Recognition of Human Activity in Domestic Environments Using Convolution Neural Networks

Tiago Ferreira,

Abstract—**With the advance of technology robots are becoming essential in humans daily lives. In order to make robots more autonomous lots of efforts are being made by researchers to create robust and efficient image processing algorithms capable of interacting with the real world. This thesis addresses a semi real-live situation by applying convolution neural networks. The case study concept is the classification of humans activities performed in a domestic scenario. The main goal is to determine if the human real-time continuous behavior is either normal or abnormal (potentially dangerous). The creation of such model would allow instant and continuously surveillance in every connected camera. In this thesis an extensive dataset was created and applied to different net architectures to determine which obtain better results. The final results are very positive showing that convolution neural networks can successfully be applied to image processing situations obtaining results that easily equalise or even overcome classic methods from computer vision.**

Keywords*—*convolution-neural-networks, image-processing, machine-learning

I. Introduction

Artificial Neural Networks appeared approximately at the same time as the field of Machine Learning but initially they were highly criticize and they were only used with the input and output layers (no hidden layers were used). In 1986 the ANNs reappeared but this time using intermediate layers (hidden layers). The most used networks were the multi-layer perceptrons with the algorithm of back propagation. The net could emulate a significant number of functions. However the algorithm was slow and not very efficient. In 2005 Deep Learning emerged. As the name suggests, the base in Deep Learning is in the use of deeper networks with much more layers and units per layer. Deep learning kept on growing with the creation of new powerful algorithms and techniques and with the increase of computational power. When trying to apply neural networks to more complex datasets, such as image processing, come the necessity of developing methods that could accept 2D inputs instead of only 1D. One of those methods is known as Convolutional Neural Networks (CNNs). Convolutional Neural Networks are mostly used in image and video processing. The main challenge in image processing is the size of the inputs. For instance, an image with a resolution of 100x100 pixels will have 10000 inputs times 3 channels (RGB). The first layers of a CNN are 2D and make the convolution between the image and the learned kernels (filters). The final layers are 1D (fully connected). The basic idea of CNNs is that the net aims to learn which kernels weights are the most effective. Each layer represents the input image in a different way. The first layers are considered more low level since they extract elementary visual features such as oriented edges, end-points and corners. In deeper layers this features are then combined in order to detect high-order features

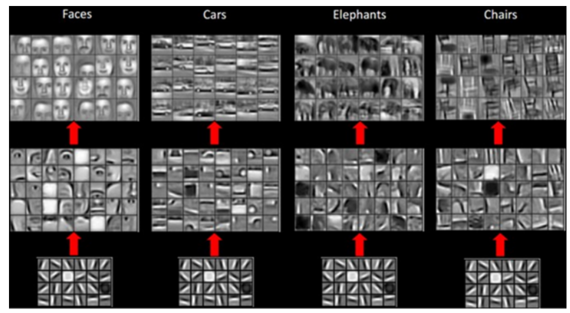

Fig. 1. Feature extraction of kernels along the multiple layers [1]

Manuscript received April 15, 2016.

A. Related Work

Microsoft released in 2012 a new version of their Microsoft Audio Video Indexing Service speech system based on deep learning[2]. These authors managed to reduce the word error rate on four major benchmarks by about 30 percent compared to state-of-the-art models based on Gaussian mixtures for the acoustic modeling and trained on the same amount of data (309 hours of speech).

The main focus of this study relies on applying Convolution Neural Networks into a real dataset recorded in the ISR test bed. These techniques have been applied in a big variety of datasets; not only in image processing but also in other areas, for example speech recognition. The Convolution Neural Networks were introduced by Fukushima et al.[3] and further developed by Lecun et al. [4] in 1998. Since then, many other researchers have published papers presenting results that outperformed the older methods that used hand-designed features selectors like SIFT [5] and HOG [6]. In Krizhevsky et al. [7], a deep Convolutional Neural Network was trained to classify 1.2 million high resolution images in the ImageNet LSVRC/2010 contest into 1000 different classes. Their network achieved top-1 and top-5 [1] test error rates of 37.5% and 17.0% where the best performance achieved during the ILSVRC-2010 competition was 47.1% and 28.2% with an approach that averages the predictions produced from

[1]top-x error, also known as rank-x error, is the error taking in consideration if the correct classification is present in the top x predictions with the higher level of confidence

six sparse/coding models trained on different features [8]. Krizhevsky et al. [7] used a network containing eight learned layers (five convolutional and three fully-connected). In order to train the network it was needed two NVIDIA GTX 580 GPU with 3GB of memory. The network was trained for 90 cycles through a training set of 1.2 million images, which took five to six days. In Girshick et al [9] it was proposed an object detection method that consisted in three modules: The first generates category-independent region proposals, the second is a large convolutional neural network that extracts a fixed-length feature vector and the third is a set of class specific linear SVMs (Support Vector Machines). In this paper the CNNS were use as a mean of feature extraction and not classifiers. The authors used a method that combined classical tools from computer vision with deep learning, achieving results of 30% relative improvement over the best previous results on PASCAL VOC 2012.

All the above information is relative to image processing, but as the purpose of this work is to classify behaviors from video sequences it is important to refer some work that has been done in the video processing area. CNNs can be modified to accept videos sequences as input. This can be achieved by modifying the filters on the first convolutional layer in the single-frame model by extending them to be of size $i * j * 3 * T$, where i,j are the pixels positions on the 3 the channels RGB and T represents some temporal extent. Andrej Karpathy et al. [10] provides an empirical evaluation of CNNS on a large scale video classification using a new dataset of 1 million YouTube videos belonging to 487 classes. In this work the authors observed that motion aware network clearly benefits from motion information in some cases. However these seem to be relatively uncommon because in most of the cases single-frame model already displays very strong performance.

B. Challenge Description

This thesis challenges can be categorized in two major topics: First, the study and understanding of deep learning and convolution neural networks, focusing on finding the best way to implement them and the best architectures to use. Second, applying this algorithms to a realistic scenario in order to determine how likely convolution neural networks can contribute for the increasing of automation in the area of robotics.

The realistic challenge of this study is to develop a model that correctly identifies abnormal behaviors of humans in a domestic scenario. The model must learn the correct classification using a labeled dataset containing both normal and abnormal behaviors.

As described in the RoCKIn Nutshell [11]: ”RoCKIn@Home is a competition that aims at bringing together the benefits of scientific benchmarking with the attraction of scientific competitions in the realm of domestic service robotics. The objectives are to bolster research in service robotics for home applications and to raise public awareness of the current and future capabilities of such robot systems to meet societal challenges like healthy ageing and longer independent living. The basic idea of the scenario of the competition is that there is a elderly, named Granny Annie, who lives in an apartment together with some pets. Granny Annie is suffering from typical problems of aging people and the robot must be able to perform a big variety of tasks that help to increase her life quality. The environment is an ordinary European apartment equipped with a big variety of networkable devices such as: ceiling lamps, electric shutters, camera-based intercom at the front door and a net of surveillance cameras spread around the flat (figure 1.2).

Fig. 2. Virtual representation of @home competitions scenario [11]

If Granny Annie is not in conditions of standing up the robot has to be able to handle visitors, who arrive and ring the door bell. The robot must correctly decide which visitors are allowed to come in and which ones are not. Even if the visitor is allowed to come in, the robot must always be aware of his actions in order to make sure that the visitor is not doing nothing out of the ordinary.

II. Background

The architecture of a CNN can be divided in two parts. The first layers are 2D and make the convolution between the image and the learned kernels. After the convolution layers the net is constituted by 1D fully-connected layers. This last layers are basic artificial neural networks.

CNNS have four base concepts: local connectivity, parameter sharing, pooling and sub-sampling and convolution.

A. Local Connectivity

Local connectivity is the idea of connecting units to local receptive fields on the input. This idea goes back to the MLPs and was almost simultaneous with Hubel and Wiesel’s discovery of locally-sensitive, orientation-selective neurons in the cat’s visual system [12].

Local connectivity allows layers not to be fully connected to each other but only to be connected to a limited number of units in the next layer. This property allows the local processing of an image.

However distortion or shifts of the input can cause the position of salient features to vary. In addition, elementary feature detectors of salient features might be useful across the entire image. This can be accomplished by forcing a set of units, whose receptive fields are located at different places on the image, to have identical weights. This propriety is called parameter sharing.

B. Parameter Sharing

Parameter Sharing is a property that allows each feature map to be applied to all the receptive fields of the input image. Therefore the weights of each feature map are updated based on all the connected receptive fields allowing a significant reduction on the net number of weights.

C. Pooling and Sub-Sampling

Once the feature has been detected, its exact location becomes less important. It is only relevant the approximate position relative to other features. A good way to illustrate this statement is to think in the characters classification: once we know that the input image contains the endpoint of a roughly horizontal segment in the upper left area, a corner in the upper right area, and the endpoint of a roughly vertical segment in the lower portion of the image, we can tell the input image is a 7. All the other features beyond these are irrelevant and might even contribute to a poorer classification. Pooling and Sub-Sampling is a way to reduce the precision of the feature map in order to highlight the most representative features. The pooling technique takes from a patch of hidden units the most representative value. There are different kinds of pooling but the most used are the max-pooling and the average-pooling. The max-pooling as the name suggests takes the maximum value in a certain neighborhood:

$$y(i,j) = maxN(x(i,j)) \tag{1}$$

where $x(i,j)$ is the value of the pixel located in i,j and $N(x)$ is the neighborhood of $x(i,j)$

The average-pooling (equation 2), again as the name suggests, takes the average value of a certain neighborhood,

$$y(i,j) = (\frac{1}{m^2}) * \sum N(x(i,j)) \tag{2}$$

where m^2 is the squared value of the size of the reception field

The reason why it is also called sub-sampling is because there is a reduction in the feature map resolution. Besides contributing to the elimination of irrelevant features [4] this technique also contributes to save a lot of weights allowing to increase the number o kernels of the next layer without substantial increase of the processing power.

D. Convolution

The feature maps are created by the discrete convolution of the kernels with a portion of the input (receptive field).

The discrete convolution is very similar to the continuous one but with the difference that it is computed in a finite sequence of points.

$$\sum_{m=-inf}^{+inf} f[m]g[n-m] \tag{3}$$

Equation 3 describes the convolution between $f[m]$ and $g[m]$. The result will be the sum of the product of $f[m]$ and the reversed and shit n of $g[m]$ for every m.

The matrix convolution has some significant differences from the processes described above. Matrix convolution (figure 4) is used in image processing and it is between an image ($I[u,v]$) and a kernel ($k[i,j]$).

$$FM[i,j] = \sum N(I[u,v]) * k[i-u, j-v]) \tag{4}$$

The first step to calculate the convolution is to flip the rows and the columns of the Kernel, which means that instead of counting the kernel index from top right as being (0,0) we place the origin index in the bottom right corner. The next step is to overlap the center of the kernel with the image pixel that we want to calculate ($I[i,j]$) and calculate the sum of the overlapped values product.

An helpful property of convolution is that it is shift invariant so if the neighborhood stays the same the image can be shifted and the result will always be very similar. If the input has three channels the output is going to be the sum of the result of the convolution in each channels. Then the result will be the input of an non-linear function (ie: $tanh$). The new feature map will be constituted by the outputs of the non-linear function.

E. CNN backpropagation

The algorithm of backpropagation in CNN's works in a very similar way than with 1D neural networks. Let's analyse how backpropagation works in convolution units and pooling units.

1) Convolution Units: The algorithm of back-propagation in CNNs its actually very similar to the one applied to fully connected layers. The aim of each iteration is to obtain a new set of kernel weights ($k[i,j]$) that reduces the model error.

$$k(n+1) = k(n) + lr * \frac{\delta E}{\delta k} \tag{5}$$

where lr is the learning rate and $\frac{\delta E}{\delta k}$ is the derivative of the loss function in order of the weights k.

In order to calculate $\frac{\delta E}{\delta k}$ it is also needed to calculate the inputs of the backward net. Instead of determining $y(i,j)$ by making the convolution between $x(i,j)$ and $k(i,j)$ we want to calculate the opposite:

$$x(i,j) = y(i,j) * k(i,j) \tag{6}$$

Finally to calculate $\frac{\delta E}{\delta k}$ we need to calculate the convolution between the input of the unit in the feedforward net and the input of the same unit in the backward net.

2) Pooling Units: For the case of the pooling units let's analyse both the maxpooling and the average pooling situations. In the case of the maxpooling the error derivative will only be backpropagated to the units that were selected as the maximum in the forward propagation. All the other units will receive the value of zero.

```
Result: Backforward maxpooling
if i and j correspond to the max unit then
 |  x(i,j) = y(a,b)
else
 |  x(i,j) = 0;
end
```

In the case of the average pooling all the units will receive the same gradient value (upsampling) divided by the m^2 (m^2 is the squared value of the size of the reception field)

$$x(i,j) = \frac{1}{m^2} * y(a,b) \tag{7}$$

F. Dropout

Dropout was introduced by Srivastava et al.[13] as a "Simple Way to Prevent Neural Networks from Overfitting". When training a neural network, specially when having limited training data, there are some noise in the samples that might not be in the real data. This may cause the model to overfit to that noise. Dropout reduces substantially the overfitting by preventing each unit to depend on their nearest neighbours. The term "dropout" refers from dropping out units which consists in temporarily removing them from the network while the training is occurring. This is accomplished by, for each training iteration, randomly setting to zero the weights of multiple units with a probability of p (the value of p is normally 0.25 or 0.5). The higher the value of p the higher the number of dropped units and consequently units depend less on their neighbours.

At test time the entire network is used to make a prediction however its weights are a scale down version of the trained weights. That can be accomplished by multiplying them by the value of p used during the training process. Dropout units can also be used with the backprogation algorithm with gradient descent. In this case some of the units are dropped out and a thinned version of the network is used for the back and forward propagation. Each dropped out unit contributes with a gradient of 0 for the connected parameters.

III. Convolutional Net Characteristics

n this thesis two different structures were taken in consideration: SmallNet and BigNet. As the names suggests the big difference between these two structures is their size. The SmallNet has only one convolution layer while the BigNet follows some characteristics used in the paper of Alex Krizhevsky et al.[7] having five convolution layers. The main purpose on using these two different structures was to find how the number of layers could interfere with the final results

A. SmallNet

As it is represented in the fig III-A the SmallNet convolution section is constituted by one convolution layer with only one 3x3 kernel, followed by a non-linearity layer (relu), a maxpooling unit with a reducing scale of (2,2) and finally a dropout unit.

The fully-connected section is constituted by one fully-connected layer that receives 4536 inputs and has 100 outputs, followed by a non-linearity layer (relu), another fully-connected layers that is the output layer receiving 100 inputs and having $nclasses$ outputs. Where $nclasses$ depends on the used dataset configuration.

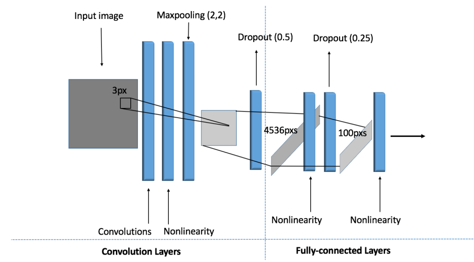

Fig. 3. Smallnet Structure

B. BigNet

As it is represented in the fig III-B the BigNet convolution section is constituted by 5 convolution layers with 1,3,6,9 and 12 kernels where the first kernel is a 9x9 the second a 6x6 and the rest are 3x3. Each convolution layer is followed by the same layers as the SmallNet. The fully-connected structure is also similar to the SmallNet one having two dense layers where the first converts 36 inputs in 100 outputs and the second 100 inputs in $nclasses$ outputs.

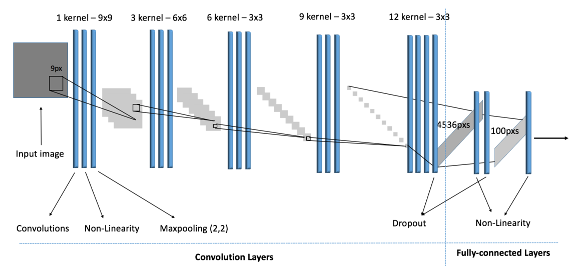

Fig. 4. Bignet Structure

IV. Case Study

The scenario of this case study occurs in the apartment kitchen and the visitor is a plumber that comes to fix the kitchen sink. While in the kitchen the subject has a lot of different actions from which some are considered normal and others abnormal. The kitchen scenario is not static and changes through the videos in order to avoid the creation of patterns that associate actions to the position of the object.

A. Obtaining dataset

The first step for building the dataset was to record a big variety of actions performed in the kitchen of the ISR test bed scenario (see figure 5). Generally all the recorded frames needed to have a big variety of actions and features in order to avoid over-fitting (model become stuck in a local minimum). In order to guarantee the variety of the dataset multiple actions were recorded with 4 different subjects (three male and one female), where the same subject would use different clothes and the disposal of the scenario would change from video to video. The result was a total of 471 actions where each action was constituted by 30 frames summing up to a total of 14130

Fig. 5. Kitchen scenario without the table

frames. In order to make the dataset flexible enough so it could be split into various classes each action was classified with a specific label such as: 'sit', 'sit on the floor','sit on the table','sit on the sink', 'walk','run','steal', 'throw a chair','throw object','throw','fight','move arms','jump','shake', 'fix', 'drink' and 'stop'. It's important to notice that each of this labels contain multiple videos where this actions were performed multiple times by the subjects. Bellow there are some examples of the collected frames.

Fig. 6. Male subject 1 performing the action of jumping

Fig. 7. Male subject 1 performing the action of fixing the sink tap

Fig. 8. Male subject 1 performing the action of fixing under the sink

Fig. 9. Male subject 2 performing action of stealing

Fig. 10. Female subject moving her arms

Of course that the fact of giving to each action such a specific label allowed the clustering of labels into different classes. In this thesis 3 different datasets were created by clustering the labels in different global classes.

1) Two classes dataset: This dataset was constituted only by to classes: normal and abnormal.

- normal = 'sit', 'sit on the floor', 'walk','run','fix', 'drink' and 'stop'
- abnormal = 'sit on the table', 'sit on the sink', 'run', 'steal', 'throw a chair', 'throw object', 'throw', 'fight',

'move', 'arms', 'jump', 'shake'.

2) Three classes dataset:

- fast movements = 'shake','throw the chair','throw object', 'run' 'throw','fight', 'move the arms', 'jump'
- slow movements = 'walk', 'fix','steal', 'drink'
- stop movements = 'stop','sit', 'sit on the floor','sit on the table','sit on the sink'

3) Four classes dataset:

- seated = 'sit', 'sit on the floor','sit on the table','sit on the sink'
- moving = 'walk','run','steal'
- fighting = 'throw chair','throw object','throw','fight','move arms','jump','shake'
- still = 'fix', 'drink', 'stop'

V. Training Process

A. Splitting the dataset

In the literature the dataset is normally divided into three types: training set (a subset that is used for training the model and that is propagated through the net), the validation set (that is also used during the training process but to implement the criteria of early stopping, generating an more independent error) and finally the test set (used after the training to measure the model error against data never used in the training process)

Since the dataset used in this thesis was not big enough the validation set was not used. Instead, it was only used the training set and the test set (the test set had both role of test set and validation set). The method used to generate both sets from the dataset is as follows:

1) Queue all frames (grouped by class) in an array
2) Start iterating through the actions, from left to right, splitting the dataset so that 1/3 was test set and the rest 2/3 training set. Which means that 1 for every 3 actions would be test set and the others training set.

It is important to notice that each action/video is constituted by 30 frames. The result is two datasets that even though both having videos with the same label each of them have different frames.

B. Avoiding Overfitting

As the dataset used is small it was really important to avoid overfitting at all costs because if not the model could only be applied to the current training set. The first technique used was dropout layers (as described before). The technique of early stopping was also used. To implement early stopping for each iteration we should measure the training error. During the training one should measure both the training error and the testing error. Both should be decreasing in each iteration. If at some point the testing error starts to increase even thought the training error is decreasing that means the we are evolving to a over fitted model. The idea behind early stopping is to stop the training when the combination of the training error and the testing error are in their lowest point. To implement that both errors were measure in all the iterations and the final model was select from the one with the lowest error combination.

VI. Results

In this section the results for the three analysed datasets will be presented for the two used net structures. In this thesis it was used the library Keras[14] to implement the training and prediction using CNNs.

A. Obtaining Error Rate

To better understand the results bellow it is important to describe how they were obtained. Even though the inputs of the nets are single frames and the output does not depend on past frames, this case study involves video processing and consequently its classification should depend on the classification of all the frames in it. To accomplished that the frames belonging to a certain video were fed as input to the model in the same order as they appear, the global classification of the video was the average of all the belonging frames classification. All the frames of the dataset are divided in sets of 30. This means that the frames from 1 to 30 belong to the same action and then frames from 31 to 60 belong to a different action.

Taking that in consideration the process to obtain the error rate was:

```
Result: Error rate
while nFrame less then size(dataset) do
    if nFrame is multiple of 30 then
        top1 = average(30 last most probable
          classifications);
        top2 = average(30 last second most probable
          classifications);
        if top1 != ground truth then
            error1 ++;
            if top2 != ground truth then
                error2 ++;
            end
        end
    else
        nFrame ++;
    end
end
Error1 = error1/Total number of frames * 100;
Error2 = error2/Total number of frames * 100;
```

B. Two Classes Dataset

This dataset was constituted by only two classes: Normal and Abnormal. And it was the most simple case used in this study. The software developed in this thesis allowed, in an automated way, to test different variations in the training parameters in order to see which combinations generated the best results. The parameters tested in the train were the learning rate (using values of 0.01, 0.05 and 0.2) and the batch size (using values of 30, 100 and 500). In this study was also tested how different weights initializations could influence the results. Two different initializations were tested:

- The "lecun uniform"[15] initialization. This algorithm of initialization states that the weights should not be to large so that the sigmoid would saturate resulting

in small gradients and consequently in a slow learning process. Neither to small causing the same effect. The weights should be intermediate so that the sigmoid is primarily active in its linear region.
- Pre-trained initialization. This type of initialization states that applying weights obtained in the training of other models, will lead to better and faster results. If both models are applied to similar datasets.

In the graphs above is shown the error percentage for the multiple training variations performed in this thesis. There are four different error types:

- Ts n p - Testset without pre-training
- Ts p - Testset with pre-training
- Trs n p - Training set without pre-training
- Trs p - Training set with pre-training

C. SmallNet Results

For the SmallNet the pre-trained was obtain by training the same net in a smaller dataset (less frames) and then applying the resulting weights as initialization of the train with the entire dataset. This process helps the net to converge faster although with more epochs the training with no initialization would probably reach the same results.

D. BigNet Results

BigNet was trained layer by layer. Which means that firstly all convolution layers but the first one were removed and trained for 10 epochs, then the second one was added and the first layer was initialized with the just learned weights. The process is repeated for the rest of the layers. As can be seen in the graphics above, without weights initialization the training has poor results and never converges however with the weights initialization the error rate decreases significantly.

E. Results Analysis

SmallNet	30	100	500	**BigNet**	30	100	500
0.01	35,2%	35,9%	40,1%	0.01	30,2%	40,1%	40,1%
0.05	33%	32,3%	43,6%	0.05	30,2%	40,1%	40,1%
0.2	40,1%	40,1%	29,5%	0.2	40,1%	40,1%	29,5%

TABLE I. SUMMARIZED RESULTS OF THE ERROR RATE OF THE TEST SET WITHOUT PRE-TRAINING

SmallNet	30	100	500	**BigNet**	30	100	500
0.01	18,3%	29,3%	41,2%	0.01	39,5%	41,2%	41,2%
0.05	4,8%	12,2%	32,2%	0.05	41,2%	41,2%	41,2%
0.2	41,2%	41,2%	11,8%	0.2	41,2%	41,2%	29,7%

TABLE II. SUMMARIZED RESULTS OF THE ERROR RATE OF THE TRAIN SET WITHOUT PRE-TRAINING

From both net results it can be concluded that in order to increase the learning rate the batch size has also to be increased. This makes sense because having a higher batch size means that the algorithm uses a bigger variety of inputs in each update so the learning rate can be reduced because each update is more moderate.

SmallNet	30	100	500	**BigNet**	30	100	500
0.01	15,4%	16,9%	23,9%	0.01	29,5%	23,9%	35,2%
0.05	18,3%	12,6%	17,6%	0.05	33,8%	28,8%	27,4%
0.2	40,1%	40,1%	15,4%	0.2	23,9%	59,8%	25,3%

TABLE III. SUMMARIZED RESULTS OF THE ERROR RATE OF THE TEST SET WITH PRE-TRAINING

SmallNet	30	100	500	**BigNet**	30	100	500
0.01	9,7%	20,4%	28,5%	0.01	6,5%	17,9%	35,9%
0.05	2,4%	6,9%	19,5%	0.05	2,8%	28,8%	25,7%
0.2	41,2%	41,2%	11,8%	0.2	19,1%	58,7%	12,6%

TABLE IV. SUMMARIZED RESULTS OF THE ERROR RATE OF THE TRAIN SET WITH PRE-TRAINING

Some cases different batch sizes have the exact same error rate. This is due to both combinations had encountered the same local minimum, not being able to reduce the error any more after that.

VII. Three Classes Dataset

In order to raise the complexity of the tests other datasets were created which included more classes. Here are described the results for the dataset with three classes: ”stop movements”, ”slow movements”, ”fast movements”. After determining that the best results were obtain with a lower batch size and learning rate the values used were 0.01 for the learning rate and 30 for the batch size. In this situation using the same methods of pre-initialization used for the two classes dataset proofed to be inefficient leading to poor results. The best results were obtain by using the *lecun uniform* initialization.

A. SmallNet Results

For the SmallNet were obtain for the training error: 2.5477% for top-1 error and 0.955% for top-2 error. Regarding the test error were: 26.7515% for top-1 error and 7.6433% for top-2 error.

This model was trained using 400 epochs using the early-stop criteria to determine when to stop.

B. BigNet Results

This model was trained using 600 epochs using the early-stop criteria to determine when to stop. The error rate obtain for the test set was, in the top-1 33,12% and for the top-2 was 8,91%. For the training set the obtain error was 12.738% for the top-1 and 3.18% for the top-2.

C. Results Analysis

SmallNet	Top 1	Top 2	**BigNet**	Top 1	Top 2
test set	26,7%	7,6	test set	33%	8,9%
training set	2,5%	0,9%	training set	12,7%	3,1%

TABLE V. SUMMARIZED RESULTS OF *Three Classes Dataset*

Even though the training error increased comparing to the two classes dataset the results are still very positive. It is also very interesting to see how the results of the SmallNet are as good or even better than the results of the Bignet meaning that this case study does not requires a lot of neurons to create a model that represents this dataset. Although the "lecun uniform" initialization tries to find intermediate weights the algorithm is still random so in some situations the training process would be stuck in a local minimum. In that case the only solution was to restart the training in order to obtain a better initialization.

VIII. Four Classes Dataset

This last dataset was also the most complex of this thesis. It was divided into 4 different classes: "still", "sited", "moving" and "fighting". The training process was identical to the one used for the "Three Classes Dataset".

A. SmallNet Results

The results of the SmallNet were very positive. The error rate for the test set was 21,65% for the top-1 and 7% for the top-2. Regarding the training set the error rate was 0,636% for the top-1 and 0,32% for the top-2.

B. BigNet Results

Although the results were a little worse than the SmallNet results they are still very positive. The error rate for the test set was 31,21% for the top-1 and 15,92% for the top-2. Regarding the training set the error rate was 20,06% for the top-1 and 7% for the top-2.

C. Results Analysis

SmallNet	Top 1	Top 2	**BigNet**	Top 1	Top 2
test set	21,6%	7%	test set	31,2%	15,9%
training set	0,6%	0,3%	training set	20%	7%

TABLE VI. SUMMARIZED RESULTS OF *Four Classes Dataset*

The results of this dataset turned out to be very positive even with a small dataset in comparison with the state of the art ones. Once again the results from the SmallNet are better than the results from the BigNet.

IX. Global Results Analyses

In this thesis the same dataset was divided into different classes, starting with only two classes until four different ones. In the "Two classes dataset" results a big variety of combinations were tested and the conclusion was that having a small learning (0.01) rate together with an also small batch size (30) turn out to be the combination that led to better results. Choosing a batch size of 30 turned out to be very efficient mostly because each action of the dataset was constituted by 30 frames. Another interesting discover was that the SmallNet structure achieved better results than the BigNet which goes against the majority of the literature where the bigger the net the better the results. This results are also very positive because the training and predicting process of the SmallNet is much faster than the BigNet. Of course that the size of the net depends on the complexity of the case study and probably bigger datasets with more classes would require bigger nets.

X. Thesis Conclusions

In this thesis the same dataset was used in three different situations by changing how the data was split. The model was tested against a dataset divided into two, three and four classes. The results obtain for each case were positive and allowed taking the following conclusions: Firstly it is definitely possible to apply CNNs into the current case study. The model will autonomously find a way to classify the dataset with less error possible. However there are no thumb rules that determine the structure of the net that should be used to each case study so in order to find the best model is required to make a several tests with different net structures and parameters. Secondly not always a deeper net obtains better results. For this thesis case study the SmallNet structure obtain better results in less time than the bigger net, bigger nets should only be used in more complex cases where the number of outputs are bigger or for a situation that requires to give more relevance to small details. Thirdly keeping both training error and test error in mind while training the dataset turned to be very useful to find the best models. Using techniques like early-stopping or cross-validation is very important to obtain the best models. Fourthly in video processing it is very important to obtain a big and miscellaneous dataset. By doing so it will avoid the appearance of patterns that lead to over-fit and prevent the model from obtaining good results in real life situations. As a last conclusion the results of this thesis show once more the potential of autonomously learning methods to solve complex problems where human intuition is involved. There is a lot of good documentation and machine learning libraries that allow researches to quickly start testing their models and obtain good results.

XI. Future Work

There are multiple ways to expand this work. In this thesis the dataset was only split into training set and test set. Using a validation set in the training process could improve the results, also the method of cross-validation could also be applied. In order to make sure that the model would obtain good classifications in real live scenarios it is important to keep improving the dataset by collecting more videos with more variety and to test the model in online mode (which means measuring the model error by making live classifications). In the future the model should be capable of doing "online learning". This means that it should be continuously learning and collecting data and could be corrected by an Human every time that it made poor classifications. Finally, as in video processing, the past frames are important for the present classifications. Using recurrent neural networks could also improve the results and allow more complex associations between frames.

REFERENCES

[1] Kernels along the layers.

[2] Frank Seide, Gang Li, and Dong Yu. Conversational speech transcription using context-dependent deep neural networks. In *Interspeech 2011*. International Speech Communication Association, August 2011.

[3] Kunihiko Fukushima. Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological Cybernetics*, 36:193–202, 1980.

[4] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, Nov 1998.

[5] David G. Lowe. Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vision*, 60(2):91–110, November 2004.

[6] Navneet Dalal and Bill Triggs. Histograms of oriented gradients for human detection. In *Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05) - Volume 1 - Volume 01*, CVPR '05, pages 886–893, Washington, DC, USA, 2005. IEEE Computer Society.

[7] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C.J.C. Burges, L. Bottou, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 1097–1105. Curran Associates, Inc., 2012.

[8] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015.

[9] Ross B. Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. *CoRR*, abs/1311.2524, 2013.

[10] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. Large-scale video classification with convolutional neural networks. In *CVPR*, 2014.

[11] Rocking@home competition.

[12] D. H. HUBEL and T. N. WIESEL. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of physiology*, 160:106–154, January 1962.

[13] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15:1929–1958, 2014.

[14] Keras documentation.

[15] Yann Lecun, Lon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. In *Proceedings of the IEEE*, pages 2278–2324, 1998.