

 Open access • Book Chapter • DOI:10.1007/978-3-030-58558-7\_35

## Large Scale Holistic Video Understanding — [Source link](#)

[Ali Diba](#), [Mohsen Fayyaz](#), [Vivek Sharma](#), [Manohar Paluri](#) ...+4 more authors

**Institutions:** [Katholieke Universiteit Leuven](#), [University of Bonn](#), [Karlsruhe Institute of Technology](#), [ETH Zurich](#)

**Published on:** 23 Aug 2020 - [European Conference on Computer Vision](#)

Related papers:

- [Quo Vadis, Action Recognition? A New Model and the Kinetics Dataset](#)
- [The Kinetics Human Action Video Dataset](#)
- [SlowFast Networks for Video Recognition](#)
- [Learning Spatiotemporal Features with 3D Convolutional Networks](#)
- [The “Something Something” Video Database for Learning and Evaluating Visual Common Sense](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/large-scale-holistic-video-understanding-39pndgftthe>

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/332669673>

# Holistic Large Scale Video Understanding

Preprint · April 2019

CITATIONS

0

READS

249

7 authors, including:



[Ali Diba](#)

University of Leuven; Sensifai

20 PUBLICATIONS 462 CITATIONS

SEE PROFILE



[Mohsen Fayyaz](#)

University of Bonn

19 PUBLICATIONS 283 CITATIONS

SEE PROFILE



[Vivek Sharma](#)

Massachusetts Institute of Technology

29 PUBLICATIONS 310 CITATIONS

SEE PROFILE



[Rainer Stiefelhagen](#)

Karlsruhe Institute of Technology

330 PUBLICATIONS 7,955 CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:



Image Generation [View project](#)



Terrain [View project](#)

# Holistic Large Scale Video Understanding

Ali Diba<sup>1,\*</sup>, Mohsen Fayyaz<sup>2,\*</sup>, Vivek Sharma<sup>3,\*</sup>,  
Manohar Paluri<sup>4</sup>, Jürgen Gall<sup>2</sup>, Rainer Stiefelhagen<sup>3</sup>, Luc Van Gool<sup>1,5</sup>

<sup>1</sup>ESAT-PSI, KU Leuven, <sup>2</sup>University of Bonn, <sup>3</sup>CV:HCI, KIT, <sup>4</sup>Facebook AI Research, <sup>5</sup>CVL, ETH Zürich

{firstname.lastname}@esat.kuleuven.be, {lastname}@iai.uni-bonn.de,

{firstname.lastname}@kit.edu, mano@fb.com

## Abstract

Action recognition has been advanced in recent years by benchmarks with rich annotations. However, research is still mainly limited to human action or sports recognition - focusing on a highly specific video understanding task and thus leaving a significant gap towards describing the overall content of a video. We fill in this gap by presenting a large-scale “Holistic Video Understanding Dataset” (HVU). HVU is organized hierarchically in a semantic taxonomy that focuses on multi-label and multi-task video understanding as a comprehensive problem that encompasses the recognition of multiple semantic aspects in the dynamic scene. HVU contains approx. 577k videos in total with 13M annotations for training and validation set spanning over 4378 classes. HVU encompasses semantic aspects defined on categories of scenes, objects, actions, events, attributes and concepts, which naturally captures the real-world scenarios.

Further, we introduce a new spatio-temporal deep neural network architecture called “Holistic Appearance and Temporal Network” (HATNet) that builds on fusing 2D and 3D architectures into one by combining intermediate representations of appearance and temporal cues. HATNet focuses on the multi-label and multi-task learning problem and is trained in an end-to-end manner. The experiments show that HATNet trained on HVU outperforms current state-of-the-art methods on challenging human action datasets: HMDB51, UCF101, and Kinetics. The dataset and codes will be made publicly available.

## 1. Introduction

Video understanding is a comprehensive problem that encompasses the recognition of multiple semantic aspects, that include: a scene or an environment, objects, actions,

\*Ali Diba, Mohsen Fayyaz and Vivek Sharma contributed equally to this work, and listed in alphabetical order.

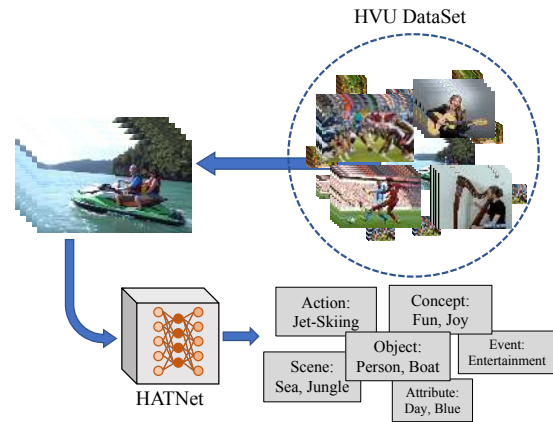


Figure 1. Holistic Video Understanding Dataset: A multi-label and multi-task dataset and HATNet as a new deep ConvNet for video classification.

events, attributes, and concepts. Even if considerable progress is made in video recognition, it is still rather limited to action recognition - this is due to the fact that there is no established video benchmark dataset that integrates joint recognition of multiple semantic aspects in the dynamic scene. While Convolutional Networks (ConvNets) have caused several sub-fields of vision to leap forward, one of the expected drawbacks of training the ConvNets for video understanding with a single-class label per task is insufficiency to describe the content of a video. This issue primarily impedes the ConvNets to learn a generic feature representation towards challenging holistic video analysis. To this end, one can easily overcome this issue by recasting the video understanding problem as multi-task classification, where multiple class labels are assigned to a video from multiple semantic aspects. Furthermore, it is possible to learn a generic feature representation for video analysis and understanding. This is in line with image classification ConvNets trained on ImageNet that facilitated the learning of generic feature representation for several vision tasks. Thus, training ConvNets on a multiple semantic as-

pects dataset can be directly applied for a holistic recognition and understanding of concepts in video data, which makes it very useful to describe the content of a video.

To address the above drawbacks, this work presents the “Holistic Video Understanding Dataset” (HVU). HVU is organized hierarchically in a semantic taxonomy that aims at providing a multi-label and multi-task large-scale video benchmarking with comprehensive list of tasks and annotations for video analysis and understanding. HVU contains approx. 577k videos in total, with 12M (11,917,055) annotations for training set and 900K for validation set spanning over 4378 classes. A full spectrum encompasses the recognition of multiple semantic aspects defined on them including 419 categories for scenes, 2651 for objects, 877 for actions, 149 for events, 160 for attributes and 122 for concepts, which naturally captures the long tail distribution of visual concepts in the real world. All these tasks are supported by rich annotations with an average of 2681 annotations per class. For instance, HVU consists of 481k, 31k, 65k samples in train, validation and test set, and is a sufficiently large dataset, which means that the scale of dataset approaches that of image datasets. The HVU action categories builds on action recognition dataset [18, 22, 24, 37, 52] and further extend then by incorporating labels of scene, objects, events, attributes, and concepts in a video. The above thorough annotations enable developments of strong algorithms for a holistic video understanding to describe the content of a video. Table 1-2 show the dataset statistics.

Furthermore, we introduce a new spatio-temporal architecture called “Holistic Appearance and Temporal Network” (HATNet) that focuses on the multi-label and multi-task learning for jointly solving multiple spatio-temporal problems simultaneously. HATNet fuses 2D and 3D architectures into one by combining intermediate representations of appearance and temporal cues, leading to a robust spatio-temporal representation.

Our HATNet achieves state-of-the-art results on the HMDB51, UCF101 and Kinetics datasets. In particular, if the model is pre-trained on HVU and fine-tuned on the corresponding datasets it outperforms models pre-trained on Kinetics. This shows the richness of our dataset as well as the importance of multi-task learning. To our best knowledge, this is the first work to focus on the training of very deep 3D ConvNets from scratch for video understanding and beating I3D pre-trained on Kinetics, thus clearly showing the impact of HVU. Further, our HATNet requires no ImageNet pre-trained image classification architecture as was the back bone of I3D. We experimentally show that HATNet achieves remarkable performance on UCF101 (96.9%), HMDB51 (74.5%) and Kinetics (73.5%).

Train	Validation	Test
481k	31k	65k

Table 1. HVU dataset statistics i.e. #videos-clips for train, validation, and test sets.

## 2. Related Work

**Action Recognition with/without ConvNets:** Over the last two decades, a multitude of action recognition techniques in videos have been proposed by the vision community. Among the hand-engineered ones that could model effectively the appearance and motion representations across frames in videos are HOG3D [23], SIFT3D [33], HOF [26], ESURF [50], MBH [6], and iDTs [44]. Several other techniques were proposed to model the temporal structure in an efficient way, such as the actom sequence model [15]; temporal action decomposition [31]; dynamic poselets [46]; ranking machines [14].

There are several approaches to end-to-end ConvNets-based action recognition [12, 21, 36, 40, 48] to exploit the appearance and the temporal information. These methods operate on 2D (individual image-level) [9, 10, 16, 38, 39, 48, 51] or 3D (video-clips or snippets of  $K$  frames) [12, 40, 41, 42]. The filters and pooling kernels for these architectures are 3D ( $x, y, \text{time}$ ) i.e. 3D convolutions ( $s \times s \times d$ ) [51] where  $d$  is the kernel’s temporal depth and  $s$  is the kernel’s spatial size. These 3D ConvNets are intuitively effective because such 3D convolution can be used to directly extract spatio-temporal features from raw videos. Carreira et al. proposed inception [20] based 3D CNNs, which they referred to as I3D [5]. More recently, some works introduced temporal transition layer that models variable temporal convolution kernel depths over shorter and longer temporal ranges, namely T3D [8]. Further in [7], Diba et al. propose spatio-temporal channel correlation that models correlations between channels of a 3D ConvNets wrt. both spatial and temporal dimensions. In contrast to these prior works, our work differs substantially in scope and technical approach. We propose an architecture, HATNet, that exploits both 2D ConvNets and 3D ConvNets to learn an effective spatio-temporal feature representation. Finally, it is worth noting the self-supervised ConvNet training works from unlabeled sources for action recognition, such as Fernando *et al.* [13] and Mishra *et al.* [29] generate training data by shuffling the video frames; Sharma et al. [34] mines labels using a distance matrix based on similarity although for video face clustering; Wei et al. [49] predict the ordering task; Ng et al. [30] estimates optical flow while recognizing actions. Self-supervised and unsupervised representation learning is beyond the scope of this paper.

The closest work to ours is by Ray et al. [32]. Ray et al. concatenates pre-trained deep features, learned independently for the different tasks, scenes, object and actions aiming to the recognition, in contrast our HATNet is




Figure 2. Left: Average #annotations per class of each category. Middle: Number of classes of each category over all classes. Right: Number of Annotations of each category over all annotations.




Figure 3. Coverage of 16 different subsets of the 6 main semantic categories in videos. 36.5% of the videos have annotations of all categories.

trained end-to-end for multi-task and multi-label recognition in videos.

**Video Classification Datasets:** Over the last decade, several video classification datasets [3, 4, 24, 25, 37] have been made publicly available with focus on action recognition, summarized in Table 3. We briefly review some of the most influential action datasets available. The HMDB51 [24] and UCF101 [37] datasets are currently the most successful in the field of action recognition. However, they are simply not large enough for training deep ConvNets from scratch. Recently, some large action recognition datasets were introduced, such as ActivityNet [4] and Kinetics [22]. ActivityNet contains 849 hours of video, including 28,000 action instances. Kinetics contains 300k videos spanning 400 human action classes with more than 400 examples for each class. The current experimental strategy is to first pre-train models on these large-scale video datasets [4, 21, 22] from scratch and then fine-tuning them on small-scale datasets [24, 37] to analyze their transfer behavior. In the last year also, a few other action datasets have been introduced with more samples, temporal duration and the diversity of category taxonomy, they are HACS [52], AVA [18], Charades [35] and Something-Something [17]. Other huge datasets such as Sports-1M [21] and YouTube-8M [2]. The

annotations of these datasets are slightly noisy as they are annotated by an automatic tagging algorithm. Furthermore, only video-level labels have been provided, and they are also limited to activity taxonomies with a focus on sports actions only. Due to these reasons, pre-training on these datasets prevents models from providing good training.

Finally, it is worth noting the work on SOA dataset [32] which is in the similar spirit of HVU. SOA is a multi-task and multi-label video dataset aiming to the recognition of different visual concepts, such as scenes, objects and actions. In comparison, HVU has multiple semantic aspects not limited to scenes, objects, actions only, but also including events, attributes, and concepts. Our HVU dataset can help the vision community and bring more attention to holistic video understanding. Further, we invite the community to help to extend this dataset that will spur research in video understanding as a comprehensive, multi-faceted problem. Note that, SOA dataset is not publicly available while we were writing this paper. Our dataset will be made publicly available in the next few months.

Motivated from the efforts to construct large-scale benchmarks for object recognition in static images, i.e. the Large Scale Visual Recognition Challenge (ILSVRC) to learn a generic feature representation is now a back-bone

Task Category	Scene	Object	Action	Event	Attribute	Concept	Total
#Classes	419	2651	877	149	160	122	4378
#Annotations	1, 485, 154	5, 944, 277	1, 552, 920	918, 696	1, 036, 308	965, 077	11, 902, 432
#Videos	366, 941	480, 821	481, 418	320, 428	368, 668	375, 664	481, 418

Table 2. Statistics of the HVU training set for different categories. The category with most number of classes and annotations is the object category.

Dataset	Scene	Object	Action	Event	Attribute	Concept	#Videos	Year
HMDB51 [24]	-	-	51	-	-	-	7K	'11
UCF101 [37]	-	-	101	-	-	-	13K	'12
ActivityNet [4]	-	-	200	-	-	-	20K	'15
AVA [18]	-	-	80	-	-	-	57.6K	'18
Something-Something [17]	-	-	174	-	-	-	108K	'17
HACS [52]	-	-	200	-	-	-	140K	'19
Kinetics [22]	-	-	600	-	-	-	500K	'17
SOA [32]	49	356	148	-	-	-	562K	'18
HVU	419	2651	877	149	160	122	577K	'19

Table 3. Comparison of the HVU dataset with other publicly available video recognition datasets in term of #classes per category. Note that SOA is not publicly available at this moment.

to support several related vision tasks. We are driven by the same spirit towards learning a generic feature representation at the video level for a holistic video understanding.

### 3. HVU Dataset

The Holistic Video Understanding dataset (HVU) is organized hierarchically in a semantic taxonomy of holistic video understanding. Almost all real-world conditioned video datasets are targeting human action recognition. However, a video is not only about an action which provides a human-centric description of the video. By focusing on human-centric description, we ignore the information about scene, objects, events, attributes of the scenes or objects available in the video. While SOA [32] is a multi-task and multi-label data set, which has classes of scenes, objects, and actions, to our knowledge it is not publicly available. Furthermore, HVU has more categories (actions, scenes, objects, events, attributes, and concepts). One of the important research questions which is not addressed well in recent works on action recognition, is leveraging the other contextual information in a video. The HVU dataset makes it possible to assess the effect of learning and knowledge transfer among different tasks, such as enabling the transfer learning of object recognition in videos to action recognition and vice-versa. In summary, HVU can help the vision community and bring more interesting solutions to holistic video understanding. Our dataset focuses on the recognition of scenes, objects, actions, attributes, events, and concepts in the real-world user generated videos.

### 3.1. HVU Statistics

HVU consists of **577k** videos. The number of samples for train, validation, and test splits are reported in Table 1. The dataset consists of trimmed videos clips. In practice, the duration of the videos are different with maximum of 10 seconds length. HVU has 6 main categories: scene, object, action, event, attribute, and concept. In total, there are 4378 classes with approx. 13M annotations for training and validation set. On average, there are  $\sim 2681$  annotations per class. We depict the distribution of categories with respect to number of annotations, classes, and annotations per class in Fig. 2. We can observe that the object category has the highest quota of classes and annotations, which is due to the abundance of objects in video. Despite having the highest quota of the classes and annotations, the object category does not have the highest annotations per class ratio. However, the average number of  $\sim 2,241$  annotations per class is a reasonable amount of training data for each class. The scene category does not have a large amount of classes and annotations which is due to two reasons: the trimmed videos of the dataset and the short duration of the videos. This distribution is somewhat the same for the action category. The dataset statistics for each category is shown in Table 2 for the training set.

### 3.2. Collecting and Annotation

Building a large-scale video understanding dataset is time-consuming task. In practice, there are two main tasks which are usually most time consuming for creating a large-scale video dataset: (a) acquisition or collecting appropriate data, and (b) annotating the data. Recent datasets,



Figure 4. Different classes tend to co-occur semantically in HVU. Here, we visualize t-SNE [27] relationship. This embedding is purely based on class co-occurrence, without using video content.

such as ActivityNet, Kinetics, and YouTube-8M have collected their data from the Internet sources to reduce the effort needed for data collection. For the annotation of these datasets, usually a semi-automatic crowdsourcing strategy is used, in which a human manually verifies the crawled videos from the web. We use a similar strategy with minor changes to reduce the cost of data collection and annotation. Since we target the real-world user generated videos, thanks to the category taxonomy diversity of Youtube8M, Kinetics-600 and HACS, we use these datasets as main source of the HVU. All of the aforementioned datasets are datasets for action recognition. Manually annotating a large number of videos with multiple semantic concepts is not feasible due to the amount of videos, and also the difficulty for a human to pay attention to every detail - which might introduce label noise that is difficult to eradicate. Therefore, we employ a semi-automatic method for annotation. We use the Sensifai Video Tagging API [1] to get rough annotations of the videos, which predicts multiple tags (or class labels) for each video. To make sure that each class has approximately around the same number of samples, we prune the classes which have less than 50 samples. Afterwards, expert human annotators verify the relevance of the tags to their corresponding video for the validation and test set. We plan to verify the training set tags in the new future versions of the HVU. Figure 4 shows the t-SNE [27] visualization of semantically related categories tend to co-occur on HVU.

### 3.3. Human Verification

As mentioned before, the tags of the videos are automatically generated. Therefore, we employ workers to verify the tags. We provide a GUI interface with a video player and the predicted tags for each video with a check-box entry for each tag. The human annotator reviews each video (often multiple times) and determines if it contains the intended classes, and removes (by marking) any irrelevant tags that




Figure 5. Multi-task neural network configuration, applied on HVU dataset.

might introduce label noise. The process takes approx. 100 seconds per clip on average for a trained worker.

### 3.4. Taxonomy

We use the Sensifai video tagging services [1] for annotating the videos. Their video tagging API is trained on internal Sensifai datasets which can recognize videos with ~10K annotations, spanning categories of scenes, objects, events, attributes, concepts, logos, emotions, and actions. As mentioned earlier, we prune tags with imbalanced distribution and finally, refine the tags to get the final taxonomy by using WordNet [28] ontology. The refinement and pruning process was aimed to preserve the true distribution of labels. Finally, we ask the human annotators to classify the tags in to 6 main semantic categories, they are scenes, objects, actions, events, attributes and concepts. Moreover, it is important to note that each video may be assigned to multiple semantic categories. There are 50 different sets of videos based on assigned semantic categories. About 36% of the videos have all of the categories. Figure 3 shows the percentage of the different subsets of the main categories.

### 3.5. Video Recognition Datasets

As covered in related-work, there exists a lot of publicly available benchmarks and datasets which focus on human action recognition. The first ones were KTH [25], HMDB51 [24], and UCF101 [37] that inspired the design of action recognition models based on hand-engineered features. However, they are simply not large enough and have insufficient variation for learning good representative features by training deep ConvNets from scratch. To solve this issue, datasets, such as, Sports1M [21], Kinetics [22], and AVA [18] were recently introduced. Although, such datasets have provided more training data, they are still highly specialized for the action recognition task only. In comparison, HVU focuses on multi-tasks. In Table 3, we compare HVU with the current publicly available video recognition datasets. Note that, SOA [32] shares the same spirit to HVUs, but to the best of our knowledge the SOA dataset is not publicly available at this moment.

## 4. Holistic Appearance and Temporal Network

In this section, we study the state-of-the-art 3D ConvNets for video classification and then describe our






Figure 6. HATNet: A new 2D/3D deep neural network with 2DConv, 3DConv blocks and merge and reduction (M&R) block to fuse 2D and 3D feature maps in intermediate stages of the network. HATNet combines the appearance and temporal cues with the overall goal to compress them into a more compact representation.

new proposed “Holistic Appearance and Temporal Network” (HATNet) for multi-task and multi-label video classification.

#### 4.1. 3D-ConvNets Baselines

3D ConvNets are designed to handle temporal information coming from video clips and have more robust performance for video classification. 3D ConvNets exploit both spatial and temporal information in one pipeline. In this work, we chose 3D-ResNet [41] and STCnet [7] as our 3D CNNs baseline which have competitive results on Kinetics and UCF101 datasets. To measure the performance on the multi-label HVU dataset, we use mean average precision (mAP) over all of the labels. We also report the performance on the category of actions and other classes (objects, scenes, events, attributes, and concepts) separately. The comparison between all of the methods can be found in Table 4. The objective function of these networks is cross entropy loss.

#### 4.2. Multi-Task Learning 3D-ConvNets

Another approach which is studied in this work to tackle the HVU dataset is to have the problem solved with multi-task learning or joint training method. As we know that the HVU dataset consists of high-level classes like objects, scenes, events, attributes, and concepts, so each of these categories can be dealt like separate tasks. In our experiments, we have defined two tasks, (a) action classification and (b) multi-label classification. So our multi-task learning network is trained with two objective functions, that is with single label action classification and multi-label classification for objects, scenes, etc. The basic network is an STCnet which has two separate Conv layer as for the last layer for each of the tasks (see Figure 5). In this experiment, we use ResNet18 as the backbone network for STCnet. The total loss of the training comes as following:

$$\mathcal{L}_{total} = \mathcal{L}_{Action} + \mathcal{L}_{Tagging} \quad (1)$$

For the tagging branch we have cross entropy loss since it is multi-label classification and softmax loss for action recognition branch as single label classification.

#### 4.3. 2D/3D HATNet

Our “Holistic Appearance and Temporal Network” (HATNet) is a spatio-temporal neural network, which extracts temporal and appearance information in a novel way to maximize engagement of the two sources of information and also the efficiency of video recognition. The motivation of proposing this method is deeply rooted in a need of handling different levels of concepts in holistic video recognition. Since we are dealing with still objects, dynamic scenes, different attributes and also different human activities, we need a deep neural network that is able to focus on different levels of semantic information. We propose a flexible method to use a 2D pre-trained model on large image dataset like ImageNet and a 3D pre-trained model on video datasets like Kinetics to fasten the process of training and of course training from scratch is still an option. The proposed HATNet is capable of learning a hierarchy of spatio-temporal feature representation using appearance and temporal neural modules.

**Appearance Neural Module.** In HATNet design, we use 2D ConvNets with 2D Convolutions (2DConv) block to extract static cues of individual frames in a video-clip. Since we aim to recognize objects, scenes and attributes alongside of actions, it is necessary to have this module in the network which can handle these concepts better. Specifically, we use 2DConv to capture the spatial structure in the frame.

**Temporal Neural Module.** In HATNet design, the 3D Convolutions (3DConv) module handles temporal cues dealing with interaction in a batch of frames. 3DConv aims to capture the relative temporal information between frames. It is crucial to have 3D Convolutions in the network to learn relational motion cues for efficiently understanding dynamic scenes and human activities. We use ResNet18 for both of 3D and 2D modules, so that they have the same spatial kernel sizes, and thus we can combine the output of the appearance and temporal branches at any intermediate stages of the network.

Figure 6 shows how we combine the 2DConv and 3DConv branches and use merge and reduction blocks to fuse feature maps in intermediate stages of HATNet. Intu-



Model	Scene	Object	Action	Event	Attribute	Concept	HVU Overall %
3D-ResNet	58.5	38.4	53.2	35.3	32.1	24.4	40.3
3D-STCNet	59.1	38.7	57.1	37.5	33.6	25.7	41.9
HATNet	<b>62</b>	<b>43.4</b>	<b>58.5</b>	<b>41.9</b>	<b>34.5</b>	<b>27.6</b>	<b>44.7</b>

Table 4. Different architecture mAP (%) performance comparison when trained on HVU dataset. The backbone ConvNet for all models is ResNet18.

itively, combining the appearance and temporal features are complementary for video understanding and this fusion step aims to compress them into a more compact and robust representation. In the experiment section, we discuss in more details the HATNet design and how we apply merge and reduction modules between 2D and 3D neural modules. Supported by our extensive experiments, we show that HATNet complements the holistic video recognition, including understanding the dynamic and static aspects of a scene and also human action recognition. In our experiments, we have also performed tests on HATNet based multi-task learning similar to 3D-ConvNets based multi-task learning discussed in Section 4.2.

## 5. Experiments

In this section, we explain the implementation details of our experiments, and then show the performance of each mentioned method on multi-label video recognition on the HVU dataset. We also compare the transfer learning ability between large scale datasets, HVU and Kinetics. Finally we talk about the results of our method and the state-of-the-art methods on three challenging human action and activities datasets. For all of our experiments and comparison, we use

Model	Action	Tags(Object, Scene, etc)
3D-ResNet (Standard)	53.2	27.5
HATNet (Standard)	<b>58.5</b>	<b>30.9</b>
3D-ResNet (Multi-Task)	54.8	29.1
3D-STCNet (Multi-Task)	56.3	30.4
HATNet (Multi-Task)	<b>60.6</b>	<b>31.7</b>

Table 5. Multi-task learning mAP (%) performance comparison of 3D-ResNet18 and HATNet, when trained on HVU: Actions and Tags categories independently. The backbone ConvNet for all models is ResNet18.

Pre-Training Dataset	UCF101	HMDB51	Kinetics
From Scratch	65.2	33.4	65.6
Kinetics	89.8	62.1	-
HVU	<b>91.1</b>	<b>64.2</b>	<b>66.9</b>

Table 6. mAP (%) performance comparison of HVU and Kinetics datasets for transfer learning generalization ability when evaluated on different action recognition dataset. The trained model for all of the datasets is 3D-ResNet18.

RGB frames as input to the ConvNet models. For our proposed methods we either use 16 or 32 frame long video clip as single input to the models for classification. We use the PyTorch for ConvNet implementation and all the networks are trained on 8 V100 NVIDIA GPUs.

### 5.1. 2D/3D HATNet Design

The HATNet includes two branches: first is the 3D-Conv blocks with merging and reduction block and second branch is 2D-Conv blocks. After each of 2D/3D blocks we merge the feature maps from each block and perform a channel reduction, which is done by applying a  $1 \times 1 \times 1$  convolutions. Given the feature maps of the first block of both 2DConv and 3DConv, that be of size 64 channels each. We first merge (or concatenate) these maps, resulting in 128 channels, and then apply  $1 \times 1 \times 1$  convolutions with 64 kernels for channel reduction, resulting in output of 64 channels. The merging and reduction is done in 3D and 2D branches, and continues independently until the last merging with two branches.

We employ 3D-ResNet and STCnet [7] with ResNet18, 50 backbone in our experiments to develop the HATNet. The STCnet is a model of 3D networks with spatio-temporal channel correlation modules which improves 3D networks performance significantly. We also had to make a small change to the 2D branch and remove pooling layer right after the first 2D Conv to maintain a similar feature map size between the 2D and 3D branches since we use  $112 \times 112$  as input resolution size.

### 5.2. HVU Results

In Table 4, we report the overall performance of different baselines, the multi-task learning baseline and also HATNet on the HVU validation set. The reported performance is mean average precision on all of the classes/tags. HATNet that exploits both appearance and temporal information in the same pipeline achieves the best performance, since recognizing objects, scenes and attributes need an appearance module which other baselines do not have. With HATNet, we show that combining the 3D (temporal) and 2D (appearance) convolutional blocks can learn a more robust reasoning ability.

Method	Pre-Trained Dataset	CNN Backbone	UCF101	HMDB51	Kinetics
Two Stream (spatial stream) [36]	Imagenet	VGG-M	73	40.5	-
Conv+LSTM [10]	Imagenet	AlexNet	68.2	-	-
TDD+FV [47]	Imagenet	VGG-M	90.3	63.2	-
RGB-I3D [5]	Imagenet	Inception v1	84.5	49.8	-
TSN [48]	Imagenet	Inception v2	86.4	53.7	-
LTC [43]	Sport1M	VGG11	82.4	48.7	-
C3D [40]	Sport1M	VGG11	82.3	51.6	-
TSN [48]	Imagenet,Kinetics	Inception v3	93.2	-	72.5
RGB-I3D [5]	Imagenet,Kinetics	Inception v1	95.6	74.8	72.1
RGB-I3D [5]	Kinetics	Inception v1	95.6	74.8	71.6
3D ResNet 101 (16 frames) [19]	Kinetics	ResNet101	88.9	61.7	62.8
3D ResNext 101 (16 frames) [19]	Kinetics	ResNext101	90.7	63.8	65.1
STC-ResNext 101 (16 frames) [7]	Kinetics	ResNext101	92.3	65.4	66.2
STC-ResNext 101 (64 frames) [7]	Kinetics	ResNext101	96.5	74.9	68.7
C3D [45]	Kinetics	ResNet18	89.8	62.1	65.6
ARTNet [45]	Kinetics	ResNet18	93.5	67.6	69.2
R(2+1)D [42]	Kinetics	ResNet50	96.8	74.5	72
SlowFast [11]	Kinetics	ResNet50	-	-	75.6
<b>HATNet (16 frames)</b>	Kinetics	ResNet18	94.1	69.2	70.4
<b>3D-ResNet18 (16 frames)</b>	HVU	ResNet18	90.4	65.1	66.9
<b>3D-ResNet18 (32 frames)</b>	HVU	ResNet18	90.9	66.6	67.3
<b>HATNet (16 frames)</b>	HVU	ResNet18	95.4	72.2	71.8
<b>HATNet (32 frames)</b>	HVU	ResNet18	96.9	74.5	73.9
<b>HATNet (16 frames)</b>	HVU	ResNet50	<b>96.5</b>	<b>73.4</b>	<b>74.6</b>
<b>HATNet (32 frames)</b>	HVU	ResNet50	<b>97.7</b>	<b>76.2</b>	<b>76.3</b>

Table 7. State-of-the-art performance comparison on UCF101, HMDB51 test sets and Kinetics validation set. The results on UCF101 and HMDB51 are average mAP over three splits, and for Kinetics is Top-1 mAP on validation set. For a fair comparison, in this table we report the performance of methods which utilize only RGB frames as input.

### 5.3. Multi-Task Learning on HVU

Since the HVU dataset is a multi-task classification dataset, it is interesting to compare the performance of different deep neural networks in the multi-task learning paradigm as well. For this, we have used the same architecture as in the previous experiment, but with a different last layer of convolution to observe multi-task learning performance, see Figure 5. We have targeted two tasks: action classification and Tagging (object, scene, attributes, events and concepts). In Table 5, we have compared standard training without multi-task learning heads versus multi-task learning networks.

The multi-task learning methods achieve higher performance on individual tasks as expected, in comparison to standard networks learning for all classes as a single task. Therefore this initial result on a real-world multi-task video dataset motivates the investigation of more efficient multi-task learning methods for video classification.

### 5.4. Transfer Learning: HVU vs Kinetics

Here, we study the ability of transfer learning with the HVU dataset. We compare the results of pre-training

3D-ResNet18 using Kinetics versus using HVU and then fine-tuning on UCF101, HMDB51 and Kinetics. Obviously, there is a large benefit from pre-training a deep 3D-ConvNets and then fine-tune it on smaller datasets (i.e. HVU, Kinetics  $\Rightarrow$  UCF101 and HMDB51). As it can be observed in Table 6, models pre-trained on our HVU dataset performed notably better than models pre-trained on the Kinetics dataset. Moreover, pre-training on HVU can improve the results on Kinetics also, although it is marginal but still effective.

### 5.5. Comparison on UCF, HMDB, Kinetics

In Table 7, we compare the HATNet performance with the state-of-the-art methods on UCF101, HMDB51 and Kinetics. For our baselines and HATNet, we employ pre-training in two separate setups: one with HVU and another with Kinetics, and then fine-tune on the target datasets. For UCF101 and HMDB51, we report the average accuracy over all three splits. We have used ResNet18,50 as backbone model for all of our networks with 16 and 32 input-frames. HATNet pre-trained on HVU with 32 frames input achieved superior performance on all three datasets even compared to pre-trained models on ImageNet and Kinetics

datasets. Note that for Kinetics dataset, HATNet even with ResNet18 as a backbone ConvNet performs almost comparable to SlowFast which is trained on ResNet50.

## 6. Conclusion

This work presents the “Holistic Video Understanding Dataset” (HVU), a large-scale multi-task, multi-label video benchmark dataset with comprehensive tasks and annotations. HVU contains 557k videos in total with 12M annotations for training set, which is richly labeled over 4378 classes encompassing scenes, objects, actions, events, attributes and concepts categorization. We believe our HVU dataset will complement computer vision in learning generic video representation that will enable many real-world applications. Furthermore, we present a novel network architecture, HATNet, that combines as well 2D and 3D ConvNets in order to learn a robust spatio-temporal feature representation via multi-task and multi-label learning in an end-to-end manner. We believe that our work will inspire new research ideas for holistic video understanding.

**Acknowledgements:** This work was supported by DBOF PhD scholarship, KU Leuven:CAMETRON project, and KIT:DFG-PLUMCOT project. Mohsen Fayyaz and Juergen Gall have been financially supported by the DFG project GA 1927/4-1 (Research Unit FOR 2535) and the ERC Starting Grant ARCA (677650). We also would like to thank Sensifai for giving us access to the Video Tagging API for dataset preparation.

## References

- [1] Sensifai video tagging api. [www.sensifai.com](http://www.sensifai.com).
- [2] S. Abu-El-Haija, N. Kothari, J. Lee, P. Natsev, G. Toderici, B. Varadarajan, and S. Vijayanarasimhan. Youtube-8m: A large-scale video classification benchmark. *arXiv:1609.08675*, 2016.
- [3] M. Andriluka, L. Pishchulin, P. Gehler, and B. Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *CVPR*, 2014.
- [4] F. Caba Heilbron, V. Escorcia, B. Ghanem, and J. Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *CVPR*, 2015.
- [5] J. Carreira and A. Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 2017.
- [6] N. Dalal, B. Triggs, and C. Schmid. Human detection using oriented histograms of flow and appearance. In *ECCV*, 2006.
- [7] A. Diba, M. Fayyaz, V. Sharma, M. M. Arzani, R. Yousefzadeh, J. Gall, and L. Van Gool. Spatio-temporal channel correlation networks for action classification. In *ECCV*, 2018.
- [8] A. Diba, M. Fayyaz, V. Sharma, A. H. Karami, M. M. Arzani, R. Yousefzadeh, and L. Van Gool. Temporal 3d convnets using temporal transition layer. In *CVPR Workshops*, 2018.
- [9] A. Diba, V. Sharma, and L. Van Gool. Deep temporal linear encoding networks. In *CVPR*, 2017.
- [10] J. Donahue, L. Anne Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *CVPR*, 2015.
- [11] C. Feichtenhofer, H. Fan, J. Malik, and K. He. Slowfast networks for video recognition. *arXiv:1812.03982*, 2018.
- [12] C. Feichtenhofer, A. Pinz, and A. Zisserman. Convolutional two-stream network fusion for video action recognition. In *CVPR*, 2016.
- [13] B. Fernando, H. Bilen, E. Gavves, and S. Gould. Self-supervised video representation learning with odd-one-out networks. In *CVPR*, 2017.
- [14] B. Fernando, E. Gavves, J. M. Oramas, A. Ghodrati, and T. Tuytelaars. Modeling video evolution for action recognition. In *CVPR*, 2015.
- [15] A. Gaidon, Z. Harchaoui, and C. Schmid. Temporal localization of actions with actoms. *PAMI*, 2013.
- [16] R. Girdhar, D. Ramanan, A. Gupta, J. Sivic, and B. Russell. Actionvlad: Learning spatio-temporal aggregation for action classification. In *CVPR*, 2017.
- [17] R. Goyal, S. E. Kahou, V. Michalski, J. Materzynska, S. Westphal, H. Kim, V. Haenel, I. Freund, P. Yianilos, M. Mueller-Freitag, et al. The” something something” video database for learning and evaluating visual common sense. In *ICCV*, 2017.
- [18] C. Gu, C. Sun, D. A. Ross, C. Vondrick, C. Pantofaru, Y. Li, S. Vijayanarasimhan, G. Toderici, S. Ricco, R. Sukthankar, C. Schmid, and J. Malik. Ava: A video dataset of spatio-temporally localized atomic visual actions. In *CVPR*, 2018.
- [19] K. Hara, H. Kataoka, and Y. Satoh. Learning spatio-temporal features with 3d residual networks for action recognition. In *ICCV*, 2017.
- [20] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015.
- [21] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei. Large-scale video classification with convolutional neural networks. In *CVPR*, 2014.
- [22] W. Kay, J. Carreira, K. Simonyan, B. Zhang, C. Hillier, S. Vijayanarasimhan, F. Viola, T. Green, T. Back, P. Natsev, M. Suleyman, and A. Zisserman. The kinetics human action video dataset. *arXiv:1705.06950*, 2017.
- [23] A. Klaser, M. Marszałek, and C. Schmid. A spatio-temporal descriptor based on 3d-gradients. In *BMVC*, 2008.
- [24] H. Kuehne, H. Jhuang, R. Stiefelhofen, and T. Serre. Hmdb51: A large video database for human motion recognition. In *High Performance Computing in Science and Engineering*. 2013.
- [25] I. Laptev, B. Caputo, et al. Recognizing human actions: a local svm approach. In *ICPR*, 2004.
- [26] I. Laptev, M. Marszałek, C. Schmid, and B. Rozenfeld. Learning realistic human actions from movies. In *CVPR*, 2008.
- [27] L. v. d. Maaten and G. Hinton. Visualizing data using t-sne. *JMLR*, 2008.

- [28] G. A. Miller. Wordnet: a lexical database for english. *Communications of the ACM*, 1995.
- [29] I. Misra, C. L. Zitnick, and M. Hebert. Shuffle and learn: unsupervised learning using temporal order verification. In *ECCV*, 2016.
- [30] J. Y.-H. Ng, J. Choi, J. Neumann, and L. S. Davis. Action-flownet: Learning motion representation for action recognition. In *WACV*, 2018.
- [31] J. C. Niebles, C.-W. Chen, and L. Fei-Fei. Modeling temporal structure of decomposable motion segments for activity classification. In *ECCV*, 2010.
- [32] J. Ray, H. Wang, D. Tran, Y. Wang, M. Feiszli, L. Torresani, and M. Paluri. Scenes-objects-actions: A multi-task, multi-label video dataset. In *ECCV*, 2018.
- [33] P. Scovanner, S. Ali, and M. Shah. A 3-dimensional sift descriptor and its application to action recognition. In *ACM MM*, 2007.
- [34] V. Sharma, M. Tapaswi, M. S. Sarfraz, and R. Stiefelhagen. Self-supervised learning of face representations for video face clustering. In *International Conference on Automatic Face and Gesture Recognition.*, 2019.
- [35] G. A. Sigurdsson, G. Varol, X. Wang, A. Farhadi, I. Laptev, and A. Gupta. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *ECCV*, 2016.
- [36] K. Simonyan and A. Zisserman. Two-stream convolutional networks for action recognition in videos. In *NIPS*, 2014.
- [37] K. Soomro, A. R. Zamir, and M. Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv:1212.0402*, 2012.
- [38] L. Sun, K. Jia, D.-Y. Yeung, and B. E. Shi. Human action recognition using factorized spatio-temporal convolutional networks. In *ICCV*, 2015.
- [39] P. Tang, X. Wang, B. Shi, X. Bai, W. Liu, and Z. Tu. Deep fishernet for object classification. *arXiv:1608.00182*, 2016.
- [40] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, 2015.
- [41] D. Tran, J. Ray, Z. Shou, S.-F. Chang, and M. Paluri. Convnet architecture search for spatiotemporal feature learning. *arXiv:1708.05038*, 2017.
- [42] D. Tran, H. Wang, L. Torresani, J. Ray, Y. LeCun, and M. Paluri. A closer look at spatiotemporal convolutions for action recognition. In *CVPR*, 2018.
- [43] G. Varol, I. Laptev, and C. Schmid. Long-term temporal convolutions for action recognition. *PAMI*, 2018.
- [44] H. Wang and C. Schmid. Action recognition with improved trajectories. In *ICCV*, 2013.
- [45] L. Wang, W. Li, W. Li, and L. Van Gool. Appearance-and-relation networks for video classification. In *CVPR*, 2018.
- [46] L. Wang, Y. Qiao, and X. Tang. Video action detection with relational dynamic-poselets. In *ECCV*, 2014.
- [47] L. Wang, Y. Qiao, and X. Tang. Action recognition with trajectory-pooled deep-convolutional descriptors. In *CVPR*, 2015.
- [48] 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 *ECCV*, 2016.
- [49] D. Wei, J. Lim, A. Zisserman, and W. T. Freeman. Learning and using the arrow of time. In *CVPR*, 2018.
- [50] G. Willems, T. Tuytelaars, and L. Van Gool. An efficient dense and scale-invariant spatio-temporal interest point detector. In *ECCV*, 2008.
- [51] J. Yue-Hei Ng, M. Hausknecht, S. Vijayanarasimhan, O. Vinyals, R. Monga, and G. Toderici. Beyond short snippets: Deep networks for video classification. In *CVPR*, 2015.
- [52] H. Zhao, Z. Yan, L. Torresani, and A. Torralba. Hacs: Human action clips and segments dataset for recognition and temporal localization. *arXiv:1712.09374*, 2019.