



HAL
open science

3D Convolutional Networks for Action Recognition: Application to Sport Gesture Recognition

Pierre-Etienne Martin, J Benois-Pineau, R Péteri, A Zemmari, J Morlier

► **To cite this version:**

Pierre-Etienne Martin, J Benois-Pineau, R Péteri, A Zemmari, J Morlier. 3D Convolutional Networks for Action Recognition: Application to Sport Gesture Recognition. Multi-faceted Deep Learning, 2021. hal-03639273

HAL Id: hal-03639273

<https://hal.science/hal-03639273>

Submitted on 12 Apr 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

3D Convolutional Networks for Action Recognition: Application to Sport Gesture Recognition

P.-E. Martin¹, J. Benois-Pineau¹, R. Péteri², A. Zemmari¹, and J. Morlier³

Abstract 3D convolutional networks is a good means to perform tasks such as video segmentation into coherent spatio-temporal chunks and classification of them with regard to a target taxonomy. In the chapter we are interested in the classification of continuous video takes with repeatable actions, such as strokes of table tennis. Filmed in a free marker less ecological environment, these videos represent a challenge from both segmentation and classification point of view. The 3D convnets are an efficient tool for solving these problems with window-based approaches.

1 Introduction

Movement is one of the most important aspects of visual perception, and stimuli associated with movement tend to be biologically significant [RMB⁺10]. Psycho-visual experiments show that a complete and exhaustive measurement of a scene is not always necessary to interpret its content. In the animal world, movement information, even partial, may be sufficient to recognize potential food: a frog can distinguish the flight of a fly from a falling leaf without needing to geometrically reconstruct the whole scene.

An illustration of recognition through movement by human beings is provided with the experiments of Johansson (*Moving Light Display experiment* [Joh73]) in which the only source of information on a moving actor is given by bright spots attached to a few joints. People shown a static image can only see meaningless dot patterns. However, as soon as they are shown

¹LaBRI, Univ. Bordeaux, 351 cours de la Libération, F33405 Talence Cedex, France.

²MIA, Univ. La Rochelle, La Rochelle, France

³IMS, Univ. Bordeaux, France

e-mail: {pierre-etienne.martin,benois-p,zemmari}@u-bordeaux.fr, renaud.peteri@univ-lr.fr

the entire sequence of images, they can recognize characteristic actions such as running or walking, and even the male or female gender of the actors. Such abilities suggest that it is possible to use movement as a means of recognition by itself.

In the field of computer vision, describing a video scene through natural language involves the narration of some key events or actions occurring in the scene. Movement analysis is obviously a central element in order to be able to take full advantage of the temporal coherence of the image sequences. Moreover, [Lap13] states that on average, on each video from the web, 35% of the pixels represent humans. Describing an image sequence from the human movements and activities performed is hence potentially discriminant and relevant for analyzing or indexing videos. These different elements, added to the complexity of the task, highlight why human action recognition in videos has been in past years a very active research topic in computer vision. Nevertheless, as shows the already large history of research for solving the problem of action recognition in video, when a real-world video scenes have to be analysed [SFC16], motion characteristics alone are not sufficient. This is why the approaches using both temporal information expressed in terms of motion characteristics, such as velocity field, i.e. Optical Flow (OF) and spatial characteristics derived from colour pixel values in video frames have shown better performances. This is hold both for methods on the basis of hand-crafted features or the ever-winning Convolutional Neural Networks (CNNs).

Focusing on Deep Learning approaches, it is nevertheless interesting to show the problem in its historical perspective for better predicting the future. In this chapter, we will briefly present approaches for action recognition with different features going from handcrafted to produced by Deep NNs. We will speak about the evolution of datasets for development and testing action recognition methods, and introduce recently created dataset in our research for fine-grained classification of Table Tennis strokes. We will also present recent contributions in fine -grained action recognition with Twin Spatio-temporal networks.

2 Highlights on Action Recognition Problem

The problem of recognition of human actions in video has a wide-range of applications, this is why the history of research is quite long-term one. From historical perspective one can distinguish two main approaches with progressive combination of them. The first one consists in designing the so-called "handcrafted" features, expressing characteristics of video frames in a local or holistic manner thus forming a new description space. Classification of features in this new space with machine learning approaches brings the solution to the action recognition problem. The second approach consists in the

"end-to-end" solution with Deep Neural Networks which extract features and then classify them.

2.1 Action classification from videos with handcrafted features

Handcrafted features extracted from videos development started to our best knowledge from feature extraction from images. Efforts were afterwards made for extracting information from the temporal domain. Such features were mainly used in the action recognition task, but also for other tasks such as scene and event recognition from videos. Most of the approaches using handcrafted features seek for their compact representation. The model of Bag of Words (BoW) or Bag-of-Visual Words (BoVW) [CDF⁺04] was introduced which allowed for quantizing a large amount of feature vectors into a set of classes - words of a dictionary of a predefined size. The classes-words of the dictionary were built by statistical clustering methods such as K-means [Mac67] or more sophisticated vector quantization techniques. This final descriptor of the images, areas in them or video frames were thus a statistical model - a histogram of class - occurrence of the "words" in the image. This model, for a long period was used for action recognition, and scene classification in video.

The use of temporal dimension of video with regard to static images was first introduced in 2003 with Spatio-Temporal Interest Points (STIP) features [SLC04]. They were an extension of the 2D corner detector to the temporal dimension for video. The equivalent to image corners in video are points which change direction over space and time. The authors of [SLC04] show that their descriptor matches with the action performed, meaning that STIPs will be located where and when the action happens. They apply their method on their newly created KTH dataset, which became one of the first widely used action datasets. Comparison is done using Support Vector Machine (SVM) classifier and Nearest Neighbor Classification (NNC) on BoW of their spatio-temporal local features. It results in a good classification score for actions which are not similar; but scores on similar actions such as "Walking", "Jogging" and "Running" remained low. The motion in those actions is very similar and the STIPs are concentrated on the same body parts.

The authors of [LP07] use jointly the Histogram of Oriented Gradients (HOG) descriptor and Motion Boundary Histogram (MBH) descriptors, using AdaBoost algorithm for recognition of two actions - smoking and drinking in their own dataset based on movies. At the same time, [GBS⁺07] introduce Space-Time Shapes (STS) features for action classification, along with the new Weizmann action dataset. Classification of STS is based on the same ideas as image shapes classification using Poisson equation. The authors classify the computed shape with Nearest Neighbor Clustering with euclidean

distance. They reach an accuracy of 97.8% on their dataset with however a low confidence on similar actions. It is important to stress that their dataset is acquired in a controlled environment and has the same complexity as the KTH dataset.

In [SG08], the authors concatenate Gabor filters features and OF features in order to perform action classification with SVM classifier on KTH and Weizmann datasets. They show the superiority of their model on both datasets, compared to other methods, and come to the conclusion that one frame is enough to get a satisfactory classification score. Indeed, using only one frame and the posture of the person, actions in both datasets are easily distinguishable.

In 2009, action recognition methods were already reaching very high accuracies on both KTH and Weizmann datasets and the introduction of UCF11 dataset gave more space for improvements. Indeed, this dataset is more challenging since it is recorded “in the wild”, that is in natural conditions, under the constraints of camera motion and flickering for example. The UCF101 [SZS12] samples are extracted from the YouTube platform. Along with their dataset, the authors propose a method for classification based on motion features and static features. They use AdaBoost learning method on the histogram-based representation and compare it with k-means clustering method. AdaBoost leads to better results and the hybrid combination resulted in 93.8% of accuracy on KTH dataset against 71.2% on UCF11 showing the higher complexity of the task for such dataset with the same number of classes.

The authors of [CRHV09] introduce at the same moment Histogram of Oriented optical Flow (HOF) features and reach 94.4% of accuracy on Weizmann actions dataset. The method is simple and easy to reproduce and will be used later on in [WKSL11], along with MBH, HOG features to compute dense trajectory features on the basis of dense optical flow field.

In [WS13], the authors improve dense trajectory features by considering camera motion. Camera motion is estimated using dense optical flow and Speeded Up Robust Features (SURF) descriptors. An homography is estimated using Random Sample Consensus (RANSAC) algorithm. The Improved Dense Trajectories (IDT) yield 91.2% of accuracy against 88.6% with regard to the original dense trajectory features. This work is used later in many applications such as action localization. The latter problem consists in defining not only temporal boundaries of actions in the video, but also spatial locus of them.

In [GHS13], the authors redefine actions as “actoms”. Actom is a short atomic action with discriminative visual information, such as opening a door. It is therefore useful for action localization but can also be applied to classification-by-localization. The definition of actoms is important in the field of action recognition to decompose an action in individual parts. Actoms thus can be present across different actions such as entering or leaving a room with “opening door” as an actom, for example. Their understanding can lead

to better video representation and accordingly to a better classification. However, a too great number of actoms might lead to the situation when they are well presented in the training set, but they would not be present in the test set and be unrelated to the action performed. The number of actoms to consider becomes then a variable to control according to the complexity of the actions to classify.

Another way to perform action recognition is developed by [JvGJ⁺14] who introduce the concept of Tubelets. It is a sampling method to produce 2D+T sequences of bounding boxes where the action is localized. This method, which tackles the localization and classification problem of actions at the same time, is based on super voxel generation through an iterative process using color, texture and motion to finally create tubelets. Those tubelets are then described by MBH features, and one BoW per class method is used for classification. The classifier with the maximum score assigns the class to the tubelet.

2.2 *The move to DNNs in action recognition problem*

Deep Neural Networks have very quickly outperformed all handcrafted feature-based methods, due to the strong generalization capacity of these classifiers. The method described in [JXYY13], was one of the first to use Deep Learning via a 3D CNN for action recognition, but they did not obtain better results than the state-of-the-art methods on the KTH dataset. It is only from 2014 and the innovative *two-stream network* approach of [SZ14] that temporal coherence will be exploited in CNN and that Deep Learning approaches will begin to supplant other methods. In the development of Deep Learning Methods for Action Recognition we could observe two trends: Deep Convolutional Neural Networks (CNNs) and recurrent neural networks (RNN) such as LSTM briefly presented in Chapter?? (??). Nevertheless, according to [VLS18] and the own experience of the authors, these networks are more difficult to train than 3D CNNs integrating spatial information along the time dimension in video. They have also difficulties to handle long term temporal interactions. Hence in this chapter we will not focus on them.

Better performances of Deep Neural Networks (DNNs) with inherent feature extraction from raw video in the end-to-end training and generalization process does not mean that engineered features have to disappear. On the contrary, human understanding of visual scenes influences the choice of the designed features such as e.g. OF. Still recent work may use handcrafted features as a baseline or fuse them with the deep features extracted by a DNN. [BGS⁺17] confirm that SVM does not perform better than a partially retrained Deep Convolutional Neural Network (DCNN) and that the learned features lead to better results than engineered ones; however the fusion outperforms the single modalities. These works confirm the findings of the com-

munity: the fusion of multiple features allows improving recognition scores for complex visual or multi-modal content understanding [IBPQ14]. Action recognition with Deep CNNs was first fulfilled with 2D CNN architectures. Here we will briefly discuss some of them.

2.3 2D Convolutional neural Networks for action classification

2D convolution refers to the fact that convolutions are performed on a 2D spatial support of the image. For RGB data, 2D convolution actually uses 3D kernels to weight each color channel differently. However the way that the kernel will move in the image will only be in 2 dimensions.

In the scope of action recognition, it is what [SZ14] perform. They introduce a Two-Stream Convolutional Network which takes one single RGB frame for one stream, and for the other stream, several frames of the computed OF. Each stream is respectively called “Spatial stream ConvNet” and “Temporal stream ConvNet”.

They notice that “Temporal stream ConvNet” reaches much better performances compared to “Spatial stream ConvNet”. This could easily be explained, as in addition to the dynamic genre of the input data, the temporal stream uses up to ten frames against only one for the spatial stream branch. Performances are much more alike when the temporal branch uses only one frame. Of course, the fusion of the two streams using a SVM method performs the best.

An action can be considered as a volume, i.e. tube in the video. This concept of tube - which expresses a homogenous content, a singular action in our case, is used in [LWLW16]. A Tube Convolutional Neural Network (T-CNN) is introduced which can be sequentially decomposed into two distinct networks. They first create motion-segmented tubes using a Residual Convolutional Neural Network (R-CNN) [RHGS17]. Then those tubes feed a VGG-like network using 20 motion amplitude frames distributed along the channel dimension.

Until now we were speaking about CNNs. Nevertheless, temporal (recurrent) neural networks also represent an alternative to CNNs with windowing approaches. The focus of our work is on CNNs hence we will just very briefly mention them.

Long-term Recurrent Convolutional Network (LRCN) models are introduced in [DHR⁺17] for action recognition. The authors extract features using 2D CNN for each image which feed a LSTM from start to end. The decision is based on the average score. This simple model is tested with OF and a single RGB image. The fusion of the two modalities perform obviously the best.

Numerous works use models based on temporal networks such as Recurrent Neural Network (RNN) and Long Short-Term Memory (LSTM) [UAM⁺18]. However, RNN may be harder to train depending on the application context. Besides, LSTMs are more efficient when they are coupled to the output of a CNN [NHV⁺15].

Now coming back to the convolutional NNs for action recognition, we can satate that 2D CNNs are often used as feature extractors and 3D convolutions are performed on the extracted features. Thus the temporal dimension is taken into account. This leads us to now focus on 3D CNN based methods for action recognition.

2.4 From 2D to 3D ConvNets in Action Classification

We can consider videos as 3D data with the third dimension along the time axis and either process them similarly to 2D images, or treat the temporal dimension differently, or extract temporal information such as dynamic data that can feed a DNN.

However, most methods need to consider extra information, which obviously leads to larger networks, greater number of parameters and the need of a greater number of GPUs with stronger capacities. This might not be possible for every research team, and brought some of them to try attaining accurate results with restrictions on the model size or computation time. This aspect thus brings many shades in the performances, and has brought to light many different methods which shall not be compared only in terms of performances, but also by their means to achieve them. In addition to such limitations, the choice of the architecture for a specific task remained open, leading to numerous implementation attempts.

3D convolutional neural networks are a good alternative as well for capturing long-term dependencies [CZ17], and involve 3D convolutions in space and time. When doing 3D convolution with C_{in} channels on a signal of depth D , width W and height H , the output value of the layer with input size (C_{in}, D, H, W) and output $(C_{out}, D_{out}, H_{out}, W_{out})$ can be precisely described as:

$$out(C_{out_j}) = bias(C_{out_j}) + \sum_{k=0}^{C_{in}-1} weight(C_{out_j}, k) \star input(k) \quad (1)$$

where C_{out_j} is the j^{th} output channel, and \star is the valid 3D cross-correlation operator.

2.5 3D Convolutional neural Networks for action classification

The Convolutional 3D (C3D) model [TBF⁺15] consists of eight consecutive convolutional layers using $3 \times 3 \times 3$ kernels and 5 max-pooling layers. In [HCS17], they use this model in a two stream T-CNN. Here videos are first divided into clips of equal length and are segmented using 3D R-CNN to create tube proposals. Tubes are then classified and linked together. By using the features extracted from the segmented video tubes with C3D model, they increase the performance compared to a direct application of the C3D model.

The authors of [FPW17] extend what was done in 2D by [SZ14], into 3D to introduce their Spatio-Temporal ResNet (ST-ResNet). They replace simple CNN branches by R-CNN with one connection between the two branches. Their results prove that RGB stream processed alone gets better performances than the OF stream. By processing them together, they reach an accuracy of 93.4% on UCF101 dataset.

A major breakthrough was proposed by the method of [CZ17], with much higher scores obtained on action classification. They present their Two-Stream I3D model as the combination of RGB-I3D and Flow-I3D models trained separately. Each of their models uses inflated inception modules, inspired from the 2D inception modules [SLJ⁺15]. The major strength of their model is the pretraining on ImageNet first and then on Kinetics-400 dataset [KCS⁺17], more complex than UCF101 with 400 classes. By using Kinetics, they boost performances from 93.4% to 98% on UCF101. They reach also 74.2% of accuracy on Kinetics-400 dataset. Inception modules have already proven their efficiency for image classification on ImageNet dataset and thus have been successfully used in the I3D models.

Long-term Temporal Convolutions (LTC) CNN were introduced by [VLS18]. They experiment different temporal sizes for input video clips, in order to improve classification. Better accuracies are obtained when considering a greater number of frames as input, especially on long-lasting actions which have a longer temporal support.

[WGW⁺18] introduce Spatial-Temporal Pyramid Pooling Layer (STPP) using 3D convolutions in a two-stream like network fed by RGB and OF streams. The output becomes the input of a LSTM network. The use of LSTM allows classification of videos of arbitrary size and length. Each modality performs similarly: 85% and 83.8% of accuracy for RGB and OF stream respectively. When fused together, the method reaches 92.6% of accuracy.

[KKA20] introduce the Bidirectional Encoder Representations from Transformers (BERT) layer to better make use of the temporal information of BERT's attention mechanism firstly used for language understanding [VSP⁺17]. The BERT layer is based on the use of the Multi-Head Attention layer, which comprises a Scaled Dot Product layer. The Multi-Head attention layer is a part of a bigger network, the Transformer model which is dedicated to

translation tasks. The incorporation of the BERT layer in the REsNeXT, R(2+1)D and I3D models, previously described, improve their performances. They reach the state-of-the-art results on UCF101 dataset with 98.7% of accuracy using the R(2+1)D architecture [TWT⁺18]. It is a ResNet-type architecture with separable temporal and spatial convolutions and a final BERT layer in order to better use the obtained features. One important point to stress is also the use of IG65M dataset [GTM19] for pre-training the model. IG65M dataset is build from the Kinetics-400 [KCS⁺17] class names. Those class names are then used as hashtags on Instagram and lead to 65M clips from 400 classes. Their dataset is however not publicly available.

2.6 Video understanding for racket sports

Our interest is fine-grained action recognition with application in table tennis. We therefore present methods focusing on video classification and/or segmentation in the domain of racket sports.

[EBMM03] propose a motion descriptor based on optical flow in order to classify actions in sports. For this purpose they consider 3 different datasets: Ballet, Football and Tennis datasets. They track the player (or the person performing the action) and build a 3D volume based on their motion. Their motion descriptor has 4 channels: the positive and negative values for horizontal and vertical motions.

Then classification is performed following a nearest neighbour approach using similarity metric:

$$S = \sum_k C_1(k)C_2(k) \quad (2)$$

with C_1 and C_2 being two cuboid samples based on the motion descriptors at coordinate k .

Even if Ballet dataset and Tennis dataset are acquired in a controlled environment, performances for the Tennis dataset are more limited. Football dataset comes from broadcast source which explains the limited performances. Moreover, the number of classes for the Tennis dataset is lesser than the two others, however, it is where their method is the less efficient. This underlines the greater complexity of racket sport and their fine-grained aspect.

Another research field in video classification aims at identifying the different parts of tennis broadcasting. To do so HMMs are applied to tennis action recognition by [KGOG06]. Their model is statistic and integrate the structure of tennis match. They combine audio and key frame features to be able to segment, with a good accuracy, the different parts of the tennis broadcasting such as the first serves, rallies, replays and breaks. The sound of the crowd such as applause, the sound of the ball or the commentator speech

combined with key frames which capture visual information lead to 86% of segmentation accuracy compared to 65% and 77% with only respectively visual features and audio features. Such applications are interesting for sport coaches who wish to comment and examine only sequences of sports.

[dCBM⁺11] present a new dataset for tennis actions. This one contains only three types of classes: “hit”, “serve” and “non-hit-class”. The dataset is build from TV broadcasts of tennis games (matches of females in the Australian Open championships). They are interested in action localization and their classification. To do so, they introduce a local BoW method on the Spatio-Temporal gradients HOG3D features which are an extension of the classical 2D HOG features in 3 dimensions. They also use STS features. Both features are from the located actor and classification is performed using Fisher discriminant analysis. They obtain an accuracy of 77.6% using STS model based. Their confusion matrix is represented in figure 1.

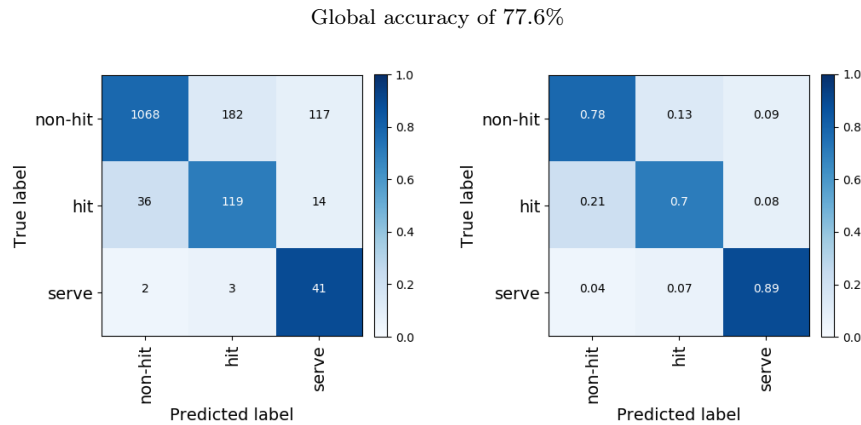


Fig. 1 Confusion matrix of tennis game from TV broadcast [dCBM⁺11]. Left with number of samples and right normalized.

One can see that “serve” samples are easier to classify than “hit” or “non-hit” samples. This is certainly due to the time that a service takes and its decomposition in time, which starts by large movement of the player when launching the ball. Hit and non-hit classes are then harder to distinguish because the hit class is very limited in time. Looking only at the player shape, the ball might not be visible, and features might look the same as when the player is simply moving in the field.

Recently, deeply related to our domain, [CPM19] use the OF Singularities with BoW and SVM in order to classify very similar actions. This task is also called fine-grained action classification. They apply their method on the TTStroke-21 dataset which contains 20 different strokes and a negative class. Their method is inspired from [BPM14] which uses the trajectory of

critical points for classification. In this case, the actions to recognize are the different types of strokes performed during table tennis training session. However, the scores remain low due to the high similarity of the different strokes and the limited amount of video samples. It makes generalization of extracted features harder.

Table Tennis stroke recognition is also performed by [LWS⁺19]. It is based on their Body Sensor Network (BSN). Their sensors collect acceleration and angular velocity information from the upper arm, lower arm and the back of the player. From the recorded signals, they extract Principal Component Analysis (PCA) features which are then fed to a SVM. They reach an accuracy of 97.4%, however they use only 5 classes: “forehand drive”, “block shot”, “forehand chop”, “backhand chop” and “smash”. Similarly, [XWX⁺20] recently proposed classification from integrated wearable sensors using K-means and DBScan clustering. Their taxonomy is more limited than in TTStroke-21 since they use only 9 classes across badminton and table tennis sports. They reach an accuracy of 86.3% when considering all the classes. This score reaches 92.5% when considering only table tennis but this classification is limited to 4 classes: “Service”, “Stroke”, “Spin” and “Picking up”. The extent of their taxonomy is thus limited and does not contribute much to the player experience. Furthermore, using such sensors, strongly limits the application possibilities and has a greater cost regarding training equipment adaptation. Also their system does not offer visualisation of the stroke performed since it is based on sensors, and it limits the feedback for the player.

Recently, a method is introduced by [WZD⁺20] to get the tactics of the players based on their performance in past matches. Their model is based on Hidden Markov Model (HMM) and aims at characterizing and simulating the competition process in table tennis. They use richer taxonomy and terms of stroke techniques than the previously presented methods : 13 different classes and 4 player positions which can be combined. Compared to TTStroke-21, we consider 10 classes with 2 player positions: “Forehand” and “Backhand”. Their goal is therefore, not to classify an input, but to simulate matches between two different players. It is not directly linked to action recognition methods, but it does give a tool for players to simulate sport encounters and give credits to the TTStroke-21 dataset which propose much richer taxonomy than previous datasets.

Thus having analysed the two kinds of approaches: with handcrafted features and with Deep Neural Networks we can state the following. Despite the use of temporal information coming from OF and derived features, the methods with handcrafted features allow a good classification on datasets that remain simple: either with a low number of classes or with classes that are easily separable. It does become more complicated when the task focuses on one particular sport with different actions within or in case, when the complexity of “in-the-wild” scenes is higher. Here, already earlier research work have given a direction to follow: the use of Deep Neural Networks. Before

we afford our solutions, it is interesting to review existing datasets which are used by the community for action recognition in video.

3 Datasets for action recognition

The need of datasets for action recognition has grown those last years. These datasets can change in terms of number of videos starting from a few videos up to millions of them. In addition to their size, the number of categories and their complexity also vary from few classes up to few hundreds, or even thousands of them. Each dataset can be labelled with annotations either by enriching the terminology, localising the action in space and time or by adding modalities information such as skeleton.

3.1 *Annotation processes*

One can distinguish two ways to annotate a dataset: automatic or "by hand", that is by a human operator. Label propagation techniques were an attempt to leverage the gap between fully "manual" annotation and automatic one [HOSS13]. Here the dataset is represented as a big graph with samples-nodes and associated similarity metric between them. The label propagation is fulfilled from manually annotated nodes to unlabelled ones by optimal search on graphs. Nevertheless, the practices in annotation of video datasets are such that, the automatic annotation by a concept recognition method at hand is a first step and then human intervention is required to filter out automatic annotation errors. Also, it is common to have a dataset split into "auto" and "clean" sets. The "auto" being the one annotated using automatic methods and "clean" the one automatically labeled, verified and adjusted by hand. The two annotation processes are first described before presentation of the datasets.

3.1.1 Automatic annotation

Tags from social network platforms can be used to collect rich datasets. It is the case, for example, with the IG65M dataset [GTM19]. To collect it, the authors use Kinetics-400 [KCS⁺17] dataset class names as hashtags on Instagram and collect 65M clips from 400 classes. Such annotation process requires filtering in order to refine the annotations.

In movies, the script can also help to automatically label the sequence. It is for example what [MLS09] have done for the Hollywood2 dataset. They generate the samples this way and clean them manually for the test set.

Similarly, datasets can be constructed from the description of the videos from online platforms hosting them. Then, according to a description, datasets can be generated in an automatic way.

3.1.2 Manual annotation

The most common way to annotate a dataset, especially when it is not large, is to label all the samples by hand. Some tools might be used to help in the process such as a pre-classification if a model already exists, or localization segment candidates of the actions when the video is untrimmed. One can distinguish two ways in the hand-labelling process: if the annotation is done by one person or several. By one person, the risk is that an inattention might lead to errors in the dataset or make it biased according to the point of view of the annotator. To overcome this issue, a crowdsourcing method can be used.

Crowdsourcing is based on the annotation of the same segments by different persons. It relies on the collective intelligence and should give better results than with only one person annotating. Outliers in annotations are not considered in the final decision. Different rules might apply, e.g take the mean of the annotators when possible or consider only the annotator that performed the best today; in order to take an annotation decision. There are also datasets which provide gross crowdsourced annotations and it is the team working on the dataset that decides which decision to make.

A new trend appeared recently: the use of Amazon Mechanical Turk (AMT) ¹. AMT, also called “MTurk”, is a crowdsourcing marketplace for individuals and businesses to outsource their processes and jobs to a distributed workforce who can perform these *micro-tasks* virtually. Here it is applied to annotation and AMT are paid according to the number of annotations performed. It started to be used with ImageNet dataset dedicated to image classification with 3.2 millions of images over 5 247 classes. The strategy is two folds:

AMT workers verify the presence of the action in the video candidates and they can also temporally annotate them. It is often coupled with a crowdsourced method meaning that each video will be annotated by several AMT workers. This allows the construction of a large dataset in a short amount of time.

¹ <https://www.mturk.com/>

3.2 Datasets for action classification

Datasets of actions classification problem can be categorized in many ways. In this section, the datasets are grouped according to the acquisition process: acquired in a controlled environment, extracted from movies or recorded “in the wild”. Obviously, such a categorization is not perfect since some datasets mix different types of videos. Without being exhaustive we will present some of them focusing mainly on sport video datasets.

3.2.1 The acquisition-controlled datasets

These datasets are often self made by the authors who decide in what type of environment the actions will be performed. It does not always mean that the dataset is easier than “in the wild” since difficulties can be added on purpose. The databases from broadcasts or recordings not meant for action recognition task are also considered in this subsection because the acquisition environment can be taken into account in the classification process.

KTH and Weizmann datasets. These two datasets were the most popular at the early ages of action recognition research. Despite their simplicity, some researchers are continuing using them as a benchmark.

KTH introduced in [SLC04]² stands for “Kungliga Tekniska Högskolan” in Swedish which is the Royal Institute of Technology (Stockholm, Sweden), institution of the authors. The dataset is composed of 6 classes: “Walking”, “Jogging”, “Running”, “Boxing”, “Handwaving” and “Hand clapping”. The acquisition was done in a controlled environment, with a homogeneous background, static camera at 25 fps, with 25 actors and has 2 391 video clips across 600 videos. Videos are recorded outdoors and indoors. Weizmann dataset³ [GBS⁺07] is quite similar but is enriched with more actions, e.g. ‘jumping- jack’, “galloping-sideways”... increasing the number of action classes to 9. It is constituted of 81 video sequences recorded at 25 fps at low resolution (180×144).

The necessity of action recognition "in-the-wild" yielded production of much more complex datasets described below.

ACASVA. [dCBM⁺11] “Adaptive Cognition for Automated Sports Video Annotation” (ACSVA) introduce a tennis action dataset ⁴. Their objectives is to evaluate classical action recognition approaches with regard to player action recognition in tennis games. The data are collected from tennis TV broadcasts. The videos are then spatially segmented on the players and tem-

² <https://www.csc.kth.se/cvap/actions/>

³ <http://www.wisdom.weizmann.ac.il/vision/SpaceTimeActions.html>

⁴ <https://www.cvssp.org/acasva/Downloads.html>

porally annotated using 3 classes: “hit”, “serve” and “non-hit”. The complexity of the the dataset remains simple.

FineGym dataset. The FineGym dataset⁵ [SZDL20], is a fine-grained action dataset with a special focus on gym sport. The authors use a rich taxonomy to decompose each actom of structured gymnastic figures. They use three-level semantics and analyse four different gymnastic routines: balance-beam, uneven-bars, vault and floor exercise. They have a total of 530 element categories but only 354 have at least one instance. This rich amount of categories is due to all the combinations of possible actoms. The authors offer two settings: Gym288 with 288 classes but of very unbalanced distribution and Gym99, more balanced but with “only” 99 classes. The total number of samples considering all classes reaches 32 697. The 708 hours of videos are hosted on YouTube with most of them in high resolution.

TUHAD. The TUHAD dataset [LJ20] is also a dataset dedicated to fine-grained action recognition on Taekwondo sport. The dataset was recorded with the help of 10 Taekwondo experts using two Kinect cameras with front and side view. The number of classes is low: with only 8 Taekwondo moves. 1 936 action samples are recorded with depth and IR images along with the RGB data. The classes are very similar in many ways but a foot position, which might be overcome with proper features.

3.2.2 Datasets from movies

Despite these datasets do not generally contain sport actions, they are also interesting as they comprise recordings of natural behaviour of actors in cluttered environments. Thus Drinking and Smoking Dataset [LP07] is composed of sequences from Jim Jarmush Movie “Coffee and Cigarettes”. It was designed for joint action detection and classification for the 2 classes: “smooking” and “drinking” with respectively 141 and 105 samples.

The Hollywood2 dataset⁶ [MLS09] was designed for action and scene classification. It contains 10 scene classes and 12 actions: “Answer phone”, “Drive car”, “Eat”, “Fight person”, “Get out car”, “Hand shake”, “Hug person”, “Kiss”, “Run”, “Sit down”, “Sit up” and “Stand up”; over 7 hours of video from 69 movies. They have a total of 1 694 actions samples. The difficulty lies in the fact that different actions can happen in the same sequence.

⁵ <https://sdolivia.github.io/FineGym/>

⁶ <https://www.di.ens.fr/~laptev/actions/hollywood2/>

3.2.3 In-the-wild datasets

“In-the-wild” means that the videos are from different sources and can be recorded by professionals or amateurs. They thus may contain camera motion, strong blur, occlusions... everything that can make the action recognition task harder. However, videos can also contain much background information, which might be an exploitable source for training of classification models.

The UCF datasets. The UCF datasets [SZS12] have become very popular for developing and benchmarking methods for action recognition in sport video. UCF title comes from the name of the university in which the datasets have been developed: University of Central Florida.

The first UCF dataset was UCF-Sports. It contains various sequences from broadcast TV channels across 9 different sports: “diving”, “golf swinging”, “kicking”, “lifting”, “horseback riding”, “running”, “skating”, “swinging a baseball bat”, and “pole vaulting”. Pole vaulting is split in 2 classes: “Swing-Bench” and “Swing-Side” totaling 10 classes. It first contained 200 sequences (reduced to 150 later) with an image resolution of 740×480 at 10 fps.

Later, the UCF YouTube Action also called UCF11 dataset⁷ is introduced. It consists of 11 classes from 1 160 videos from the online video platform YouTube.

UCF50 is an extension of UCF11 with a total of 50 action classes. This version is then extended to make UCF101 dataset⁸. UCF101 includes a total number of 101 action classes which can be divided into five domains: “Human-Object Interaction”, “Body-Motion Only”, “Human-Human Interaction”, “Playing Musical Instruments” and “Sports”. Constructed from 2500 videos “in-the-wild”, they extract a total of 13 320 clips in order to have at least 101 clips per class. The dataset is widely used by the scientific community and led to the THUMOS challenge⁹ held in 2013, 2014 and 2015. The dataset was cleaned and enriched with temporal annotations in 2015 in order to provide qualitative benchmark for different methods and be used also for spatio-temporal localization and temporal detection only.

The Kinetics datasets. The kinetics datasets¹⁰: Kinetics-400 [KCS⁺17], Kinetics-600 and Kinetics-700 [CNHZ19] consider respectively 400, 600 and 700 action classes. They are all financed by DeepMind company, specialized in AI which, from 2014, belongs to Google. In the taxonomy of actions they contain, we find sport actions as well, e.g. “playing squash or racquetball”.

The videos are collected from YouTube video platform, automatically annotated and candidates are refined using AMT. The difference between the versions of the datasets lies in:

⁷ www.crcv.ucf.edu/data/UCF_YouTube_Action.php

⁸ www.crcv.ucf.edu/data/UCF101.php

⁹ www.thumos.info

¹⁰ <https://deepmind.com/research/open-source/kinetics>

- the number of classes: the number of classes has increased over time. New classes were added and pre-existing classes were refined. Some were merged.
- the amount of videos: the number of videos started with 306 245 clips and more than doubled in the last version
- the splits between the different sets: training, validation and test sets have been modified over time. For example, samples belonging to the training set in the first version might belong to the test set in the last version.

AVA and AVA-kinetics. In [GSR⁺18], the authors introduced the AVA dataset¹¹ in order to perform joint localization and classification of actions. It contains 437 videos gathered from YouTube, 15 minutes are extracted from them and annotated every second. They use a vocabulary of 80 atomic actions. The difficulty in this dataset is the overlapping actions in time and their localization. They offer a split of the dataset by extracting 900 video segments of 3 seconds from all the 15 minutes videos. By doing so, the 55 hours of video are split in 392 000 overlapping segments. The AVA-Kinetics dataset [LTR⁺20]¹². Is the merger of the two: Kinetics-700 and AVA datasets. Kinetics videos were annotated using AVA protocol. The dataset thus contains over 230 000 video clips spatially and temporally annotated using the 80 AVA action classes.

SAR4. [FVD⁺19] present the SAR4 dataset which focuses on action in football sport (or soccer). They track and label the players from available videos on YouTube. The actions performed by the tracked players are then annotated using a taxonomy of 4 classes: “dive”, “shoot”, “pass received” and “pass given”. The total number of sequences is 1 292 with actions lasting from 5 up to 59 frames. The discussed datasets are summarized in table 1

Table 1 Presentation of the different action datasets in terms of number of classes, acquisition process, the amount of videos and the number of extracted clips.

Datasets	# classes	Acquisition	# videos	# clips
KTH [SLC04]	6	Controlled	600	2 391
Weizmann [GBS ⁺ 07]	9	Controlled	-	81
Coffee and Cigarettes [LP07]	2	Film	1	246
UCFSports	10	Broadcast	-	150
UCF11	11	In the wild	-	1 160
Hollywood2 [MLS09]	12	Films	69	1 694
UCF101 [SZS12]	101	In the wild	2 500	13 320
AVA [GSR ⁺ 18]	80	In the wild	437	392 000
SAR4 [FVD ⁺ 19]	4	Broadcast	-	1 292
Kinetics-700 [CNHZ19]	700	In the wild	-	650 317
FineGym [SZDL20]	354	Broadcast	-	32 697
AVA-Kinetics [LTR ⁺ 20]	80	In the wild	-	230 000

¹¹ <https://research.google.com/ava/>

¹² <https://deepmind.com/research/open-source/kinetics>

3.3 The *TTStroke-21* dataset

The TT-Stroke21 dataset was recorded for fine-grained recognition of sport actions, in the context of the improvement of sport performance for amateurs or professional athletes. Our case study is table tennis, and our goal is the temporal segmentation and classification of strokes performed. The low inter-class variability makes the task more difficult for this content than for more general action databases such as UCF or Kinetics.

Twenty stroke classes and an additional rejection class have been established based on the rules of table tennis. The filmed athletes are students, and their teachers supervise the exercises performed during the recorded sessions. The objective of table tennis stroke recognition is to help the teachers to focus on some of these strokes to help the students in their practice.

Table tennis strokes are most of the time visually similar. Action recognition in this case requires not only a tailored solution, but also a specific expertise to build the ground truth. This is the reason why annotations were carried out by professional athletes. They use a rather rich terminology that allows the fine-grained stroke definition. Moreover, the analysis of the annotations shows that, for the same video and the same stroke, professionals do not always agree. The same holds for defining temporal boundaries of a stroke, which may differ for each annotator. This variability cannot be considered as noise, but shows the ambiguity and complexity of the data and has to be taken into account. We call this new database TTStroke-21, TT standing for Table Tennis and 21 for the number of classes.

3.3.1 TTStroke-21 acquisition

TTStroke-21 is composed of videos of table tennis games with 17 different players. The sequences are recorded indoors without markers using artificial light and light-weight cameras. The recording setting is illustrated in figure 2.a.

The player is filmed in three situations:

- performing repetition of the same stroke. However those repetitions might fail once or several times in the video and the player might do another stroke than the one expected.
- simple exchanges between two players: those exchanges are meant to practise the different techniques.
- in match conditions: the players are meant to mark points. The game speed is much faster and strokes are shorter in time.



Fig. 2 Overview of the TTStroke-21 dataset.

3.3.2 TTStroke-21 annotation

The annotation process was designed as a crowdsourcing method. The annotation sessions were supervised by professional table tennis players and teachers. A user-friendly web platform has been developed by our team for this purpose (figure 2.b), where the annotator spots and labels strokes in videos: starting frame, end frame and the stroke class. The annotator also indicates if the player is right-handed or left-handed. The taxonomy is built upon a shake-hand grip of the racket leading to forehand and backhand stroke according to the side of the racket used.

The taxonomy comprises 20 table tennis stroke classes. All the strokes can, as well, be divided in two super-classes: **Forehand** and **Backhand**. The linguistic analysis of annotations shows that for the same video and the same stroke, professionals do not employ the same degree of details in their annotations. The same problem occurs with temporal analysis: for instance, a service (first stroke when the player releases the ball) might be considered to start i) when the player is in position, ii) when the ball is released or iii) when the racket is moving.

Since a video can be annotated by several annotators, temporal annotations needed to be filtered. An overlap between each annotation of 25% of the annotated stroke duration is allowed. Above this percentage, the annotations are considered to be part of the same stroke and are temporally fused.

Another filter is applied by checking if labels of the same stroke are consistent. If not, this portion of video is not considered in our classification task. This filtering, based on multiple annotations for the same recorded video, can still leave some labeling errors since multiple labeling of the same clip by different annotators was not always easy to meet in practice.

3.3.3 Negative samples extraction

Negative samples are created from videos with more than 10 detected strokes. This was decided after noticing how some videos were poorly annotated and could lead to include actual strokes as negative samples.

The negative samples are video sub-sequences between each detected stroke. We allow the overlap with the previous and the subsequent stroke of 10% of our target time window length: 0.83 seconds, which allows to capture short strokes without considering another one. This represents 100 frames at 120 fps.

3.3.4 Data for evaluation

129 videos at 120 fps have been considered. This content represents 94 minutes of table tennis games, totalling 675 000 video frames and 1 387 annotations. After filtering, 1 074 annotations were retained. The peak statistics of stroke duration are $min = 0.64$ seconds, $max = 2.27$ seconds and the average duration is 1.46 seconds with standard deviation of 0.36. Accordingly, a total of 1 048 strokes were extracted with a min duration of 0.83 seconds, a max duration of 2.31 seconds and an average duration of 1.47 seconds with standard deviation of 0.36. Some annotations were merged making the statistical duration a bit longer. After these steps, 681 negative (non-stroke) samples were extracted. They have a mean duration of 2.34 seconds and standard deviation of 2.66 seconds. This high standard deviation comes from the non game activity of long period between strokes, which can be due to a ball lost or talks of players between games. However, as presented in table 2 representing the distribution over the split of the dataset, not all negative samples are considered to avoid biases in the training and evaluation processes.

4 TSTCNN - A Twin Spatio-Temporal 3D Convolutional Neural Network for action recognition

After having reviewed a bunch of methods for action recognition in video and of reference datasets, we introduce here our solution to the problem of fine-grained action recognition in video with a 3D CNN we call TSTCNN - a twin spatio-temporal CNN.

A two stream architecture

The Twin Spatio-Temporal Convolutional Neural Network (TSTCNN), denoted as illustrated in figure 3, is a two stream 3D Convolutional Network

Table 2 Datasets Taxonomy of TTStroke-21

Table tennis strokes	# Samples				# Frames		
	Train	Val	Test	Sum	Min	Max	Mean*
Def. Backhand Backspin	22	6	3	31	121	233	189 ± 25
Def. Backhand Block	19	5	3	27	100	261	131 ± 37
Def. Backhand Push	6	2	1	9	121	229	155 ± 31
Def. Forehand Backspin	29	8	4	41	129	229	177 ± 25
Def. Forehand Block	8	2	2	12	100	137	115 ± 14
Def. Forehand Push	23	7	3	33	105	177	143 ± 19
Off. Backhand Flip	25	7	3	35	100	265	195 ± 49
Off. Backhand Hit	28	8	4	40	100	173	134 ± 21
Off. Backhand Loop	21	6	3	30	100	229	155 ± 32
Off. Forehand Flip	31	9	5	45	113	269	186 ± 44
Off. Forehand Hit	45	13	6	64	100	233	158 ± 34
Off. Forehand Loop	23	7	3	33	101	277	177 ± 43
Serve Backhand Backspin	56	16	8	80	133	261	188 ± 31
Serve Backhand Loop	43	12	6	61	100	265	186 ± 42
Serve Backhand Sidespin	60	17	9	86	129	269	193 ± 33
Serve Backhand Topspin	57	16	8	81	100	273	175 ± 48
Serve Forehand Backspin	58	17	8	83	125	269	182 ± 35
Serve Forehand Loop	56	16	8	80	100	273	171 ± 51
Serve Forehand Sidespin	57	16	9	82	101	273	192 ± 39
Serve Forehand Topspin	67	19	9	95	100	273	184 ± 52
Non strokes samples	74	21	11	106	100	1255	246 ± 154
Total	808	230	116	1154	100	1255	182 ± 65

* in the form: mean value ± standard deviation

constituted of 2 branches. Each branch follows the same structure: 3 blocks constituted of a 3D convolutional layer using kernels of size $(3 \times 3 \times 3)$, with stride and padding 1 in all directions and “ReLU” as activation function, feeding a 3D Max-Pooling layer using kernels of size $(2 \times 2 \times 2)$ and floor function. From input to output, the convolutional layers use 30, 60 and 80 filters. Each branch ends with a fully connected layer of size 500. The two branches are combined using a bilinear transformation with Softmax function to output a classification score of size 21 corresponding to the number of classes considered in our task.

Dynamic data

The use of dynamic data, such as optical flow, gives extra information to the network and an understanding of the physical world that the network does not have with only the rgb video stream. The OF, which represents the movement through the displacement of the pixels from one image to another, is encoded in the Cartesian coordinate system $V = (v_x, v_y)^T$ with v_x the

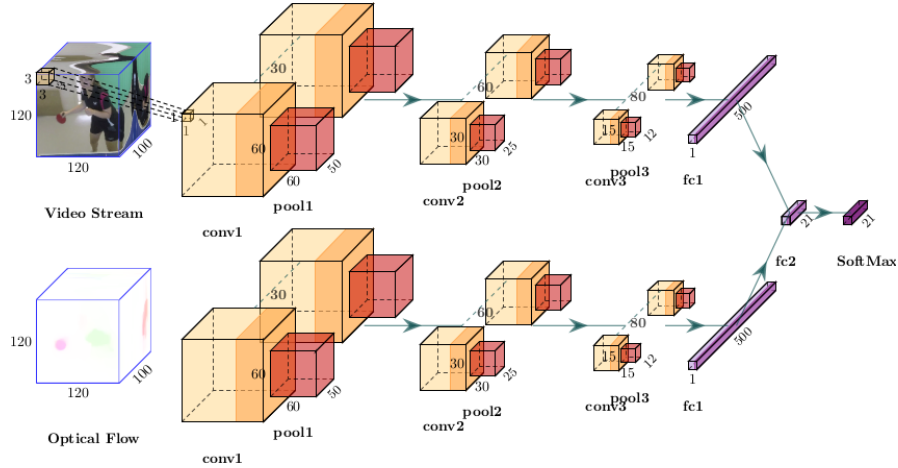


Fig. 3 TSTCNN - Twin Spatio-Temporal Convolutional Neural Network architecture. The video stream and its optical flow are processed in parallel through the network composed of successive 3D convolutional layers and 3D pooling layers and the extracted features a merge at the end through a bilinear transformation giving as output a classification score.

horizontal displacement and v_y the vertical displacement. The optical flow values are then normalized between -1 and 1.

In [MBPM19], several optical flow estimation methods are compared using Mean Squared Error (MSE) of motion compensation, angular error (AE) and end-point error (EPE) quality metrics on two datasets:

- Sintel Benchmark [MIH⁺16], dataset of synthetic videos with available reference optical flows (compared with MSE, AE and EPE),
- and TTStroke-21 [MBPM20], recorded in natural conditions with strong flickering due to synthetic light (compared with MSE only since reference flow is not available).

In the same work, different normalization methods are also tested for classification. According to their results, both of these modalities, optical flow estimator and normalization method, are of primary importance for classification. Indeed, the accuracy of the classification on the test sets of TTStroke-21 varies, according to the best performances for each optical flow estimator, from 41.4% for DeepFlow method [WRHS13] to 74.1% for Beyond Pixel method [Liu09] (BP). Even if BP is sensitive to flickering, it is able to capture fine details, such as the motion of the ball, contrarily to DeepFlow method, explaining such gap between performances of those two estimators. On the other hand, the normalization methods allow for boosting performances from 44% with “Max” normalization method to 74.1% with “Normal” normalization method. The “Max” normalization method strongly reduces the magnitude of most of motion vectors and therefore increases also the inter-

similarity of the strokes; while the “Normal” normalization method increases the magnitude of most vectors and leaves room for inter-dissimilarity.

In the light of these results, we use the Beyond Pixel method [Liu09], based on iterative re-weighted least square solver, to estimate the optical flow and normalize it using the “Normal” normalization method as described in eq. 3.

$$v' = \frac{v}{\mu + 3 \times \sigma}$$

$$v^N(i, j) = \begin{cases} v'(i, j) & \text{if } |v'(i, j)| < 1 \\ \text{SIGN}(v'(i, j)) & \text{otherwise.} \end{cases} \quad (3)$$

where v and v^N represent respectively one component of the optical flow \mathbf{V} and its normalization, μ and σ are mean value and standard deviation of the component.

The optical flow is then filtered by considering only the optical flow of the foreground using the method of Zivkovic and Van der Heijden [ZvdH06].

3D Attention - what could it bring?

Attention mechanisms, in classification problem from rgb images, are used to determine which part of information is useful and/or needed to classify an image. Such attention can be obtained by recording the gaze fixations of individuals when looking at the image and performing the same classification task [OBGR19] as a CNN classifier, to create saliency map on image and then propagate it through the layers of the CNN. In DNNs, internal attention mechanisms have become popular. We distinguish two of them: i) global attention which expresses the contribution of feature channels along convolutional layers into decision making for image classification task [HSA⁺20] and ii) local attention, which focuses on important features in the channels.

When importance of feature channels has to be computed, the processing consists of three steps: i) *squeeze* (synthesis), ii) *excitation* and iii) *feature scaling*. Thus a small network of neurons learns a weighting coefficient for each channel at each layer and outputs the characteristic channels thus weighted to the next layer.

For the local attention, the authors of [WJQ⁺17] use the principles of residual neural networks to propose “residual” learning of the attention masks incorporated in the convolution layers in both forward and backward propagation, which leads to better robustness to noise. By minimizing the objective function by gradient descent, the attention mechanisms are implicitly introduced via the derivative calculation where the weighted characteristics are used. “Teacher-student” networks is another way to introduce attention in the layers of a classification network. In [ZK17] the “Teacher” network is the one that learns attention and guides the “student” network for the image classification task. Such works which focus on image classification inspire the design of spatio-temporal attention to tackle the action classification problem.

In [DYL⁺18], the authors use spatial attention mechanism based on feature pyramids and construct the temporal attention by aggregation of the attention maps obtained spatially, in the temporal domain. Then, based on their work, the authors of [DK19] build 3D convolutional blocks and incorporate them in a 3DResNet network for 3D gesture recognition. However, as in [DYL⁺18], the authors do not use the motion information explicitly. In [MBPM21], we introduce the attention blocks through our TSTCNN into the two video streams: the branch containing the spatial information (RGB) and the branch containing the temporal information (optical flow) as depicted in fig. 6. The optical flow plays the discriminating role for classification and does not act only as localization information.

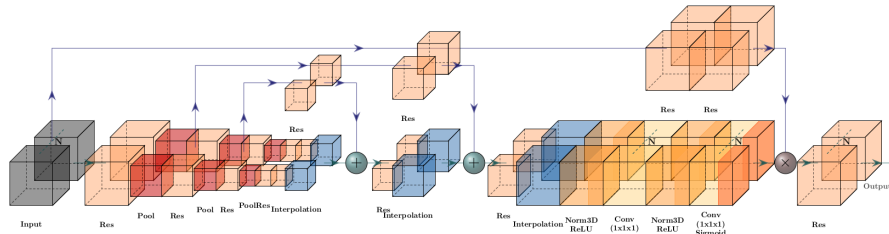


Fig. 4 TSTCNN with 3D attention blocks after each Max Pool layer

As depicted in fig. 4, the attention mechanism uses several 3D residual blocks, illustrated in fig. 5.

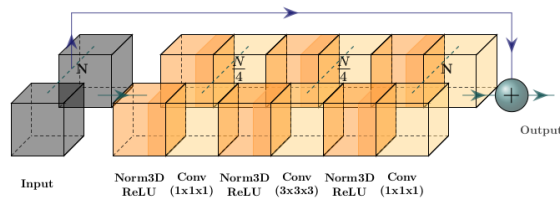


Fig. 5 3D residual block architecture

The implemented residual block is inspired from the work carried out in 2D in [HZRS16] and has been extended and adapted in 3D with a 4D data block of size $(N \times W \times H \times T)$ representing respectively the number of channels, the two spatial dimensions and the temporal dimension. Input data are then processed by 3 successive layers $f_{conv_i}, i = 1, \dots, 3$, see eq. 10, with respectively $\frac{N}{4}, \frac{N}{4}$ and N filters of size $(1 \times 1 \times 1)$, $(3 \times 3 \times 3)$ and $(1 \times 1 \times 1)$. Their output is then summed with the input data to build the output of the 3D residual block.

$$res(x) = f_{conv_3}(f_{conv_2}(f_{conv_1}(x))) + x \quad (4)$$

The 3D batch normalization is performed channel by channel over the batch of data. If we have $x = (x_1, x_2, \dots, x_{N_{channels}})$, then the normalization is $F_n(x) = (f_n(x_1), f_n(x_2), \dots, f_n(x_{N_{channels}}))$ with:

$$f_n(x_i) = \frac{x_i - \mu_i}{\sqrt{\sigma_i^2 \oplus \epsilon}} * \gamma_i \oplus \beta_i \quad (5)$$

with $i = 1, \dots, N_{channels}$, μ_i and σ_i the mean and standard deviation vectors of x_i computed over the batch, \oplus is addition of a scalar to each vector coordinate, γ_i and β_i are learnable parameters per channel and the division by $\sqrt{\sigma_i^2 \oplus \epsilon}$ is element-wise. Here, $N_{channels} = N$ or $\frac{N}{4}$, depending on the normalization position in the residual block.

The 3D attention blocks illustrated in Figure 4 are inspired from the work carried out in 2D in [WJQ⁺17].

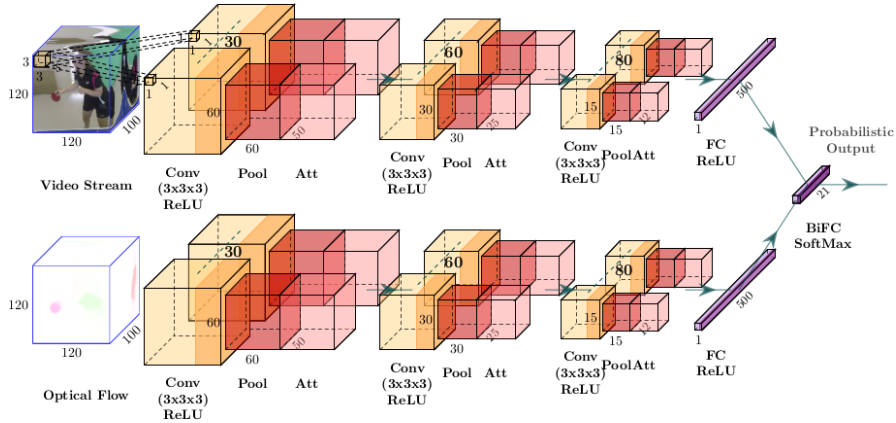


Fig. 6 3D attention block architecture

A 3D attention block takes as input a 4D data block of size $(N \times W \times H \times T)$. In this block, all convolutions are performed with the same number of filters, N , to maintain the dimension of the processed data. The input data are processed by the first 3D residual block, denoted as “ res ”. The network then splits in two branches: the trunk branch consisting of 2 successive 3D residual blocks, described by equation 6; and the soft floating mask branch, see lowest position in Figure 4), described through the equations 7, 8, 10. Its role is to accentuate the features generated by the trunk branch. Those two branches are merged as described in equation 11.

$$branch_{trunk}(\cdot) = res(res(\cdot)) \quad (6)$$

The soft mask branch is constituted of several 3D residual blocks followed by Max Pooling layers, denoted as “*MaxP*”. It increases the reception field of convolutions using a bottom-up architecture, denoted as $f_{bu}(\cdot) = res(MaxP(\cdot))$. The lowest resolution is obtained after 3 Max Pooling steps.

$$\begin{aligned} x_1 &= f_{bu}(res(Input)) \\ x_2 &= f_{bu}(x_1) \\ x_3 &= f_{bu}(x_2) \end{aligned} \quad (7)$$

The information is then extended by a symmetrical top-down architecture, $f_{td}(\cdot) = Inter(res(\cdot))$, to project the input features of each resolution level. “*Inter*” denotes the trilinear interpolations used for up-sampling. Two skipped connections are used for collecting information at different scales.

$$\begin{aligned} y_1 &= f_{td}(x_3) + res(x_2) \\ y_2 &= f_{td}(y_1) + res(x_1) \\ y_3 &= f_{td}(y_2) \end{aligned} \quad (8)$$

The soft mask branch is then composed of 2 successive layers. Each includes a 3D batch normalization, denoted as $F_n(\cdot)$ as described by equation 5, followed by a ReLU activation function and a convolution layer with kernel sizes $(1 \times 1 \times 1)$. This is expressed by equation 9:

$$f_{conv}(\cdot) = conv(ReLU(F_n(\cdot))) \quad (9)$$

It ends with a sigmoid function, denoted as “*Sig*”, to scale values between 0 and 1. These two layers are depicted on the right of the lowest branch in figure 4 and are expressed by equation 10.

$$branch_{f_{mask}}(Input) = Sig(f_{conv}(f_{conv}(y_3))) \quad (10)$$

The output of the trunk branch is then multiplied term by term by $(1 \oplus branch_{f_{mask}}(Input))$ where $branch_{f_{mask}}(Input)$ is the output of the mask branch. The result is then processed by the last 3D residual block $res(\cdot)$ which ends the attention block, see equation 11.

$$y = res(branch_{trunk}(Input) \odot (1 \oplus branch_{f_{mask}}(Input))) \quad (11)$$

Here the \odot is an element-wise multiplication and \oplus is an addition of a scalar to each vector component as defined above.

4.1 Results

In Table 3, we compare the models in terms of accuracy for the pure classification task. In order to have an overall view, comparison is done with the

models using three attention blocks (one after each max pooling layer) with the models and without them.

Table 3 Comparison of the classification performances for models using attention mechanism after convergence in terms of accuracy.

Models	Epochs	Accuracies in %		
		Train	Val	Test
RGB-I3D	778	98.3	72.6	69.8
RGB-STCNN	1665	96.7	88.7	89.8
RGB-STCNN with Attention	524	96.5	88.3	93.2
Flow-I3D	1112	98.8	74.8	73.3
Flow-STCNN	1449	97.5	79.6	75.9
Flow-STCNN with attention	732	96.4	83.5	90.7
Two-Stream I3D	-	99.2	76.2	75.9
LF-STCNN	-	97	88.7	89.8
LF-STCNN with attention	-	97	88.7	94.9
T-STCNN	1784	95.8	87.8	93.2
T-STCNN with attention	591	97.3	87.8	95.8

The presented models using spatio-temporal convolutions shallower than the I3D models prove to be more efficient to classify fine-grained actions on a challenging dataset with limited amount of samples. The fine-grained aspect of the task seems to be dealt better if the model is not too deep. Indeed, the deepness of the I3D models is efficient on task were a great number of objects and scenes need to be recognized in order to classify coarse-grained actions, but it seems to be less effective for the fine-grained task.

Furthermore, the amount of samples and the fact that the models are not pre-trained, makes the task even harder. The overfitting problem is more noticeable on the very deep models I3D. However, the STCNN models using attention mechanisms do not suffer from this: the performances are improved and the convergence is faster. Best performances are observed for the Twin model using attention mechanism. The intermediate fusion allows a better combination of each branches in order to perform classification. A late fusion approach seems efficient but is 1% behind.

5 Conclusion and Perspectives

In this chapter, we have shown that the features developed to perform for action classification in videos have evolved from handcrafted to deep learnt with the increasing complexity of datasets and their increasing number of video samples. Furthermore, the increasing capacity of CPUs and GPUs to

process large amount of data have allowed deep learning methods to move from 2D to 3D convolutions. 3D convolutions for action classification from videos have proven to be efficient tools to capture efficiently spatio-temporal features in order to perform classification.

The incorporation of attention mechanisms, through attention blocks, helps in the process by highlighting the discriminant features, boosting the convergence speed and classification performances. Such behaviour was observed on the fine-grained dataset TTStroke-21. It has also been noted that the deepness of the implemented models needs to be adapted to the classification task and the dataset in order to avoid overfitting.

References

- [BGS⁺17] Mateusz Budnik, Efrain-Leonardo Gutierrez-Gomez, Bahjat Safadi, Denis Pellerin, and Georges Quénot. Learned features versus engineered features for multimedia indexing. *Multim. Tools Appl.*, 76(9):11941–11958, 2017.
- [BPM14] Cyrille Beaudry, Renaud Péteri, and Laurent Mascarilla. Action recognition in videos using frequency analysis of critical point trajectories. In *2014 IEEE International Conference on Image Processing, ICIP 2014, Paris, France, October 27-30, 2014*, pages 1445–1449, 2014.
- [CDF⁺04] Gabriella Csurka, Christopher R. Dance, Lixin Fan, Jutta Willamowski, and Cédric Bray. Visual categorization with bags of keypoints. In *In Workshop on Statistical Learning in Computer Vision, ECCV*, pages 1–22, 2004.
- [CNHZ19] João Carreira, Eric Noland, Chloe Hillier, and Andrew Zisserman. A short note on the kinetics-700 human action dataset. *CoRR*, abs/1907.06987, 2019.
- [CPM19] Jordan Calandre, Renaud Péteri, and Laurent Mascarilla. Optical flow singularities for sports video annotation: Detection of strokes in table tennis. In *Working Notes Proceedings of the MediaEval 2019 Workshop, Sophia Antipolis, France, 27-30 October 2019*, 2019.
- [CRHV09] Rizwan Chaudhry, Avinash Ravichandran, Gregory D. Hager, and René Vidal. Histograms of oriented optical flow and binet-cauchy kernels on nonlinear dynamical systems for the recognition of human actions. In *2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2009), 20-25 June 2009, Miami, Florida, USA*, pages 1932–1939, 2009.
- [CZ17] João Carreira and Andrew Zisserman. Quo vadis, action recognition? A new model and the kinetics dataset. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 4724–4733, 2017.
- [dCBM⁺11] Teofilo de Campos, Mark Barnard, Krystian Mikolajczyk, Josef Kittler, Fei Yan, William J. Christmas, and David Windridge. An evaluation of bags-of-words and spatio-temporal shapes for action recognition. In *IEEE Workshop on Applications of Computer Vision (WACV 2011), 5-7 January 2011, Kona, HI, USA*, pages 344–351, 2011.
- [DHR⁺17] Jeff Donahue, Lisa Anne Hendricks, Marcus Rohrbach, Subhashini Venugopalan, Sergio Guadarrama, Kate Saenko, and Trevor Darrell. Long-term recurrent convolutional networks for visual recognition and description. *IEEE Trans. Pattern Anal. Mach. Intell.*, 39(4):677–691, 2017.
- [DK19] Naina Dhingra and Andreas M. Kunz. Res3atn - deep 3d residual attention network for hand gesture recognition in videos. In *2019 International Confer-*

- ence on *3D Vision, 3DV 2019, Québec City, QC, Canada, September 16-19, 2019*, pages 491–501, 2019.
- [DYL⁺18] Yang Du, Chunfeng Yuan, Bing Li, Lili Zhao, Yangxi Li, and Weiming Hu. Interaction-aware spatio-temporal pyramid attention networks for action classification. In *ECCV (16)*, volume 11220 of *Lecture Notes in Computer Science*, pages 388–404. Springer, 2018.
- [EBMM03] Alexei A. Efros, Alexander C. Berg, Greg Mori, and Jitendra Malik. Recognizing action at a distance. In *9th IEEE International Conference on Computer Vision (ICCV 2003), 14-17 October 2003, Nice, France*, pages 726–733, 2003.
- [FPW17] Christoph Feichtenhofer, Axel Pinz, and Richard P. Wildes. Spatiotemporal multiplier networks for video action recognition. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 7445–7454, 2017.
- [FVD⁺19] Mehrnaz Fani, Kanav Vats, Christopher Dulhanty, David A. Clausi, and John S. Zelek. Pose-projected action recognition hourglass network (PARHN) in soccer. In *16th Conference on Computer and Robot Vision, CRV 2019, Kingston, ON, Canada, May 29-31, 2019*, pages 201–208, 2019.
- [GBS⁺07] Lena Gorelick, Moshe Blank, Eli Shechtman, Michal Irani, and Ronen Basri. Actions as space-time shapes. *IEEE Trans. Pattern Anal. Mach. Intell.*, 29(12):2247–2253, 2007.
- [GHS13] Adrien Gaidon, Zaïd Harchaoui, and Cordelia Schmid. Temporal localization of actions with actoms. *IEEE Trans. Pattern Anal. Mach. Intell.*, 35(11):2782–2795, 2013.
- [GSR⁺18] Chunhui Gu, Chen Sun, David A. Ross, Carl Vondrick, Caroline Pantofaru, Yeqing Li, Sudheendra Vijayanarasimhan, George Toderici, Susanna Ricco, Rahul Sukthankar, Cordelia Schmid, and Jitendra Malik. AVA: A video dataset of spatio-temporally localized atomic visual actions. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 6047–6056, 2018.
- [GTM19] Deepti Ghadiyaram, Du Tran, and Dhruv Mahajan. Large-scale weakly-supervised pre-training for video action recognition. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 12046–12055, 2019.
- [HCS17] Rui Hou, Chen Chen, and Mubarak Shah. Tube convolutional neural network (T-CNN) for action detection in videos. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pages 5823–5832, 2017.
- [HOSS13] Michael E. Houle, Vincent Oria, Shin’ichi Satoh, and Jichao Sun. Annotation propagation in image databases using similarity graphs. *ACM Trans. Multim. Comput. Commun. Appl.*, 10(1):7:1–7:21, 2013.
- [HSA⁺20] Jie Hu, Li Shen, Samuel Albanie, Gang Sun, and Enhua Wu. Squeeze-and-excitation networks. *IEEE Trans. Pattern Anal. Mach. Intell.*, 42(8):2011–2023, 2020.
- [HZRS16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 770–778, 2016.
- [IBPQ14] Bogdan Ionescu, Jenny Benois-Pineau, Tomas Piatrik, and Georges Quénot, editors. *Fusion in Computer Vision - Understanding Complex Visual Content*. Advances in Computer Vision and Pattern Recognition. Springer, 2014.
- [Joh73] G. Johansson. Visual perception of biological motion and a model for its analysis. *Perception and Psychophysics*, 14:pp. 201–211, 1973.
- [JvGJ⁺14] Mihir Jain, Jan C. van Gemert, Hervé Jégou, Patrick Bouthemy, and Cees G. M. Snoek. Action localization with tubelets from motion. In *2014*

- IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2014, Columbus, OH, USA, June 23-28, 2014*, pages 740–747, 2014.
- [JXY13] Shuiwang Ji, Wei Xu, Ming Yang, and Kai Yu. 3d convolutional neural networks for human action recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 35(1):221–231, 2013.
- [KCS⁺17] Will Kay, João Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, Mustafa Suleyman, and Andrew Zisserman. The kinetics human action video dataset. *CoRR*, abs/1705.06950, 2017.
- [KGOG06] Ewa Kijak, Guillaume Gravier, Lionel Oisel, and Patrick Gros. Audiovisual integration for tennis broadcast structuring. *Multim. Tools Appl.*, 30(3):289–311, 2006.
- [KKA20] M. Esat Kalfaoglu, Sinan Kalkan, and A. Aydin Alatan. Late temporal modeling in 3d CNN architectures with BERT for action recognition. *CoRR*, abs/2008.01232, 2020.
- [Lap13] Ivan Laptev. *Modeling and visual recognition of human actions and interactions*. Habilitation à diriger des recherches, Ecole Normale Supérieure de Paris - ENS Paris, July 2013.
- [Liu09] Ce Liu. *Beyond Pixels: Exploring New Representations and Applications for Motion Analysis*. PhD thesis, Massachusetts Institute of Technology, 5 2009.
- [LJ20] Jinkue Lee and Hoeryong Jung. Tuhad: Taekwondo unit technique human action dataset with key frame-based cnn action recognition. *Sensors*, 20(17):4871, 2020.
- [LP07] Ivan Laptev and Patrick Pérez. Retrieving actions in movies. In *IEEE 11th International Conference on Computer Vision, ICCV 2007, Rio de Janeiro, Brazil, October 14-20, 2007*, pages 1–8, 2007.
- [LTR⁺20] Ang Li, Meghana Thotakuri, David A. Ross, João Carreira, Alexander Vostrikov, and Andrew Zisserman. The ava-kinetics localized human actions video dataset. *CoRR*, abs/2005.00214, 2020.
- [LWLW16] Zhihao Li, Wenmin Wang, Nannan Li, and Jinzhuo Wang. Tube convnets: Better exploiting motion for action recognition. In *2016 IEEE International Conference on Image Processing, ICIP 2016, Phoenix, AZ, USA, September 25-28, 2016*, pages 3056–3060, 2016.
- [LWS⁺19] Ruichen Liu, Zhelong Wang, Xin Shi, Hongyu Zhao, Sen Qiu, Jie Li, and Ning Yang. Table tennis stroke recognition based on body sensor network. In *Internet and Distributed Computing Systems - 12th International Conference, IDCSS 2019, Naples, Italy, October 10-12, 2019, Proceedings*, pages 1–10, 2019.
- [Mac67] J. Macqueen. Some methods for classification and analysis of multivariate observations. In *In 5-th Berkeley Symposium on Mathematical Statistics and Probability*, pages 281–297, 1967.
- [MBPM19] Pierre-Etienne Martin, Jenny Benois-Pineau, Renaud Péteri, and Julien Morlier. Optimal choice of motion estimation methods for fine-grained action classification with 3d convolutional networks. In *2019 IEEE International Conference on Image Processing, ICIP 2019, Taipei, Taiwan, September 22-25, 2019*, pages 554–558, 2019.
- [MBPM20] Pierre-Etienne Martin, Jenny Benois-Pineau, Renaud Péteri, and Julien Morlier. Fine grained sport action recognition with twin spatio-temporal convolutional neural networks. *Multim. Tools Appl.*, 79(27-28):20429–20447, 2020.
- [MBPM21] Pierre-Etienne Martin, Jenny Benois-Pineau, Renaud Péteri, and Julien Morlier. 3d attention mechanisms in twin spatio-temporal convolutional neural networks. application to action classification in videos of table tennis games. In *25th International Conference on Pattern Recognition (ICPR2020) - MiCo Milano Congress Center, Italy, 10-15 January 2021*, 2021.
- [MIH⁺16] Nikolaus Mayer, Eddy Ilg, Philip Häusser, Philipp Fischer, Daniel Cremers, Alexey Dosovitskiy, and Thomas Brox. A large dataset to train convolutional

- networks for disparity, optical flow, and scene flow estimation. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 4040–4048, 2016.
- [MLS09] Marcin Marszałek, Ivan Laptev, and Cordelia Schmid. Actions in context. In *2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2009), 20-25 June 2009, Miami, Florida, USA*, pages 2929–2936, 2009.
- [NHV⁺15] Joe Yue-Hei Ng, Matthew J. Hausknecht, Sudheendra Vijayanarasimhan, Oriol Vinyals, Rajat Monga, and George Toderici. Beyond short snippets: Deep networks for video classification. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 4694–4702, 2015.
- [OBGR19] Abraham Montoya Obeso, Jenny Benois-Pineau, Mireya Saraf García-Vázquez, and Alejandro Alvaro Ramírez-Acosta. Forward-backward visual saliency propagation in deep nns vs internal attentional mechanisms. In *Ninth International Conference on Image Processing Theory, Tools and Applications, IPTA 2019, Istanbul, Turkey, November 6-9, 2019*, pages 1–6, 2019.
- [RHGS17] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. Faster R-CNN: towards real-time object detection with region proposal networks. *IEEE Trans. Pattern Anal. Mach. Intell.*, 39(6):1137–1149, 2017.
- [RMB⁺10] A Roksztin, Z Márkus, G Braunitzer, A Berényi, G Benedek, and A Nagy. Visual pathways serving motion detection in the mammalian brain. *Sensors*, 10(4):3218–3242, 2010.
- [SFBC16] Andrei Stoian, Marin Ferecatu, Jenny Benois-Pineau, and Michel Crucianu. Fast action localization in large-scale video archives. *IEEE Trans. Circuits Syst. Video Techn.*, 26(10):1917–1930, 2016.
- [SG08] Konrad Schindler and Luc Van Gool. Action snippets: How many frames does human action recognition require? In *2008 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2008), 24-26 June 2008, Anchorage, Alaska, USA*, 2008.
- [SLC04] Christian Schüldt, Ivan Laptev, and Barbara Caputo. Recognizing human actions: A local SVM approach. In *17th International Conference on Pattern Recognition, ICPR 2004, Cambridge, UK, August 23-26, 2004*, pages 32–36, 2004.
- [SLJ⁺15] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott E. Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 1–9, 2015.
- [SZ14] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 568–576, 2014.
- [SZDL20] Dian Shao, Yue Zhao, Bo Dai, and Dahua Lin. Finegym: A hierarchical video dataset for fine-grained action understanding. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 2613–2622, 2020.
- [SZS12] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. UCF101: A dataset of 101 human actions classes from videos in the wild. *CoRR*, abs/1212.0402, 2012.
- [TBF⁺15] Du Tran, Lubomir D. Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pages 4489–4497, 2015.

- [TWT⁺18] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 6450–6459, 2018.
- [UAM⁺18] Amin Ullah, Jamil Ahmad, Khan Muhammad, Muhammad Sajjad, and Sung Wook Baik. Action recognition in video sequences using deep bi-directional LSTM with CNN features. *IEEE Access*, 6:1155–1166, 2018.
- [VLS18] Gül Varol, Ivan Laptev, and Cordelia Schmid. Long-term temporal convolutions for action recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 40(6):1510–1517, 2018.
- [VSP⁺17] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 5998–6008, 2017.
- [WGW⁺18] Xuanhan Wang, Lianli Gao, Peng Wang, Xiaoshuai Sun, and Xianglong Liu. Two-stream 3-d convnet fusion for action recognition in videos with arbitrary size and length. *IEEE Trans. Multimedia*, 20(3):634–644, 2018.
- [WJQ⁺17] Fei Wang, Mengqing Jiang, Chen Qian, Shuo Yang, Cheng Li, Honggang Zhang, Xiaogang Wang, and Xiaoou Tang. Residual attention network for image classification. In *CVPR*, pages 6450–6458. IEEE Computer Society, 2017.
- [WKS11] Heng Wang, Alexander Kläser, Cordelia Schmid, and Cheng-Lin Liu. Action recognition by dense trajectories. In *The 24th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2011, Colorado Springs, CO, USA, 20-25 June 2011*, pages 3169–3176, 2011.
- [WRHS13] Philippe Weinzaepfel, Jérôme Revaud, Zaïd Harchaoui, and Cordelia Schmid. Deepflow: Large displacement optical flow with deep matching. In *IEEE International Conference on Computer Vision, ICCV 2013, Sydney, Australia, December 1-8, 2013*, pages 1385–1392, 2013.
- [WS13] Heng Wang and Cordelia Schmid. Action recognition with improved trajectories. In *IEEE International Conference on Computer Vision, ICCV 2013, Sydney, Australia, December 1-8, 2013*, pages 3551–3558, 2013.
- [WZD⁺20] Jiachen Wang, Kejian Zhao, Dazhen Deng, Anqi Cao, Xiao Xie, Zheng Zhou, Hui Zhang, and Yingcai Wu. Tac-simur: Tactic-based simulative visual analytics of table tennis. *IEEE Trans. Vis. Comput. Graph.*, 26(1):407–417, 2020.
- [XWX⁺20] Kun Xia, Hanyu Wang, Menghan Xu, Zheng Li, Sheng He, and Yusong Tang. Racquet sports recognition using a hybrid clustering model learned from integrated wearable sensor. *Sensors*, 20(6):1638, 2020.
- [ZK17] Sergey Zagoruyko and Nikos Komodakis. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*, 2017.
- [ZvdH06] Zoran Zivkovic and Ferdinand van der Heijden. Efficient adaptive density estimation per image pixel for the task of background subtraction. *Pattern Recognit. Lett.*, 27(7):773–780, 2006.