012. 1
- [23] H. Kuehne, H. Jhuang, E. Garrote, T. Poggio, and T. Serre. HMDB: a large video database for human motion recognition. In *International Conference on Computer Vision (ICCV)*, pages 2556–2563. (IEEE), 2011. 2
- [24] S. Lapuschkin, A. Binder, G. Montavon, K.-R. Müller, and W. Samek. The lrp toolbox for artificial neural networks. *The Journal of Machine Learning Research*, 17(1):3938–3942, 2016. 2
- [25] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *nature*, 521(7553):436, 2015. 1
- [26] Y. A. LeCun, L. Bottou, G. B. Orr, and K.-R. Müller. Efficient backprop. In *Neural networks: Tricks of the trade*, pages 9–48. Springer, 2012. 1
- [27] Y. Li, M. Liu, and J. M. Rehg. In the Eye of Beholder: Joint Learning of Gaze and Actions in First Person Video. In *European Conference on Computer Vision (ECCV)*, pages 639–655. (Springer), 2018. 2
- [28] J. Liu, A. Shahroudy, D. Xu, and G. Wang. Spatio-temporal LSTM with trust gates for 3D human action recognition. In *European Conference on Computer Vision (ECCV)*, pages 816–833. (Springer), 2016. 2
- [29] J. Lu, C. Xiong, D. Parikh, and R. Socher. Knowing when to look: Adaptive attention via a visual sentinel for image captioning. In *Computer Vision and Pattern Recognition (CVPR)*, pages 375–383. (IEEE), 2017. 1
- [30] J.-H. Luo, J. Wu, and W. Lin. Thinet: A filter level pruning method for deep neural network compression. In *International Conference on Computer Vision (ICCV)*, pages 5058–5066. (IEEE), 2017. 2

- [31] G. Montavon, W. Samek, and K.-R. Müller. Methods for interpreting and understanding deep neural networks. *Digital Signal Processing*, 73:1–15, 2018. 1
- [32] C. Olah, A. Mordvintsev, and L. Schubert. Feature visualization. *Distill*, 2017. <https://distill.pub/2017/feature-visualization>. 2
- [33] E. Park, X. Han, T. L. Berg, and A. C. Berg. Combining multiple sources of knowledge in deep cnns for action recognition. In *Applications of Computer Vision (WACV)*, pages 1–8. (IEEE), 2016. 2
- [34] Z. Qiu, T. Yao, C.-W. Ngo, X. Tian, and T. Mei. Learning spatio-temporal representation with local and global diffusion. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12056–12065. (IEEE), 2019. 5
- [35] W. Samek, T. Wiegand, and K.-R. Müller. Explainable artificial intelligence: Understanding, visualizing and interpreting deep learning models. *arXiv preprint arXiv:1708.08296*, 2017. 2
- [36] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, D. Batra, et al. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *International Conference on Computer Vision (ICCV)*, pages 618–626. (IEEE), 2017. 2
- [37] A. Shrikumar, P. Greenside, and A. Kundaje. Learning important features through propagating activation differences. In *International Conference on Machine Learning (ICML)*, pages 3145–3153, 2017. 2
- [38] A. Shrikumar, P. Greenside, A. Shcherbina, and A. Kundaje. Not just a black box: Learning important features through propagating activation differences. *arXiv preprint arXiv:1605.01713*, 2016. 2
- [39] K. Simonyan, A. Vedaldi, and A. Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*, 2013. 2
- [40] K. Simonyan and A. Zisserman. Two-stream convolutional networks for action recognition in videos. In *Advances in Neural Information Processing Systems (NIPS)*, pages 568–576, 2014. 2
- [41] K. Soomro, A. R. Zamir, and M. Shah. UCF101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. 2
- [42] J. Springenberg, A. Dosovitskiy, T. Brox, and M. Riedmiller. Striving for simplicity: The all convolutional net. In *International Conference on Learning Representations workshop track (ICLR)*, 2015. 2
- [43] A. Stergiou, G. Kapidis, G. Kalliatakis, C. Chrysoulas, R. Veltkamp, and R. Poppe. Saliency tubes: Visual explanations for spatio-temporal convolutions. *arXiv preprint arXiv:1902.01078*, 2019. 3
- [44] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Computer Vision and Pattern Recognition, (CVPR)*. (IEEE), 2015. 5
- [45] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna. Rethinking the inception architecture for computer vision. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2818–2826. (IEEE), 2016. 5
- [46] M. Tan and Q. Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning (ICML)*, pages 6105–6114, 2019. 2
- [47] D. Tran, H. Wang, L. Torresani, and M. Feiszli. Video classification with channel-separated convolutional networks. *arXiv preprint arXiv:1904.02811*, 2019. 3, 5
- [48] D. Tran, H. Wang, L. Torresani, J. Ray, Y. LeCun, and M. Paluri. A closer look at spatiotemporal convolutions for action recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6450–6459. (IEEE), 2018. 3, 5
- [49] L. Van Der Maaten. Accelerating t-sne using tree-based algorithms. *The Journal of Machine Learning Research*, 15(1):3221–3245, 2014. 2
- [50] L. Wang, Y. Xiong, Z. Wang, Y. Qiao, D. Lin, X. Tang, and L. Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *European Conference on Computer Vision (ECCV)*, pages 20–36. (Springer), 2016. 2
- [51] L. Wang, J. Zang, Q. Zhang, Z. Niu, G. Hua, and N. Zheng. Action recognition by an attention-aware temporal weighted convolutional neural network. *Sensors*, 18(7):1979, 2018. 5
- [52] S. Woo, J. Park, J.-Y. Lee, and I. So Kweon. Cbam: Convolutional block attention module. In *European Conference on Computer Vision (ECCV)*, pages 3–19, 2018. 1
- [53] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He. Aggregated residual transformations for deep neural networks. In *Computer Vision and Pattern Recognition (CVPR)*, pages 5987–5995. (IEEE), 2017. 5
- [54] S. Zagoruyko and N. Komodakis. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. In *International Conference on Learning Representations (ICLR)*, 2017. 1
- [55] M. D. Zeiler and R. Fergus. Visualizing and understanding convolutional networks. In *European conference on computer vision (ECCV)*, pages 818–833. (Springer), 2014. 1, 2
- [56] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba. Learning deep features for discriminative localization. In *Computer Vision and Pattern Recognition, (CVPR)*, pages 2921–2929. (IEEE), 2016. 1
- [57] W. Zhu, C. Lan, J. Xing, W. Zeng, Y. Li, L. Shen, and X. Xie. Co-occurrence feature learning for skeleton based action recognition using regularized deep LSTM networks. In *Association for the Advancement of Artificial Intelligence (AAAI)*, volume 2, page 8, 2016. 2
- [58] B. Zoph, V. Vasudevan, J. Shlens, and Q. V. Le. Learning transferable architectures for scalable image recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2