



HAL
open science

3D attention mechanism for fine-grained classification of table tennis strokes using a Twin Spatio-Temporal Convolutional Neural Networks

Pierre-Etienne Martin, Jenny Benois-Pineau, Renaud Péteri, Julien Morlier

► To cite this version:

Pierre-Etienne Martin, Jenny Benois-Pineau, Renaud Péteri, Julien Morlier. 3D attention mechanism for fine-grained classification of table tennis strokes using a Twin Spatio-Temporal Convolutional Neural Networks. 25th International Conference on Pattern Recognition (ICPR2020), Jan 2021, Milano, Italy. hal-02977646

HAL Id: hal-02977646

<https://hal.science/hal-02977646>

Submitted on 19 Nov 2020

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 attention mechanism for fine-grained classification of table tennis strokes using a Twin Spatio-Temporal Convolutional Neural Networks

Pierre-Etienne Martin^{*}, Jenny Benois-Pineau^{*}, Renaud Péteri[†] and Julien Morlier[‡]

^{*}Univ. Bordeaux, CNRS, Bordeaux INP, LaBRI, UMR 5800, F-33400, Talence, France

Emails: pierre-etienne.martin@u-bordeaux.fr, jenny.benois-pineau@u-bordeaux.fr

[†]MIA, University of La Rochelle, La Rochelle, France

Email: renaud.peteri@univ-lr.fr

[‡]IMS, University of Bordeaux, Talence, France

Abstract—The paper addresses the problem of recognition of actions in video with low inter-class variability such as Table Tennis strokes. Two stream, “twin” convolutional neural networks are used with 3D convolutions both on RGB data and optical flow. Actions are recognized by classification of temporal windows. We introduce 3D attention modules and examine their impact on classification efficiency. In the context of the study of sportsmen performances, a corpus of the particular actions of table tennis strokes is considered. The use of attention blocks in the network speeds up the training step and improves the classification scores up to 5% with our twin model. We visualize the impact on the obtained features and notice correlation between attention and player movements and position. Score comparison of state-of-the-art action classification method and proposed approach with attentional blocks is performed on the corpus. Proposed model with attention blocks outperforms previous model without them and our baseline.

I. INTRODUCTION

Action recognition in videos is one of the key problems in computer vision. Despite intensive research, recognition and discrimination of visually very similar actions remains a challenge [1]. Several methods have been developed in the literature, the most recent ones using deep learning approaches [2]. Numerous works use models based on temporal networks such as RNN and LSTM [3]. However, these networks are difficult to train and lack of stability [4]. 3D convolutional neural networks are a good alternative for capturing long-term dependencies [5]. They involve 3D convolutions: convolutions in space and time. These deep network architectures of multilayer perceptron type include - in the first layers - an extraction of features, and classifiers in the last layers. This type of approach translates in a powerful way through the prism of deep learning, what we knew beforehand: to extract features from time windows and use them for classification [6]. Recent methods also improve performances of Inflated 3D ConvNet [5] by either capturing simultaneously slow and fast features using different video frame rates [7] or by adding

This work was supported by the New Aquitaine Region and Bordeaux Idex Initiative grant within the framework of the CRISP project - ComputeR vIsion for Sport Performance.

non local operations in the network [8]. In [9], the authors also obtain 3D CNN performances while keeping the 2D CNN complexity through Temporal Shift Modules.

Recognition of similar actions is required quite frequently, and belongs to the fine-grained classification problem. In sport for instance, such as table tennis or gymnastics [10], exercises are filmed in the same environment and movements can be quite similar. Hence the recognition problem becomes harder: the classifier cannot be helped by background information where an action is performed. The classifier has to focus on meaningful regions and changes in the video to be efficient. This is a subject of a recent trend in Deep Learning, that is the introduction of “attention mechanisms”. The latter are designed to reinforce the contribution of meaningful features and channels into the decision and thus to increase the target accuracy. Recently we proposed a comparative study of these attention mechanisms inherent to convolutional networks, as described in [11]. Selection of the most relevant characteristics in different layers is very similar to the human attention mechanisms measured in psycho-visual experiments. While these attention mechanisms in 2D networks have been intensively studied [12], this question remains to be further explored for a spatio-temporal content analysis using 3D convolutional networks.

The target application of our research is fine grained recognition of sport actions, in the context of the improvement of sport performances for amateur 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-101 [13] or DeepMind Kinetics [14]. Twenty stroke classes and an additional rejection class have been established based on the rules of table tennis. This taxonomy has been discussed and designed with table tennis professionals. We are working on videos recorded at the Faculty of Sports of the University of Bordeaux (STAPS). The filmed athletes are students, and their teachers supervise the exercises performed during the recording sessions. These recordings are done without mark-

ers, which allows the players to play in natural conditions. 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.

We can mention that action recognition in Table Tennis videos is recently getting interest in the research community. In [15], the authors try to visualize and characterize tactics in table tennis competitions using a Markov chain model for comparing the profile of different players. Other works only focus on the ball tracking and trajectory estimation [16]. In [17], the authors propose an advanced real-time solution for scene segmentation, ball trajectory estimation and event detection but are not considering stroke classification.

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 that has to be taken into account. This new database, called TTStroke-21 has been introduced in [18], [19] as well as the twin network architecture used - Twin Spatio-Temporal CNN (TSTCNN).

Attention mechanisms for action recognition have been recently introduced in LSTM [20] in a recognition approach based on the analysis of joints of a human skeleton. In 3D CNNs both global channel attention and spatial attention maps for different feature layers have been proposed [21]. We also follow this trend and design attentional blocks for our TSTCNN model. In this paper, we propose spatio-temporal attention mechanisms in 3D convolution networks for recognition of challenging similar actions: Table Tennis strokes.

The rest of the paper is organized as follows: in section II, works using attention mechanisms are presented. The section III presents the proposed method with attentional mechanisms and in details the attention block. The results are drawn in section IV through feature analysis and classification performances. The conclusion and perspectives are given in the section V.

II. STATE OF THE ART ON ATTENTION MECHANISMS

In this section, we present a brief state of the art on attention mechanisms introduced in convolutional networks for the classification of images and videos. One can distinguish two classes: 2D attention models, which concern images, and 3D models (2D +T) concerning videos. Although such a separation may seem artificial as the same principles govern the design of the models in both cases, we prefer to treat the spatio-temporal content separately.

A. 2D Attention Models

One of the pioneering works introducing the use of an attention model in neural networks for image classification is presented in [22]. The authors are interested in the contribution of feature channels along convolutional layers into decision making. The attention model here is “global”: a channel weighting mechanism is introduced by “attention blocks”. The processing consists of three steps: i) synthesis (*squeeze*), ii) excitation (*excitation*) and iii) feature scaling (*scale*). Thus, for each channel, a block is a small network of neurons that learns a weighting coefficient. The next layers of the network ingest the characteristic channels thus weighted. This global weighting has been used as a basis for the authors of [23] who propose “double attention” blocks, i.e. ensuring a global and spatial weighting of the characteristics in convolution network layers. The authors of [12], [24] use the principles of residual neural networks to propose “residual” learning of the attention masks incorporated in the convolution layers. Their experiments on CIFAR data bases show that on CIFAR 10 the residual attention network with depth of 452 has the best error rate compared to all the basic residual networks (3.90%). The authors propose the incorporation of attention mechanisms in both forward (*forward*) and backward (*backward*) runs. This is also the approach we had in [11], but by selecting important characteristics and not by weighting features and channels. Note that when minimizing the objective function by gradient descent, the attention mechanisms are implicitly introduced via the derivative calculation where the weighted characteristics are used. The authors of [12] report that this use in back propagation makes the training data robust to noise. This is also our approach in this work.

Other works such as [25] propose “Teacher-student” networks where the “Teacher” network is the one that learns attention and guides the student network for the image classification task. In our approach we also use a kind of attention transfer as in our architecture the attention branch and trunk branch will join together for selection of important features. Also, in [26], attention mechanism is coupled with LSTM to learn the correlation between different data modalities such as text and image and therefore leads to better embedding.

B. 3D Attention Models

We focus here on the contribution of “3D” (2D+T) spatio-temporal attention models in deep networks for the action recognition problem.

Attention mechanism is used in [20] on joint skeleton and coupled with LSTM for 3D action recognition task. They reported better accuracy with attention mechanism than without. They also propose a recurrent attention mechanism on their model which strengthens the attention effect but might not lead to better performances if iterated too many times. In [27], the author consider attention on the channels of aggregated temporal features extracted from videos on appearance and motion streams. Similarly, the authors in [28] construct a spatial attention model for each image by introducing feature pyramids. The temporal extension is obtained by a simple

aggregation of the attention maps estimated for each of the K images of the pyramid extended to the spatio-temporal domain. Motion information is not taken into account. We differ from this approach by introducing attention blocks in our twin network at the level of the two branches: RGB and the optical flow.

In [29], motion information, via the optical flow, acts as the attention map for locating actions in the video. The authors introduce the “motion condition” layer to train the network on RGB appearance components conditional to this optical flow based map. The motion weighting layer allows to modify the spatial characteristics in the convolution layers. We find here the philosophy of using motion as an indicator of areas of interest [30]. Once more, the difference of our approach consists in introducing attention blocks in the two branches (RGB and optical flow) of the twin network.

In [31], attention clusters over the temporal dimension are used on image features extracted using Inception-ResNet-v2 [32] for RGB and Flow modalities. The Inception-ResNet-v2 models are pretrained on ImageNet [33] and are fine-tuned for the optical flow model. A third branch processes the audio signal using VGG-16 [34] on extracted spectrogram samples and is processed similarly to an image [35].

3D attention blocks have also recently been introduced in ResNet 3D type networks for the recognition of 3D hand gestures from videos [24] or from action recognition dataset [21] such as HMDB-51 [36], UCF-101 [13] and Kinetics [14]. The authors of [24] build on the work of [12], and propose a convolution network using the RGB image for feature extraction, and another coupled network to determine a soft attention mask with the derivable sigmoid function. The values of the extracted mask are then combined with the characteristics extracted from the RGB array. As in [28], the authors do not use the motion information explicitly.

Hence, our approach differs from current methods in the literature in the following:

- we introduce the 3D attention blocks into the two video streams: the branch containing the spatial information (RGB) and the branch containing the temporal information (optical flow).
- movement (optical flow) plays the discriminating role in our fine-grained classification context, our problem being to recognize actions and not to locate them.

In the following section, we first introduce the basis of our approach - the 3D (2D+T) convolutional twin arrays that we previously proposed in [18], we then detail the spatio-temporal attention blocks introduced in our network.

III. 3D ATTENTION MECHANISM IN TWIN SPACE-TIME NETWORKS

We first introduce our Twin Spatio-Temporal Convolutional Neural Network - TSTCNN used for classification and then detail the 3D attention and residual blocks developed and tested.

A. Twin Spatio-Temporal Convolutional Neural Network - TSTCNN

In order to perform action classification in videos, we use a two streams convolutional neural network (*twin*) with attention mechanism. Its architecture without attention blocks and early results on TTStroke21 are described in [19] and the analysis of the impact on classification of optical flow normalization is done in [37]. Its architecture with attention mechanism is presented in figure 1. The difference from other Two stream networks [38], [39], [40] lies in: i) the symmetries of our network, ii) the input 4D data type (horizontal, vertical, temporal and channel) and iii) the final fusion step with a bilinear layer at the end of our two branches.

Our twin network, the so-called TSTCNN - (Twin Spatio-Temporal Convolutional Neural Network) consists of two individual branches: one branch takes as input the values of the RGB images of the sequence, the other branch uses the optical flow estimated by the method of [41]. It thus allows the incorporation of both spatial and temporal features [18]. The played stroke is predicted from the RGB images of the sequence and the estimated motion vectors $\mathbf{v} = (v_x, v_y)^T$.

Each branch consists of three convolutional layers comprising successively 30, 60 and 80 3D filters, followed by a fully connected layer of size 500. The 3D convolutional layers use space-time filters of size $3 \times 3 \times 3$. The two branches are merged through a final bilinear fully connected layer of size 21, followed by a Softmax function to obtain an output class membership probability. A detailed implementation and description of our network is available on GitHub¹ to facilitate reproducibility or the use of our method for other applications.

1) *Learning Phase*: learning of our TSTCNN network is done by stochastic gradient descent (SGD) with Nesterov momentum [42]. In order to avoid overfitting, data augmentation is performed in the spatial domain using rotations, homotheties and scale transformations. Data augmentation is also performed in the time domain in order to add variability around the temporal boundaries of the played stroke. We refer the reader to [18] for more details.

2) *Performance without attention*: the classification on the test data of the database TTStroke-21 gives an accuracy of 91.4% in [18]. For comparison, the two-branch I3D model [5], used as a baseline, gives an accuracy of only 43.1% for the same dataset using a temporal window of 64 frames. Even if the results are satisfactory, it should be noted that the trained network gives comparable results in terms of accuracy: the gain compared to a network using image information only is 2.7%. There are therefore still room for improvement, notably by making better use of motion information during the training phase, as in [4], [37], or by introducing information from 3D attention models.

Detailed results, including experimental settings and further analysis, are presented in [18].

3) *TSTCNN with residual and attention mechanism*: to test the efficiency of residual and attention blocks, we first

¹<https://github.com/P-eMartin/crisp>

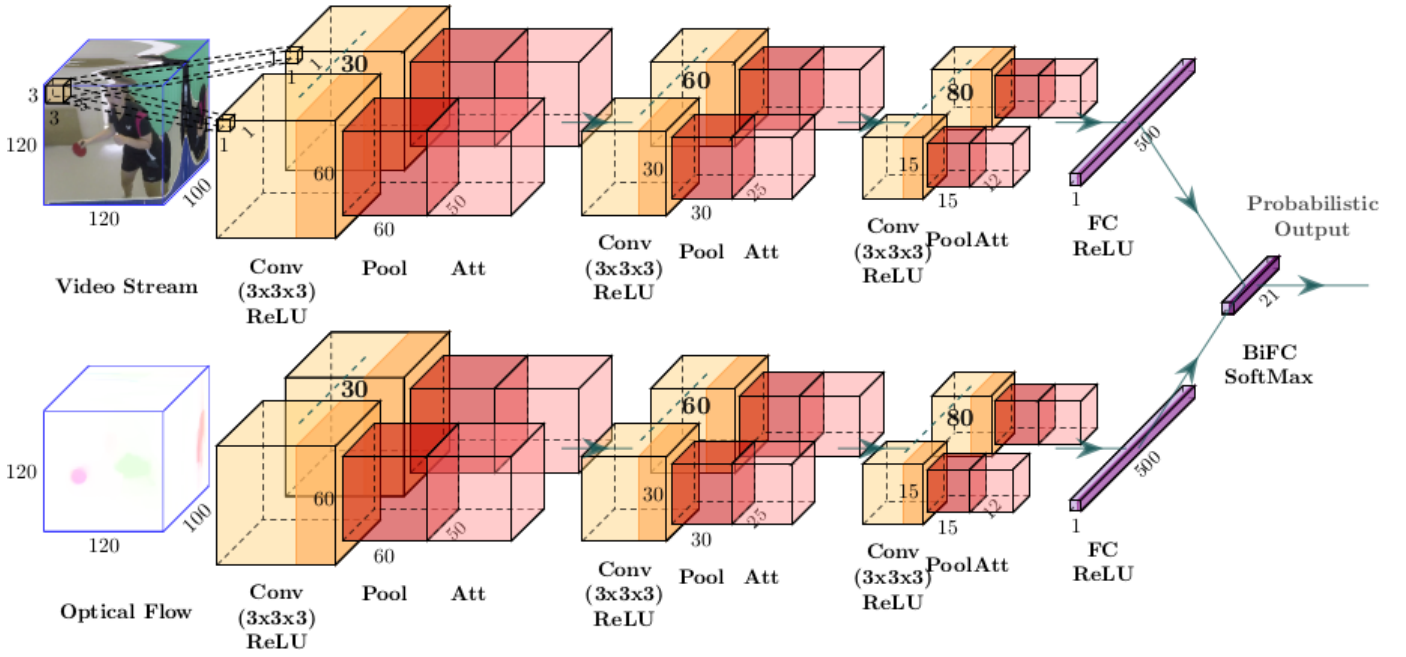


Fig. 1. Twin network architecture with spatio-temporal convolutions and attention mechanism. The number of filters for each convolution are depicted above them.

added residual block after the max pooling layers starting from the first max pooling layer until reaching all the max pooling layers. We did the same with attention blocks so to see the impact of each type of blocks and the impact of their number in the network. We also experimented using separated branches of the network and training them separately, RGB branch denoted as “RGB” model and Flow branch denoted as “Flow” model, to analyse the contribution on each stream. The TSTCNN, denoted as “Twin” model, with 3 attention blocks, is presented on figure 1.

B. 3D attention block

3D attention block, inspired by the work carried out in 2D [12], takes as input a 4D data block of size $(N \times W \times H \times T)$ as illustrated in Figure 2. In this block, all convolution presented uses the same number of filters, N , to maintain the dimension of the processed data. Our input data are processed by a first 3D residual block, denoted as “*res*”, presented in section III-C. Our network then splits in two branches: the trunk branch consisting of 2 successive 3D residual blocks (equation 1) and the soft floating mask branch (lowest position in Figure 2), described through the equations 2, 3, 5. Its role is to accentuate the features generated by the trunk branch. Those two branches are merged as described in equation 6.

$$branch_{trunk}(\cdot) = res(res(\cdot)) \quad (1)$$

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.

$$\begin{aligned} x_1 &= f_{bu}(res(Input)) \\ x_2 &= f_{bu}(x_1) \\ x_3 &= f_{bu}(x_2) \end{aligned} \quad (2)$$

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 [43] used for up-sampling. Two skipped connections are used for collecting information at different scales.

$$\begin{aligned} y_1 &= f_{td}(x_3) + res(x_2) \\ y_2 &= f_{td}(y_1) + res(x_1) \\ y_3 &= f_{td}(y_2) \end{aligned} \quad (3)$$

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 8, followed by a ReLU activation function and a convolution layer with kernel sizes $(1 \times 1 \times 1)$. This is expressed by equation 4:

$$f_{conv}(\cdot) = conv(ReLU(F_n(\cdot))) \quad (4)$$

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 2 and are expressed by equation 5.

$$branch_{fmask}(Input) = Sig(f_{conv}(f_{conv}(y_3))) \quad (5)$$

The output of our trunk branch is then multiplied term by term by $(1 \oplus branch_{fmask}(Input))$ where

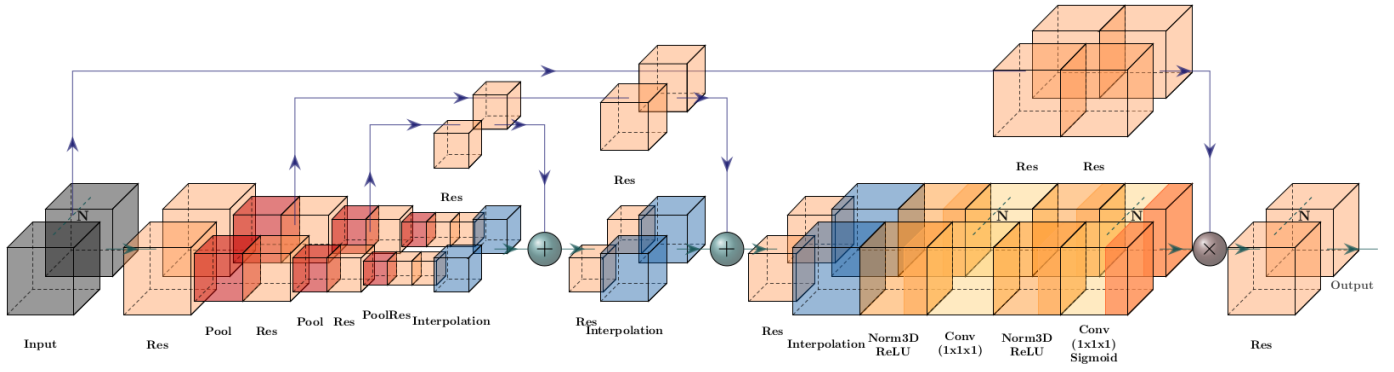


Fig. 2. 3D attention block architecture

$branch_{fmask}(Input)$ 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 6.

$$y = res(branch_{trunk}(Input) \odot (1 \oplus branch_{fmask}(Input))) \quad (6)$$

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

C. 3D residual block

Implemented 3D residual block (Fig. 3) inspired by the work carried out in 2D in [44], takes as input 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 $f_{conv_i}, i = 1, \dots, 3$ (eq. 5). The result of residual block is the sum of the output of these 3 successive layers and our input data:

$$res(x) = f_{conv_3}(f_{conv_2}(f_{conv_1}(x))) + x \quad (7)$$

Here, the first layer f_{conv_1} uses $\frac{N}{4}$ convolution filters of size $(1 \times 1 \times 1)$, the second layer f_{conv_2} uses $\frac{N}{4}$ convolution filters of size $(3 \times 3 \times 3)$. Finally, the third layer f_{conv_3} employs N convolution filters of size $(1 \times 1 \times 1)$.

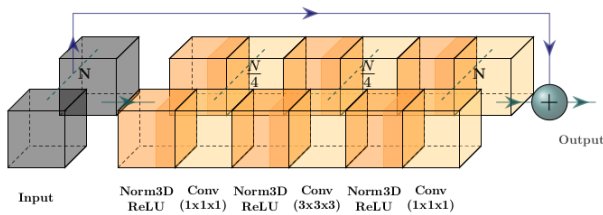


Fig. 3. 3D residual block architecture

The 3D batch normalization, described in [45], 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:

$$f_n(x_i) = \frac{x_i - \mu_i}{\sqrt{\sigma_i^2 \oplus \epsilon}} * \gamma_i \oplus \beta_i \quad (8)$$

with $i = 1, \dots, N_{channels}$, μ_i and σ_i the mean and standard deviation vectors of x_i computed over the batch, γ_i and β_i 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.

IV. EXPERIMENTS AND RESULTS

To assess the efficiency of the attention block for capturing qualitative features for classification task, we compare the classification results of the model with and without attention blocks on the TTStroke-21 dataset [18].

A. TTStroke-21 dataset

Our dataset, entitled TTStroke-21 [18], is composed of recorded table tennis game videos. These sequences are recorded indoors at different frame rates, with artificial light and without markers. The player is filmed in game or training situations, see figure 4a. These videos have been annotated by table tennis players and experts from the Faculty of Sports (STAPS) of the University of Bordeaux, France. A web platform has been developed by our team for this purpose where the annotator locates in time and labels the strokes performed. The annotation platform is presented in figure 4b.



a. Video acquisition

b. Annotation platform

Fig. 4. TTStroke-21 acquisition and annotation platform

In the following experiments, 129 of videos recorded at 120 frames per second have been considered. They represent a total of 1048 table tennis strokes. From these time-segmented table tennis strokes in these 129 videos, 681 negative samples were selected but only 106 were used to maintain class balance.

B. Visualizing the impact of the attention mechanism on features

The attention block highlights features that contribute the most to the classification. In this way, the model can learn faster meaningful features in the classification task. Figure 5 shows outputs of floating mask branch of each attention block for a RGB image input to the “Twin” model (Fig. 1). Feature values range from 0 to 1, but are normalized using min-max normalization and resized for better visualization.




Fig. 5. Visualisation of soft mask branch output from each attention block of the RGB branch of the “Twin” model. The RGB input segmented from the original frame and its class is represented in light blue on a.

RGB input is of size $(100 \times 120 \times 120)$, each dimension representing respectively time, width and height. These parameters were fixed experimentally as a function of video resolution, frame-rate and stroke speed. The output size of the soft mask branch of attention blocks decreases by a factor 2. Figure 5b, c and d represents 3 channels at a specific time. It can be noticed how the network focuses on the table’s edges, on the player and even on the ball (more visible in figure 5c). To classify a stroke, it is important to observe the posture of the player but also his/her position with respect to the table. The ball position and trajectory can also be of high importance to classify the stroke. On the whole training set, output values of the soft mask branch range between 0.35 and 0.65, meaning no features are totally left out or overrated, on the contrary.

C. Convergence of the models

Conducted experiments required to change the hyper parameters used in our previous work [18]. Indeed, the number of parameters to train, which depends on the number of attention or residual blocks, greatly increased compared to our first experiments without attention mechanisms, see table I.

We compared our results with the I3D model [5] which contains around 25M parameters to train. Their model uses

TABLE I
NUMBER OF PARAMETERS TO LEARN ACCORDING TO THE ARCHITECTURE OF THE MODEL*

Models	without attention blocks	with 3 attention blocks
RGB	180 800	498 420
Flow	179 990	497 610
Twin	360 790	996 030

* parameters of the fully connected layers are not considered.

Inception modules introduced in [46] which are combination of different 3D convolutional layers using different filter sizes and concatenating their output. We trained their model according to their instructions with rgb data and optical flow trained separately. The training process differs according to the type of data: a larger number of iterations is required for optical flow, with a specific scheduled learning rate. The output of the two models on the test set can be combined together to improve the performances as shown in table III. We train their model using a time window of $T = 100$ frames which we have selected after several experiments. Note that in our previous work, T was set to 64 frames and therefore performances were limited.

Furthermore, in our case the type of model trained (“RGB”, “Flow” or “Twin”) also influences the training process. Since different combinations and number of blocks were tested, the learning rate during training had to be adapted. A learning rate scheduler was then used, which reduced and increased the learning rate when the observed metric reached a plateau. Weights and state of the model were saved when it was performing the best and we re-loaded when the learning rate was changed. This allowed to re-start the optimization process from the past state with a new step-size in the gradient descent optimizer.

We started training with a learning rate of 0.01. A number of epochs: *patience*, set to 50, was considered before updating the learning rate, unless the performance drastically dropped (in our case: 0.7 of the best validation accuracy obtained).

The metric of interest was the training loss: if its average on the last 25 epochs was greater than its average on the 35 epochs before, the process was re-started from the past state and the learning rate divided by 10 until reaching 10^{-5} . After this step, the learning rate was set back to 0.01 and process continued. These numbers of epochs were set empirically after preliminary experiments.

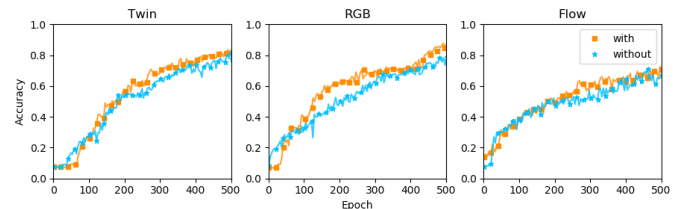


Fig. 6. Evolution of the validation accuracy for the different models with and without attention blocks

It worth mentioning that convergence is slower when us-

ing this method of leaning rate re-scheduler with our past architectures introduced in [18], [19]. Here, this strategy is however efficient to adapt and find adequate learning rates during training for different architecture configurations.

When comparing models with and without attention blocks as illustrated in 6, it can be noticed, that our training process requires less epochs to adapt to our models with attention blocks. The convergence is faster and after the same number of epochs (500) models with attention blocks outperform models without attention (see table II, where best results are depicted in bold for each type of model).

TABLE II
COMPARISON OF THE CLASSIFICATION PERFORMANCES AFTER 500 EPOCHS IN TERMS OF ACCURACY IN %

Models	I3D [5]		without attention		with attention	
	Train	Val.	Train	Val.	Train	Val.
RGB	95.1	59.6	83.9	78.7	96.3	87.8
Flow	97.7	58.7	93	71.8	87.4	72.6
Twin	-	-	88.9	82.6	92.7	83.5

Analysing results further, table II, attention mechanisms seem to be more efficient with the one branch model “RGB” which is 11% better in term of classification accuracy. With other models, the superiority is less obvious but still noticeable with 1% gain. We can notice the gap between training and validation accuracy for the I3D models. This model is far more deeper and has much more parameters to learn. It is therefore more subject to over-fitting on our dataset. The attention blocks applied to our networks, even if greatly increasing the number of parameters, seem to not over-fit. At 500 epochs, the Twin model performs worse than the RGB model and still needs to be trained to outperform, see section IV-D. Its slower convergence can be due to the increased number of parameters to train. It may also come from the batch size used during training, which had to be decreased from 10 to 5 because of resource limitations.

D. Classification performances

Our implemented attention blocks have shown to lead to faster convergence. The different models are also compared for the classification task. This comparison is done after reaching a stable convergence, which is around 1000 epochs using the attention blocks and 1500 without. All the results are reported in table III), along with a comparison with the I3D model [5] and the model introduced in our last work [37]. The Twin model without attention was retrained using the method presented in section IV-C and better performances were also noticeable.

The I3D model still suffers from over-fitting but its performances have greatly improved compared to our previous work [19]. The reason is that the temporal dimension of input data increased from 64 frames to 100 frames. The combination of RGB and Flow models allows a 3% rise of the accuracy on the test set compared to the RGB model alone. The use of attention blocks allows to gain up to 12% with our RGB model and 5% with our Flow and Twin models. It may be

TABLE III
COMPARISON OF THE CLASSIFICATION PERFORMANCES AFTER CONVERGENCE IN TERMS OF ACCURACY %

Models	Train	Val.	Test
RGB - I3D [5]	98	72.6	69.8
RGB [37]	98.6	87	76.7
RGB with Attention	96.9	88.3	85.6
Flow - I3D [5]	98.9	73.5	73.3
Flow [37]	88.5	73.5	74.1
Flow with Attention	96.4	83.5	79.7
RGB + Flow - I3D [5]	-	-	75.9
Twin [37]	89.4	79.1	72.4
Twin (Retrained)	99	86.1	81.9
Twin with Attention	97.3	87.8	87.3

surprising that the Twin could not benefit more from the attention mechanism contrary to its contribution for the other models. From the visualization analysis, figure 5, and results in table III, it can be argued that since the attention mechanism learns to focus on areas where RGB data are changing with respect to time, its contribution is lesser for models fed with temporal information such as the optical flow.

V. CONCLUSION

In this work, we have extended the work carried out in 2D [44], [12] to implement 3D residual blocks and 3D attention blocks. We applied these architectures to fine grained action recognition in video on TTStroke-21.

We have shown that 3D attention blocks enable faster convergence of the models in terms of epochs. They outperform the models without attention blocks compared to our previous work with twin deep neural networks. However the amount of parameters to learn increases and the size of the network increases too, slowing down the training process for each epoch. The new training method also improves our Twin model without attention and performances of our baseline increased when considering longer input samples. After convergence, the Twin model with attention mechanism outperforms all other modalities. According to the visualization of the soft mask branch, it is safe to say that the attention blocks focus on meaningful features such as motion, body parts, position of the player with respect to the table, rackets and ball. We also noticed a greater efficiency of the attention mechanism on RGB data.

We are limited by the number of samples of our datasets and we are still working on enriching TTStroke-21 to add shades to our results. To deeper analyse the contribution of the attention blocks, we provide our implementation online and we are planning to test it on other publicly available datasets for fine grained action recognition [10].

ACKNOWLEDGMENT

The work has been supported by research Grants CRISP of New Aquitania Region.

REFERENCES

- [1] C. Gu, C. Sun, S. Vijayanarasimhan, C. Pantofaru, D. A. Ross, G. Toderici, Y. Li, S. Ricco, R. Sukthankar, C. Schmid, and J. Malik, "AVA: A video dataset of spatio-temporally localized atomic visual actions," *CoRR*, vol. abs/1705.08421, 2017.
- [2] H. Rahmani, A. S. Mian, and M. Shah, "Learning a deep model for human action recognition from novel viewpoints," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 3, pp. 667–681, 2018.
- [3] A. Ullah, J. Ahmad, K. Muhammad, M. Sajjad, and S. W. Baik, "Action recognition in video sequences using deep bi-directional LSTM with CNN features," *IEEE Access*, vol. 6, pp. 1155–1166, 2018.
- [4] G. Varol, I. Laptev, and C. Schmid, "Long-term temporal convolutions for action recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 6, pp. 1510–1517, 2018.
- [5] J. Carreira and A. Zisserman, "Quo vadis, action recognition? A new model and the kinetics dataset," *CoRR*, vol. abs/1705.07750, 2017.
- [6] A. Stoian, M. Ferecatu, J. Benois-Pineau, and M. Crucianu, "Fast action localization in large-scale video archives," *IEEE Trans. Circuits Syst. Video Techn.*, vol. 26, no. 10, pp. 1917–1930, 2016.
- [7] C. Feichtenhofer, H. Fan, J. Malik, and K. He, "Slowfast networks for video recognition," *CoRR*, vol. abs/1812.03982, 2018.
- [8] X. Wang, R. B. Girshick, A. Gupta, and K. He, "Non-local neural networks," in *IEEE CVPR 2018*. IEEE Computer Society, 2018, pp. 7794–7803.
- [9] J. Lin, C. Gan, and S. Han, "TSM: temporal shift module for efficient video understanding," in *IEEE/CVF ICCV 2019*. IEEE, 2019, pp. 7082–7092.
- [10] D. Shao, Y. Zhao, B. Dai, and D. Lin, "Finegym: A hierarchical video dataset for fine-grained action understanding," in *CVPR*. IEEE, 2020, pp. 2613–2622.
- [11] A. M. Obeso, J. Benois-Pineau, M. S. García-Vázquez, and A. A. Ramírez-Acosta, "Forward-backward visual saliency propagation in deep nns vs internal attentional mechanisms," in *IPTA*. IEEE, 2019, pp. 1–6.
- [12] F. Wang, M. Jiang, C. Qian, S. Yang, C. Li, H. Zhang, X. Wang, and X. Tang, "Residual attention network for image classification," in *CVPR*. IEEE Computer Society, 2017, pp. 6450–6458.
- [13] K. Soomro, A. R. Zamir, and M. Shah, "UCF101: A dataset of 101 human actions classes from videos in the wild," *CoRR*, vol. 1212.0402, 2012.
- [14] 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," *CoRR*, vol. abs/1705.06950, 2017.
- [15] J. Wang, K. Zhao, D. Deng, A. Cao, X. Xie, Z. Zhou, H. Zhang, and Y. Wu, "Tac-simur: Tactic-based simulative visual analytics of table tennis," *IEEE Trans. Vis. Comput. Graph.*, vol. 26, no. 1, pp. 407–417, 2020.
- [16] H. Lin, Z. Yu, and Y. Huang, "Ball tracking and trajectory prediction for table-tennis robots," *Sensors*, vol. 20, no. 2, p. 333, 2020.
- [17] R. Voeikov, N. Falaleev, and R. Baikulov, "Ttnet: Real-time temporal and spatial video analysis of table tennis," *CoRR*, vol. abs/2004.09927, 2020.
- [18] P. Martin, J. Benois-Pineau, R. Péteri, and J. Morlier, "Sport action recognition with siamese spatio-temporal cnns: Application to table tennis," in *CBMI 2018*. IEEE, 2018, pp. 1–6.
- [19] —, "Fine grained sport action recognition with siamese spatio-temporal convolutional neural networks," in *Multimedia Tools and Applications*, 2020.
- [20] J. Liu, G. Wang, P. Hu, L. Duan, and A. C. Kot, "Global context-aware attention LSTM networks for 3d action recognition," in *CVPR*. IEEE Computer Society, 2017, pp. 3671–3680.
- [21] J. Cai and J. Hu, "3d rans: 3d residual attention networks for action recognition," *Vis. Comput.*, vol. 36, no. 6, pp. 1261–1270, 2020.
- [22] J. Hu, L. Shen, S. Albanie, G. Sun, and E. Wu, "Squeeze-and-excitation networks," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 42, no. 8, pp. 2011–2023, 2020.
- [23] Y. Chen, Y. Kalantidis, J. Li, S. Yan, and J. Feng, "A²-Nets: Double Attention Networks," in *Advances in Neural Information Processing Systems*, 2018, pp. 352–361.
- [24] N. Dhingra and A. M. Kunz, "Res3atn - deep 3d residual attention network for hand gesture recognition in videos," in *2019 International Conference on 3D Vision, 3DV 2019, Québec City, QC, Canada, September 16-19, 2019*. IEEE, 2019, pp. 491–501. [Online]. Available: <https://doi.org/10.1109/3DV.2019.00061>
- [25] S. Zagoruyko and N. Komodakis, "Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer," in *ICLR (Poster)*. OpenReview.net, 2017.
- [26] F. Huang, X. Zhang, C. Li, Z. Li, Y. He, and Z. Zhao, "Multimodal network embedding via attention based multi-view variational auto-encoder," in *ICMR*. ACM, 2018, pp. 108–116.
- [27] J. Lei, Y. Jia, B. Peng, and Q. Huang, "Channel-wise temporal attention network for video action recognition," in *ICME*. IEEE, 2019, pp. 562–567.
- [28] Y. Du, C. Yuan, B. Li, L. Zhao, Y. Li, and W. Hu, "Interaction-aware spatio-temporal pyramid attention networks for action classification," in *ECCV (16)*, ser. Lecture Notes in Computer Science, vol. 11220. Springer, 2018, pp. 388–404.
- [29] J. Zhao and C. G. M. Snoek, "Dance with flow: Two-in-one stream action detection," in *CVPR*. Computer Vision Foundation / IEEE, 2019, pp. 9935–9944.
- [30] F. Manerba, J. Benois-Pineau, R. Leonardi, and B. Mansencal, "Multiple moving object detection for fast video content description in compressed domain," *EURASIP J. Adv. Sig. Proc.*, vol. 2008, 2008.
- [31] X. Long, C. Gan, G. de Melo, J. Wu, X. Liu, and S. Wen, "Attention clusters: Purely attention based local feature integration for video classification," in *IEEE CVPR 2018*. IEEE Computer Society, 2018, pp. 7834–7843.
- [32] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the inception architecture for computer vision," in *IEEE CVPR 2016*. IEEE Computer Society, 2016, pp. 2818–2826.
- [33] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. S. Bernstein, A. C. Berg, and F. Li, "Imagenet large scale visual recognition challenge," *International Journal of Computer Vision*, vol. 115, no. 3, pp. 211–252, 2015.
- [34] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *CoRR*, vol. abs/1409.1556, 2014.
- [35] J. Cheng, L. Dong, and M. Lapata, "Long short-term memory-networks for machine reading," in *EMNLP 2016*, J. Su, X. Carreras, and K. Duh, Eds. The Association for Computational Linguistics, 2016, pp. 551–561.
- [36] H. Kuehne, H. Jhuang, E. Garrote, T. A. Poggio, and T. Serre, "HMDB: A large video database for human motion recognition," in *ICCV*. IEEE Computer Society, 2011, pp. 2556–2563.
- [37] P. Martin, J. Benois-Pineau, R. Péteri, and J. Morlier, "Optimal choice of motion estimation methods for fine-grained action classification with 3d convolutional networks," in *ICIP 2019*. IEEE, 2019, pp. 554–558.
- [38] K. Simonyan and A. Zisserman, "Two-stream convolutional networks for action recognition in videos," in *NIPS*, 2014, pp. 568–576.
- [39] C. Feichtenhofer, A. Pinz, and A. Zisserman, "Convolutional two-stream network fusion for video action recognition," in *IEEE CVPR 2016*, 2016, pp. 1933–1941.
- [40] X. Chen, K. Kundu, Y. Zhu, H. Ma, S. Fidler, and R. Urtasun, "3d object proposals using stereo imagery for accurate object class detection," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 5, pp. 1259–1272, 2018.
- [41] C. Liu, "Beyond pixels: Exploring new representations and applications for motion analysis," Ph.D. dissertation, Massachusetts Institute of Technology, 5 2009.
- [42] I. Sutskever, J. Martens, G. E. Dahl, and G. E. Hinton, "On the importance of initialization and momentum in deep learning," in *Proceedings of the 30th International Conference on Machine Learning, ICML 2013, Atlanta, GA, USA, 16-21 June 2013*, ser. JMLR Workshop and Conference Proceedings, vol. 28. JMLR.org, 2013, pp. 1139–1147.
- [43] B. Kenwright, "Free-form tetrahedron deformation," in *ISVC (2)*, ser. Lecture Notes in Computer Science, vol. 9475. Springer, 2015, pp. 787–796.
- [44] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," *CoRR*, vol. abs/1512.03385, 2015.
- [45] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *ICML*, ser. JMLR Workshop and Conference Proceedings, vol. 37. JMLR.org, 2015, pp. 448–456.
- [46] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. E. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *IEEE CVPR 2015*, 2015, pp. 1–9.