Res3ATN – Deep 3D Residual Attention Network for Hand Gesture Recognition in Videos

Conference Paper**Author(s):**

Dhingra, Naina; Kunz, Andreas

Publication date:

2019-09

Permanent link:

<https://doi.org/10.3929/ethz-b-000365762>

Rights / license:

[In Copyright - Non-Commercial Use Permitted](#)

Originally published in:

<https://doi.org/10.1109/3DV.2019.00061>

Funding acknowledgement:

177542 - Barrierefreie Besprechungszimmer für sehbehinderte Menschen (SNF)

Res3ATN - Deep 3D Residual Attention Network for Hand Gesture Recognition in Videos

Naina Dhingra, Andreas Kunz
 ETH Zurich, Innovation Center Virtual Reality
 Leonhardstrasse 21, 8092 Zurich, Switzerland
 (ndhingra, kunz)@iwf.mavt.ethz.ch

Abstract

Hand gesture recognition is a strenuous task to solve in videos. In this paper, we use a 3D residual attention network which is trained end to end for hand gesture recognition. Based on the stacked multiple attention blocks, we build a 3D network which generates different features at each attention block. Our 3D attention based residual network (Res3ATN) can be build and extended to very deep layers. Using this network, an extensive analysis is performed on other 3D networks based on three publicly available datasets. The Res3ATN network performance is compared to C3D, ResNet-10, and ResNext-101 networks. We also study and evaluate our baseline network with different number of attention blocks. The comparison shows that the 3D residual attention network with 3 attention blocks is robust in attention learning and is able to classify the gestures with a better accuracy, thus outperforming existing networks.

1. Introduction

In conversations with other people, we use different types of gestures. In such discussions, non-verbal communication (NVC) is an important part since it could carry up to 55% of the overall communication [31, 19, 24]. In total, 136 different gestures were listed by [7]. Thus, recognizing these gestures - and in particular hand gestures - is crucial to understand the implicit communication in the conversation.

While such NVCs are intuitively understood by human beings, they turn out to be difficult to be recognized and interpreted by machines. The application fields of such automated gesture recognition by machines are manifold, they reach from automated gesture recognition for robots [40], understanding the psychological factors, up to the evaluation and recognition of the sign language [41, 30]. Gesture recognition also plays a huge role in advanced driver assistance systems (ADASs) [6]. Here, vision-based hand

gesture detection systems are employed for the interaction of the driver with vehicle. They are used for implementing touch-less sensors [36]. These sensors help the drivers to interact with secondary functions such as music, heating, etc. which improves the safety and comfort while driving. Furthermore, hand gesture recognition have wide applications in various fields, such as, robotic imitation learning, virtual/augmented reality, tele-operations, security operations, etc.

Hand gestures form a part of our conversation and are important to fully understand the topic being discussed. However, visually impaired people are not be able to access these hand gestures and consequently may not be able to easily follow the conversation [23]. Here, a real-time implementations of the hand gesture detection and the corresponding output to the interface of the blind user could help in addressing this issue.

In this paper, we address the hand gesture detection using a deep learning framework. Motivated by the success of the attention networks and recent advances in deep residual networks, we use a special kind of network, i.e. 3D residual attention network which is an end to end trainable network to classify the hand gestures given in the video frames. Similar to the residual attention network for image classification [54], we use multiple layers of attention blocks to generate features which capture different attention at each block.

The reason for using residual networks is to have significant influence on increasing the network depth [27]. This gives the advantage of deeper networks which are easier to train with an increase in accuracy. The reason for using attention mechanisms is that they give relative importance to particular sub-parts of the scenarios. For instance, compared to the human-eye mechanism, our brain ignores the information captured during the saccades and accepts the given useful information during fixation. From the given fixations, the brain gives the importance to the segmented useful information. This mechanism is referred to as an attentive behaviour of brain. Similarly, adapting this attention mechanism can help the residual neural network to classify

hand gestures with better accuracy than without any attention.

The main contributions of this paper can be summarized as follows: (1) We develop an end to end trainable 3D attention-based deep residual neural network which has stacked attention blocks. The attention features adaptively change with the depth of the network. (2) We provide insights in the number of frames from videos to be input in our network to achieve better results. (3) We perform an evaluation of our network on three different data-sets and illustrate the comparison results, (4) We compare our network with other state of the art networks for hand gesture recognition and also perform ablation study on the number and position of the attention blocks in the baseline network. (5) We also provide suggestions to improve the accuracy of our network by varying the parameters. (6) Finally, we also release our PyTorch implementation¹ as an open-source. We also expect that our work will give further advances in hand gesture recognition using attention blocks.

The paper is structured as follows. Related work is discussed in the Section 2. The proposed 3D attention block and residual attention network are described in Section 3. The experiments and ablation study are elaborately illustrated in section 4. Finally, Section 5 concludes our work.

2. Related Work

Machine learning methods such as support vector machines (SVMs), Hidden Markov Models (HMMs), decision trees, random forests, and conditional random fields have been implemented in the state of art to classify the hand gestures [36, 14, 33, 46, 57, 51]. Several feature extraction techniques such as hand crafted spatio-temporal features, histogram of gradient features, classical descriptors, etc. were used in the classical machine learning techniques to recognize the gestures [33, 37].

Deep Learning techniques have been successfully used in a number of different applications of computer vision [15, 44] such as object recognition, image segmentation, image classification, image registration, etc. These techniques have outperformed the classical methods by achieving high benchmark results [55]. Deep learning techniques have also found to be performing well on videos analysis and 3D images analysis for medical purposes. Considering the performance and success of these techniques, we have used them to detect and classify hand gestures on different open source video data sets.

Video analysis and gesture recognition based on deep learning includes mainly *three* techniques based on how the temporal dimensions of the video data are treated [2].

First is a two or more stream data input approach, where two or more streams of data are fed into the 2D convolutional

neural network (2D CNNs) [44, 56, 8, 47]:

- RGB images are encoded as input of one of the two streams.
- Optical flow is encoded to input as the second stream.

Some of the recent work includes feeding more than two streams of the data where the third stream is encoded depth maps, or, extra features [45, 34, 50].

Second is the end to end approach, where 3D video data is fed into the network which uses 3D convolutional layers having 3D filters to capture the features from the 3D data along temporal and spatial dimensions. Compare to the 2D CNNs, 3D CNNs are able to extract discriminative features. 2D CNNs have the advantage of using the pre-trained network which is trained on large available 2D datasets [18]. In this paper, we use a 3D CNN for hand gesture recognition which can take advantage of discriminative features along with temporal and spatial ones.

Third is the combination of 2D or 3D CNNs and temporal sequence modelling [2]. This combination is then applied to a single frame or stacks of frames using recurrent neural networks or long short term memory (LSTM) as they can capture temporal features using recurrent connections. Several variants such as Hierarchical RNNs (H-RNNs), bidirectional RNNs (B-RNNs), etc and HMMs have been successfully used for temporal modelings [20, 17, 58, 53].

There are supervised and unsupervised ways of learning the features from 3D video datasets. Supervised learning uses ground truth data for optimizing the training process. Unsupervised learning includes extracting the invariant spatio-temporal features from the videos using independent subspace analysis (ISA), autoencoders, or some other variant network [3, 39, 35]. Convolutional Restricted Boltzmann Machines (RBMs) have also been used for generating feature representations of the video frames [9].

Attention Mechanism: An attention mechanism works based on the functions of the human eye. As we pay attention to a particular region of the total field of view, our brain is trained to interpret the information based on the higher attention area [30] rather than treating the field of view homogeneously. Similarly, attention mechanism have been explored in deep learning specifically for the combination of image analysis and natural language processing applications such as answer selection tasks [49], image captioning [59], handwriting synthesis [21], machine translation [4], or phoneme recognition [11, 12]. They have also shown to perform well for speech recognition [5] when supplemented with location-awareness [12].

Attention mechanism are of two different types, namely, soft attention and hard attention. *Soft attention* uses the differentiable function such as softmax function, sigmoid function, etc., and can be trained using backpropagation algorithms [43, 54, 59, 30, 47, 10, 29]. Spatial transformer

¹<https://github.com/nainadhingra2012/Res3ATN>

network is an example of soft attention which uses the spatial manipulation of the data [29]. *Hard attention* is based on the non-differentiable stochastic functions such as heavy-side step function, switch functions, etc. which have abrupt changes and discontinuities in the functions [22].

Our network is based on the soft attention developed for an inclusive feed-forward network in a 2D residual attention network for image classification [54]. It uses bottom-up and top-down feed-forward structures which provide soft weights to the features [54].

There has been work done with 3D residual networks [32, 26] in the past for gesture recognition which can be scaled to hundreds of layers. To our knowledge, there is no state-of-the-art attention-based 3D residual network used for hand gesture recognition. The 3D residual attention network will have the benefit from residual learning to have deeper networks along with the positives of attention aware features. The attention blocks will help the network to concentrate on the useful region of the video along the spatial as well as the temporal domain. To achieve a better accuracy for recognizing the gestures using the combination of residual network along with attention mechanism, we built our Res3ATN network.

3. Methodology

3.1. 3D Convolutional Neural Network

3D CNNs capture the information from the spatial content as well as the temporal relationship between the different frames of a video. We use 3D CNNs in the 3D residual attention network and other various networks for comparison. The C3D network is used as a network to compare with and to evaluate the results of 3D residual attention network (Res3ATN).

3.2. Residual Network

The residual network have short connections, which are direct connection between two non-consecutive layers [27]. Since it is easier to optimize the deeper networks using residual connections, it will result in an increased accuracy. Fig. 1 shows the residual blocks used for our residual attention network. We use a ResNet-10 and ResNext-101 architecture which is build with the ResNet and ResNext block as shown in Fig. 1.

For the C3D, ResNet-10 and ResNext-101 models, we use the same architecture as proposed by [32]. But instead of using pre-trained models on Jester datasets, we train our model from scratch on the 3 datasets, i.e., EgoGesture, Jester, and NVIDIA Dynamic Hand Gesture dataset.

3.3. 3D Residual Attention Network

Deep residual networks have shown to perform well for very deep networks [28] with good convergence character-

Figure 1. Basic residual blocks used in ResNet-10, Res3ATN and ResNext-101, respectively.

istics. Our network is a 3D residual network with multiple 3D attention blocks. These attention blocks consists of two parts, i.e., trunk and mask layer. We use the trunk layer which consist of residual units [28, 54], but it can be replaced by any other 3D model units. The mask layer consists of the soft attention differentiable function which creates the 3D mask $M(x)$ of the same dimension as the 3D features $T(x)$ generated by the trunk layer. The output O of attention block can be described by equation (1). The attention block is differentiable and hence it can be easily trained end to end as we use differentiable soft attention.

$$O_{i,c,f} = M_{i,c,f} * T_{i,c,f} \quad (1)$$

where $i \in \{1, \dots, H * W\}$ is the spatial index, H is the height, and W is the width of the 2D frame, $c \in \{1, \dots, C\}$ is the channel index, and $f \in \{1, \dots, F\}$ is the frame index.

The mask layer uses bottom-up and top-down mechanism to get the mask which can manipulate the output features generated by the trunk layer. When the output mask M is multiplied with the T , it gives us the weighted features and hence behaves like a feature selector. During the back-propagation, due to its property of differentiability, it updates the gradient. This network also has a noise filter feature for the noisy labels [54] as shown in equation (2). If the trunk features T are not correct, mask can decide weather to include those T features as there is a multiplication factor of the mask M with partial derivative of T while updating the gradients as shown in equation (2).

$$\frac{\partial M(x, \theta_m) T(x, \theta_t)}{\partial \theta_t} = M(x, \theta_m) \frac{\partial T(x, \theta_t)}{\partial \theta_t} \quad (2)$$

where x is the input, θ_t are trunk layer parameters and θ_m are mask layer parameters. We stack number of attention blocks to filter out different useful features. For the

Figure 2. The 3D Residual Attention Network (Res3ATN) used for hand gesture recognition. The input video is fed to the Res3ATN in the form of a stack of 2D images. BN corresponds to batch normalisation, ReLu to rectified linear units and Maxpool to max pooling. The stacks of residual blocks and attention blocks are used to capture the features from the videos followed by two fully connected layers. The last fully connected layer gives outputs equal to the number of classes in each dataset.

task of hand gesture detection, the network should detect where it should concentrate its attention and then recognize the hand gesture class such as moving hand up, down, right or left, etc. So, the attention block helps the residual attention network to identify the region where the gesture is performed in each frame of a video. The use of multiple attention blocks makes the network more robust as it is able to capture different types of attention focusing at different types of features at each attention block. Since the gesture detection in videos is a difficult task considering the number of features the network has to learn for a single video as compared to a single image, the multiple attention network alleviates this problem by learning multiple masks.

Motivated by the techniques used by [54], to keep the identity function of the residual network, instead of using equation (1) as it is, we add "1" to the generated mask M . This leads to revive the functionality of the residual network and at the same time it does not hamper the performance due to the dot product between trunk output T with the zeros from the mask M . So, it modifies equation (1) to equation (3), which is similarly done by [54] for the 2D images classification.

$$O_{i,c,f} = (1 + M_{i,c,f}) * T_{i,c,f} \quad (3)$$

Similar to the bottom-up and top-down structure in the famous U-Net [42, 13, 16] network used for segmentation tasks, we use the same approach in our attention block which is used for achieving a good mask that could act as a filter to the trunk layer features. The output of the mask layer has the same dimensions as the output from the trunk layer. So, the number of 3D max-pool layers [48], which are used for down-sampling, are the same as for the 3D interpolation layers. As suggested in [54], we used skip connections to connect the bottom-up and top-down structures to get the feature information from various scale levels.

Regarding the 2D results from [54], we have mixed attention, i.e., spatial, channel, and frame attention by using a Sigmoid function [25] for each point in the 3D data x , given by:

$$f(x_{i,c,f}) = \frac{1}{1 + \exp(-x_{i,c,f})} \quad (4)$$

where $i \in \{1, \dots, H * W\}$ is the spatial index, $c \in \{1, \dots, C\}$ is the channel index, and $f \in \{1, \dots, F\}$ is the frame index.

The 3D residual attention network adaptively changes attention as the feature changes. Each attention block learns and captures different types of features which avoids errors or wrong focus of attention. Moreover, a wrong attention predicted by one block can be corrected by the other attention blocks. Thus, multiple blocks make the network quite robust on the attention prediction.

Our architecture of Res3ATN (see Fig. 2) is similar to [54], but we added one extra fully connected layer in their 2D residual attention network. We implemented our network for 3D application of hand gesture in videos so we used 3D-CNN instead of 2D-CNN, 3D Max-pool instead of 2D Max-pool and likewise 3D interpolation and 3D batch-normalisation. Also, the number of filters, kernel size and strides used for each layer are different from the ones used by them as shown in Table 1. The 3 different attention blocks used in our Res3ATN are shown in Fig. 3. We reduce the number of residual blocks, skip connections, up-sampling and downsampling blocks in attention 2 block as compared to attention 1 block and further reduce for attention 3 block.

4. Experiments

In this section, we will elaborately explain our experiments and their results on our 3D residual attention network

Figure 3. The attention blocks used in the Res3ATN Network. There are multiple residual blocks with the same number of up-sampling as down-sampling blocks. There are four skip connections used in the first attention block. We take skip connections as a parameter, so we reduce the skip connections to two in the second attention block in Res3ATN and zero skip connections in the last attention block in the Res3ATN network. For each max pooling 3D layer, we use a kernel size= 3x3x3 and stride size=2, padding=1. For each residual block, we have the same number of output and input channels as well as 3D-CNN. Each residual block has a kernel size equal to 1x1x1, 3x3x3, 1x1x1, respectively, with stride=1. Two consecutive 3D-CNN layers (green colour box in the figure) in the trunk layer have a kernel size=1x1x1, stride=1x1x1.

for hand gesture recognition in different video datasets. The performance of the Res3ATN network is tested on three open sourced datasets: EgoGesture [60], Jester [1], and NVIDIA Dynamic Hand Gesture dataset [38]. We compare and evaluate the performance of the Res3ATN with three other networks, i.e., C3D, ResNet-10, ResNext-101.

4.1. Training Details

To evaluate all the compared networks fairly, we used the same experimental conditions for all of them. We used a learning rate of 0.01, which was kept constant throughout the training. Further, we used Nesterov Stochastic Gradient Descent (SGD) optimizer with a momentum of 0.9 and weight decay of 0.001. We trained the networks for 30 epochs. Next, we used data augmentation techniques for training our networks similar to ones used by [32]. We ran-

Figure 4. Receptive field through the 3D attention block is similar to the receptive field for 2D attention block in [54]. The receptive field decreases from input to the output mask in the mask layer and also similarly decreases for the trunk layer. The output of the trunk and mask layer have the same dimensions.

domly crop the image to a size of 112 x 112 px. Each image is scaled with either 1 , $1/2^{1/4}$, $1/2^{3/4}$ or $1/2$ randomly and then we apply spatial elastic displacement having $\sigma = 2$ and $\alpha = 1$ after cropping the image. We also select the set of defined numbers of frames from the video containing the gestures randomly. The normalization of images is performed to have the pixel values between 0-1 scale value. We use the same training details for all the networks which are used for comparison and evaluation of Res3ATN network.

4.2. Evaluation using EgoGesture Datasets

EgoGesture dataset has 83 classes depicting 83 different hand gestures in videos which are captured in a mixture of diverse indoor and outdoor scenes with 50 different subjects. These videos show interaction of wearable devices with hands. A total of 6 scenes was used, 4 scenes were outdoor and 2 scenes were indoor. The total dataset contains 2,081 RGB-D videos which is divided into training, validation, and testing sets with the 3:1:1 ratio randomly based on the subjects. 1,239 training videos have 14,416 gesture samples, 411 validation videos have 4,768 gesture samples and 431 testing videos have 4,977 gesture samples. Some of gesture examples in the videos are: scroll hand downward, number 0-9, applaud, walk, move hand towards right,

scroll fingers towards left, etc.

EgoGesture is a large egocentric dataset being publicly available. We used depth modality images for our training and evaluation of all the four networks because [32] stated that they achieved better results with EgoGesture depth images than EgoGesture RGB images. The reason for this is that the depth sensors concentrate on the hand motion and neglect the background motion. The hand gesture detection results by all the four networks is described in Table 3. It can be seen that Res3ATN outperforms all the 3 other networks in our experiments. We also investigated the number of input frames for the Res3ATN. The detailed results of input frames is shown in Table 4. It shows that the 32 frames have better results than 8 or 16 input frames. We use 32 frames as input to all the networks.

4.3. Evaluation using Jester Datasets

Jester dataset is a large set of videos with hand gestures performed by humans in front of the webcam or a laptop camera. It is a crowd-sourced data base with large number of people contributing to it. It is publicly available for research purposes. It has 27 classes with total of 148,092 RGB videos. from which 118,562 are used for training, 14,787 videos for validation and 14,743 videos for testing.

Layer	Res3ATN Network
Conv	3x3x3, 64, stride=1
MaxPool3D	3x3x3, stride=2
Residual Block	1x1x1, 32, stride=1 3x3x3, 32, stride=2 1x1x1, 128, stride=1
Attention Block 1	128
Residual Block	1x1x1, 64, stride=1 3x3x3, 64, stride=2 1x1x1, 256, stride=1
Attention Block 2	256
Residual Block	1x1x1, 128, stride=1 3x3x3, 128, stride=2 1x1x1, 512, stride=1
Attention Block 3	512
Residual Block	1x1x1, 256, stride=1 3x3x3, 256, stride=2 1x1x1, 1028, stride=1
Residual Block	1x1x1, 256, stride=1 3x3x3, 256, stride=1 1x1x1, 1028, stride=1
Residual Block	1x1x1, 256, stride=1 3x3x3, 256, stride=1 1x1x1, 1028, stride=1
Residual Block	1x1x1, 512, stride=1 3x3x3, 512, stride=1 1x1x1, 2048, stride=1
Average3D Pool	2x2x2, stride=2
Fully Connected layer	512
Fully connected layer	output classes

Table 1. Res3ATN Network configuration.

	EgoGesture	Jester	NVIDIA dataset
Classes	83	26	25
Total	2,081	148,092	1532
Train	1239	118,562	1050
Valid	411	14,787	-
Test	431	14,743	482

Table 2. Overview of EgoGesture, Jester and NVIDIA dynamic hand gesture dataset.

Model	Frames	Modality	Top-1 acc	Top-5 acc
C3D	32	Depth	92.11	98.20
ResNet-10	32	Depth	91.56	97.84
ResNext-101	32	Depth	92.24	98.32
Res3ATN-56	32	Depth	93.63	98.65

Table 3. Comparison of the Res3ATN model with different network configurations using the EgoGesture dataset.

Some of the examples of the classes are: Pulling Hand In, Pushing Two Fingers Away, Shaking Hand, Sliding Two Fingers Up, Swiping Down, etc.

We investigated the performance of C3D, ResNet-10, ResNext-101 and Res3ATN using the Jester dataset. The

No. of frames	8	16	32
Top-1 acc	79.28	88.38	93.63

Table 4. Comparison of performance of Res3ATN with different number of input frames.

Model	Frames	Modality	Top-1 acc	Top-5 acc
C3D	32	RGB	75.20	85.16
ResNet-10	32	RGB	78.22	87.03
ResNext-101	32	RGB	82.24	89.01
Res3ATN	32	RGB	84.56	91.95

Table 5. Comparison of Res3ATN with different models for the Jester dataset.

Model	Frames	Modality	Top-1 acc	Top-5 acc
C3D	32	RGB	53.94	71.16
ResNet-10	32	RGB	56.74	76.03
ResNext-101	32	RGB	51.24	64.00
Res3ATN	32	RGB	62.65	81.95

Table 6. Comparison of Res3ATN with different network configurations for NVIDIA Dynamic Hand Gesture dataset.

Jester dataset is the largest of all the three datasets. We use the RGB modality videos of this dataset. We noted that ResNext-101 performed better than the ResNet-10 and C3D because of the larger network depth in ResNext-101, which is able to learn more descriptive features for the large dataset. Our network, Res3ATN performs better than the other three networks as shown in Table 5.

4.4. Evaluation using NVIDIA Dynamic Hand Gesture Dataset

NVIDIA Dynamic Hand Gestures dataset has 25 hand gesture types. Each video is recorded by multiple sensors from different viewpoints. A total of 1532 dynamic hand gesture videos are captured by 20 subjects in an indoor environment, varying the intensity of the light using a car simulator with different lightning conditions. The dataset contains 70 % training videos, i.e., 1050 videos and 30 % testing videos, i.e., 482 videos. The split is created randomly on the data based on the subjects. Some of the gestures are: showing the two or three fingers, shaking the hand, pushing the hand up, rotating two fingers clockwise or counterclockwise, etc.

We investigated the performance of Res3ATN network using NVIDIA Dynamic Hand Gesture dataset, which is a very small dataset compared to the EgoGesture and Jester datasets. Because of the small size, the performance of all four networks is low as compared to the performance of the other two datasets. The detailed results are shown in Table 6. However, the evaluation accuracy can be increased by using pre-trained networks on large datasets and then train on the NVIDIA Dynamic Hand Gesture dataset.

Model	Attention blocks	Top-1 acc	Top-5 acc
Baseline with 0 ATN	0	51.45	64.10
Baseline with 1 ATN	1	52.38	67.73
Baseline with 2 ATN	2	52.70	68.26
Res3ATN	3	62.65	81.95

Table 7. Comparison of networks having different number of attention blocks. The baseline network is the Res3ATN without any of the attention blocks, Baseline with 1 ATN is the baseline network with 1 attention block. Baseline with 2 ATN is the baseline network with 2 attention blocks.

Model	Attention blocks	Top-1 acc	Top-5 acc
1_0	1	47.30	62.03
1_1	1	52.38	67.73
1_2	1	51.86	65.87

Table 8. Comparison of networks having 1 attention block at different positions in a baseline network. The model name 1_0 refers to the baseline network the having first attention block of the Res3ATN. 1_1 refers to the baseline network having the second attention block of the Res3ATN. Similarly, 1_2 refers to the network having the third attention block of the Res3ATN.

5. Ablation study on multiple attention blocks

In this section, we compare Res3ATN with its baseline network which has same structure as Res3ATN but has zero attention blocks. We evaluate networks having different numbers of attention blocks and also investigate results with different positions of attention blocks for the same network. The results are shown in Table 7. It is evident that the Res3ATN performs better than the network having 1 or 2 blocks. The reason for the better performance of Res3ATN is that multiple attention blocks tend to capture the correct corresponding attention even if one of the attention blocks is capturing wrong features. Hence, results into making Res3ATN robust for the prediction of attention.

We also compared the location of the attention block in case of 1 attention block in the baseline network. Table 8 indicates that the location of second attention block of the network in Figure 3 performs better than the first and the third attention, when used only one attention block at a time.

While evaluating the networks having 2 attention blocks at a time, we analyzed that the baseline network having the first and the third attention block performs better than the one having the first and the second attention block, and the network having the second and the third attention block as described in Table 9.

6. Further Improvements

Further improvements are possible by using larger input frames than just 112 x 112 px size, since an increase of the spatial resolution and the number of frames can improve

Model	Attention blocks	Top-1 acc	Top-5 acc
2_0	2	50.83	66.49
2_1	2	52.70	68.26
2_2	2	51.66	65.45

Table 9. Comparison of networks with 2 attention blocks at different positions in a baseline network. The model name 2_0 refers to the baseline network having the first two attention blocks of the Res3ATN. 2_1 refers to the baseline network having the first and third attention blocks of the Res3ATN, while 2_2 refers to the network having the second and third attention blocks of the Res3ATN.

the classification accuracy [52, 26]. Due to GPU limitations and considering the size of the Res3ATN, we only used 112 x 112 px size. With the use of batch-normalization in the networks, it is important to train models with large batch-sizes [8, 26]. We use batch-normalization in our network architecture, so it is better to have larger batch-size for training. However, we only used a small batch-size of 6 due to GPU limitations. Since activity recognition is a very related task to hand gesture recognition, using pre-trained models trained on activity recognition datasets such as activity-net or UCF-101 might help to increase the accuracy of the hand gesture recognition network.

7. Conclusion

In this paper, we have successfully developed Res3ATN, i.e., 3D attention based residual network for hand gesture recognition. We validated the performance of Res3ATN on three publicly available hand gesture datasets. The proposed techniques performs better than the three other widely used networks, i.e., C3D, ResNet-10, ResNext-101 for 3D activity and gesture recognition. We investigated the number of frames which should be input to the Res3ATN. We inspected the number of attention blocks and the position of the attention blocks to be used in the network. Our analysis shows that the stacked multiple soft attention blocks help the network to recognize the hand gestures with better accuracy. For future work, we will use pre-trained networks which are trained on different activity recognition datasets and evaluate the performance of Res3ATN.

Acknowledgements

This work has been supported by the Swiss National Science Foundation (SNF) under the grant no. 200021E_177542 / 1. It is part of a joint project between TU Darmstadt, ETH Zurich, and JKU Linz with the respective funding organizations DFG (German Research Foundation), SNF (Swiss National Science Foundation) and FWF (Austrian Science Fund).

References

- [1] <https://20bn.com/datasets/jester/v1>.
- [2] M. Asadi-Aghbolaghi, A. Clapes, M. Bellantonio, H. J. Escalante, V. Ponce-López, X. Baró, I. Guyon, S. Kasaei, and S. Escalera. A survey on deep learning based approaches for action and gesture recognition in image sequences. In *2017 12th IEEE international conference on automatic face & gesture recognition (FG 2017)*, pages 476–483. IEEE, 2017.
- [3] M. Baccouche, F. Mamalet, C. Wolf, C. Garcia, and A. Baskurt. Spatio-temporal convolutional sparse auto-encoder for sequence classification. In *British Machine Vision Conference 2012*, pages 1–12, 2012.
- [4] D. Bahdanau, K. Cho, and Y. Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- [5] D. Bahdanau, J. Chorowski, D. Serdyuk, P. Brakel, and Y. Bengio. End-to-end attention-based large vocabulary speech recognition. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4945–4949. IEEE, 2016.
- [6] F. Baradel, C. Wolf, J. Mille, and G. W. Taylor. Glimpse clouds: Human activity recognition from unstructured feature points. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 469–478, 2018.
- [7] C. R. Brannigan and D. A. Humphries. Human non-verbal behavior, a means of communication. *Ethological studies of child behavior*, pages 37–64, 1972.
- [8] J. Carreira and A. Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017.
- [9] B. Chen. *Deep learning of invariant spatio-temporal features from video*. PhD thesis, University of British Columbia, 2010.
- [10] L.-C. Chen, Y. Yang, J. Wang, W. Xu, and A. L. Yuille. Attention to scale: Scale-aware semantic image segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3640–3649, 2016.
- [11] J. Chorowski, D. Bahdanau, K. Cho, and Y. Bengio. End-to-end continuous speech recognition using attention-based recurrent nn: First results. *arXiv preprint arXiv:1412.1602*, 2014.
- [12] J. K. Chorowski, D. Bahdanau, D. Serdyuk, K. Cho, and Y. Bengio. Attention-based models for speech recognition. In *Advances in neural information processing systems*, pages 577–585, 2015.
- [13] Ö. Çiçek, A. Abdulkadir, S. S. Lienkamp, T. Brox, and O. Ronneberger. 3d u-net: learning dense volumetric segmentation from sparse annotation. In *International conference on medical image computing and computer-assisted intervention*, pages 424–432. Springer, 2016.
- [14] N. H. Dardas and N. D. Georganas. Real-time hand gesture detection and recognition using bag-of-features and support vector machine techniques. *IEEE Transactions on Instrumentation and Measurement*, 60(11):3592–3607, 2011.
- [15] J. Donahue, L. Anne Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2625–2634, 2015.
- [16] H. Dong, G. Yang, F. Liu, Y. Mo, and Y. Guo. Automatic brain tumor detection and segmentation using u-net based fully convolutional networks. In *annual conference on medical image understanding and analysis*, pages 506–517. Springer, 2017.
- [17] Y. Du, W. Wang, and L. Wang. Hierarchical recurrent neural network for skeleton based action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1110–1118, 2015.
- [18] C. Feichtenhofer, A. Pinz, and A. Zisserman. Convolutional two-stream network fusion for video action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1933–1941, 2016.
- [19] C. Frith. Role of facial expressions in social interactions. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1535):3453–3458, 2009.
- [20] F. A. Gers, N. N. Schraudolph, and J. Schmidhuber. Learning precise timing with lstm recurrent networks. *Journal of machine learning research*, 3(Aug):115–143, 2002.
- [21] A. Graves. Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850*, 2013.
- [22] K. Gregor, I. Danihelka, A. Graves, D. J. Rezende, and D. Wierstra. Draw: A recurrent neural network for image generation. *arXiv preprint arXiv:1502.04623*, 2015.
- [23] S. Guenther, R. Koutny, N. Dhirga, M. Funk, C. Hirt, K. Miesenberger, M. Muehlhaeuser, and A. Kunz. Mapvi: Meeting accessibility for persons with visual impairments. In *PERvasive Technologies Related to Assistive Environments*, pages 343–352, New York, USA, 2019. ACM.
- [24] O. Gupta, D. Raviv, and R. Raskar. Deep video gesture recognition using illumination invariants. *arXiv preprint arXiv:1603.06531*, 2016.
- [25] J. Han and C. Moraga. The influence of the sigmoid function parameters on the speed of backpropagation learning. In *International Workshop on Artificial Neural Networks*, pages 195–201. Springer, 1995.
- [26] K. Hara, H. Kataoka, and Y. Satoh. Learning spatio-temporal features with 3d residual networks for action recognition. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3154–3160, 2017.
- [27] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [28] K. He, X. Zhang, S. Ren, and J. Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer, 2016.
- [29] M. Jaderberg, K. Simonyan, A. Zisserman, et al. Spatial transformer networks. In *Advances in neural information processing systems*, pages 2017–2025, 2015.
- [30] C. Kingan, J. Owoyemi, and K. Hashimoto. Point attention network for gesture recognition using point cloud data. In *British Machine Vision Conference 2018, BMVC*

- 2018, Northumbria University, Newcastle, UK, September 3-6, 2018, page 118, 2018.
- [31] M. L. Knapp, J. A. Hall, and T. G. Horgan. *Nonverbal communication in human interaction*. Cengage Learning, 2013.
 - [32] O. Köpüklü, A. Gunduz, N. Kose, and G. Rigoll. Real-time hand gesture detection and classification using convolutional neural networks. *CoRR*, abs/1901.10323, 2019.
 - [33] J. J. LaViola Jr. An introduction to 3d gestural interfaces. In *ACM SIGGRAPH 2014 Courses*, page 25. ACM, 2014.
 - [34] M. Ma, H. Fan, and K. M. Kitani. Going deeper into first-person activity recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1894–1903, 2016.
 - [35] F.-J. H. MarcAurelio Ranzato, Y.-L. Boureau, and Y. LeCun. Unsupervised learning of invariant feature hierarchies with applications to object recognition. In *Proc. Computer Vision and Pattern Recognition Conference (CVPR07)*. IEEE Press, volume 127, 2007.
 - [36] P. Molchanov, S. Gupta, K. Kim, and J. Kautz. Hand gesture recognition with 3d convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 1–7, 2015.
 - [37] P. Molchanov, S. Gupta, K. Kim, and K. Pulli. Multi-sensor system for driver’s hand-gesture recognition. In *2015 11th IEEE international conference and workshops on automatic face and gesture recognition (FG)*, volume 1, pages 1–8. IEEE, 2015.
 - [38] P. Molchanov, X. Yang, S. Gupta, K. Kim, S. Tyree, and J. Kautz. Online detection and classification of dynamic hand gestures with recurrent 3d convolutional neural network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4207–4215, 2016.
 - [39] K. Nandakumar, K. W. Wan, S. M. A. Chan, W. Z. T. Ng, J. G. Wang, and W. Y. Yau. A multi-modal gesture recognition system using audio, video, and skeletal joint data. In *Proceedings of the 15th ACM on International conference on multimodal interaction*, pages 475–482. ACM, 2013.
 - [40] K. Nickel and R. Stiefelhausen. Visual recognition of pointing gestures for human–robot interaction. *Image and vision computing*, 25(12):1875–1884, 2007.
 - [41] L. Pigou, S. Dieleman, P.-J. Kindermans, and B. Schrauwen. Sign language recognition using convolutional neural networks. In *European Conference on Computer Vision*, pages 572–578. Springer, 2014.
 - [42] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
 - [43] P. H. Seo, Z. Lin, S. Cohen, X. Shen, and B. Han. Progressive attention networks for visual attribute prediction. *arXiv preprint arXiv:1606.02393*, 2016.
 - [44] K. Simonyan and A. Zisserman. Two-stream convolutional networks for action recognition in videos. In *Advances in neural information processing systems*, pages 568–576, 2014.
 - [45] S. Singh, C. Arora, and C. Jawahar. First person action recognition using deep learned descriptors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2620–2628, 2016.
 - [46] T. Starner, J. Weaver, and A. Pentland. Real-time american sign language recognition using desk and wearable computer based video. *IEEE Transactions on pattern analysis and machine intelligence*, 20(12):1371–1375, 1998.
 - [47] S. Sudhakaran and O. Lanz. Attention is all we need: Nailing down object-centric attention for egocentric activity recognition. *arXiv preprint arXiv:1807.11794*, 2018.
 - [48] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015.
 - [49] M. Tan, C. d. Santos, B. Xiang, and B. Zhou. Lstm-based deep learning models for non-factoid answer selection. *arXiv preprint arXiv:1511.04108*, 2015.
 - [50] Y. Tang, Y. Tian, J. Lu, J. Feng, and J. Zhou. Action recognition in rgb-d egocentric videos. In *2017 IEEE International Conference on Image Processing (ICIP)*, pages 3410–3414. IEEE, 2017.
 - [51] P. Trindade, J. Lobo, and J. P. Barreto. Hand gesture recognition using color and depth images enhanced with hand angular pose data. In *2012 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI)*, pages 71–76. IEEE, 2012.
 - [52] G. Varol, I. Laptev, and C. Schmid. Long-term temporal convolutions for action recognition. *IEEE transactions on pattern analysis and machine intelligence*, 40(6):1510–1517, 2017.
 - [53] V. Veeriah, N. Zhuang, and G.-J. Qi. Differential recurrent neural networks for action recognition. In *Proceedings of the IEEE international conference on computer vision*, pages 4041–4049, 2015.
 - [54] F. Wang, M. Jiang, C. Qian, S. Yang, C. Li, H. Zhang, X. Wang, and X. Tang. Residual attention network for image classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3156–3164, 2017.
 - [55] L. Wang, Y. Qiao, and X. Tang. Action recognition with trajectory-pooled deep-convolutional descriptors. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4305–4314, 2015.
 - [56] L. Wang, Y. Xiong, Z. Wang, Y. Qiao, D. Lin, X. Tang, and L. Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In B. Leibe, J. Matas, N. Sebe, and M. Welling, editors, *Computer Vision – ECCV 2016*, pages 20–36. Cham, 2016. Springer International Publishing.
 - [57] S. B. Wang, A. Quattoni, L.-P. Morency, D. Demirdjian, and T. Darrell. Hidden conditional random fields for gesture recognition. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’06)*, volume 2, pages 1521–1527. IEEE, 2006.

- [58] D. Wu, L. Pigou, P.-J. Kindermans, N. D.-H. Le, L. Shao, J. Dambre, and J.-M. Odobez. Deep dynamic neural networks for multimodal gesture segmentation and recognition. *IEEE transactions on pattern analysis and machine intelligence*, 38(8):1583–1597, 2016.
- [59] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhutdinov, R. Zemel, and Y. Bengio. Show, attend and tell: Neural image caption generation with visual attention. *arXiv preprint arXiv:1502.03044*, 2015.
- [60] Y. Zhang, C. Cao, J. Cheng, and H. Lu. Egogesture: a new dataset and benchmark for egocentric hand gesture recognition. *IEEE Transactions on Multimedia*, 20(5):1038–1050, 2018.