

Emotion and Gender Classification using Convolutional Neural Networks

¹V ILANGO, ²Naresh Kumar

¹Head of Department, ²Student

¹Master of Computer Application,

¹New Horizon college of Engineering, Bengaluru, India

Abstract: In this paper we propose an implement a general convolutional neural network (CNN) building framework for designing real-time CNNs. We validate our models by creating a real-time vision system which accomplishes the tasks of face detection, gender classification and emotion classification simultaneously in one blended step using our proposed CNN architecture. After presenting the details of the training procedure setup we proceed to evaluate on standard benchmark sets. We report accuracies of 96% in the IMDB gender dataset and 66% in the FER-2013 emotion dataset. Along with this we also introduced the very recent real-time enabled guided back-propagation visualization technique. Guided back-propagation uncovers the dynamics of the weight changes and evaluates the learned features. We argue that the careful implementation of modern CNN architectures, the use of the current regularization methods and the visualization of previously hidden features are necessary in order to reduce the gap between slow performances and real-time architectures.

I. INTRODUCTION

The success of service robotics decisively depends on a smooth robot to user interaction. Thus, a robot should be able to extract information just from the face of its user, e.g. identify the emotional state or deduce gender. Interpret-ing correctly any of these elements using machine learning (ML) techniques has proven to be complicated due the high variability of the samples within each task [4]. This leads to models with millions of parameters trained under thousands of samples [3]. Furthermore, the human accuracy for classifying an image of a face in one of 7 different emotions is 65% 5% [4]. One can observe the difficulty of this task by trying to manually classify the FER-2013 dataset images in Figure 1 within the following classes “angry”, “disgust”, “fear”, “happy”, “sad”, “surprise”, “neutral”.



Fig. 1: Samples of the FER-2013 emotion dataset [4].



Fig. 2: Samples of the IMDB dataset [9].

The implementations have been validated in a real-time facial expression system that provides face-detection, gender classification and that achieves human-level performance when classifying emotions..

Furthermore, CNNs are used as black-boxes and often

learned features remain hidden, making it complicated to establish a balance between their classification accuracy and unnecessary parameters. Therefore, we implemented a real-time visualization of the guided-gradient back-propagation proposed by Springenberg [11] in order to validate the features learned by the CNN.

II. RELATED WORK

Commonly used CNNs for feature extraction include a set of fully connected layers at the end. Fully connected layers tend to contain most of the parameters in a CNN. Specifically, VGG16 [10] contains approximately 90% of all its parameters in their last fully connected layers. Recent architectures such as Inception V3 [12], reduced the amount of parameters in their last layers by

including a Global Average Pooling operation. Global Average Pooling reduces each feature map into a scalar value by taking the average over all elements in the feature map. The average operation forces the network to extract global features from the input image. Modern CNN architectures such as Xception [1] leverage from the combination of two of the most successful experimental assumptions in CNNs: the use of residual modules [6] and depth-wise separable convolutions [2]. Depth-wise separable convolutions reduce further the amount of parameters by separating the processes of feature extraction and combination within a convolutional layer.

Furthermore, the state-of-the-art model for the FER2013 dataset is based on CNN trained with square hinged loss. This model achieved an accuracy of 71% [4] using approximately 5 million parameters. In this architecture 98% of all parameters are located in the last fully connected layers.

The second-best methods presented in [4] achieved an accuracy of 66% using an ensemble of CNNs.

III. MODEL

We propose two models which we evaluated in accordance to their test accuracy and number of parameters. Both models were designed with the idea of creating the best accuracy over number of parameters ratio. Reducing the number of parameters help us overcoming two important problems. First, the use of small CNNs alleviate us from slow performances in hardware-constrained systems such robot platforms. And second, the reduction of parameters provides a better generalization under an Occam’s razor framework. Our first model relies on the idea of eliminating completely the fully connected layers. The second architecture combines the deletion of the fully connected layer and the inclusion of the combined depth-wise separable convolutions and residual. modules. Both architectures were trained with the ADAM optimizer [8].

Following the previous architecture schemas, our initial architecture used Global Average Pooling to completely remove any fully connected layers.

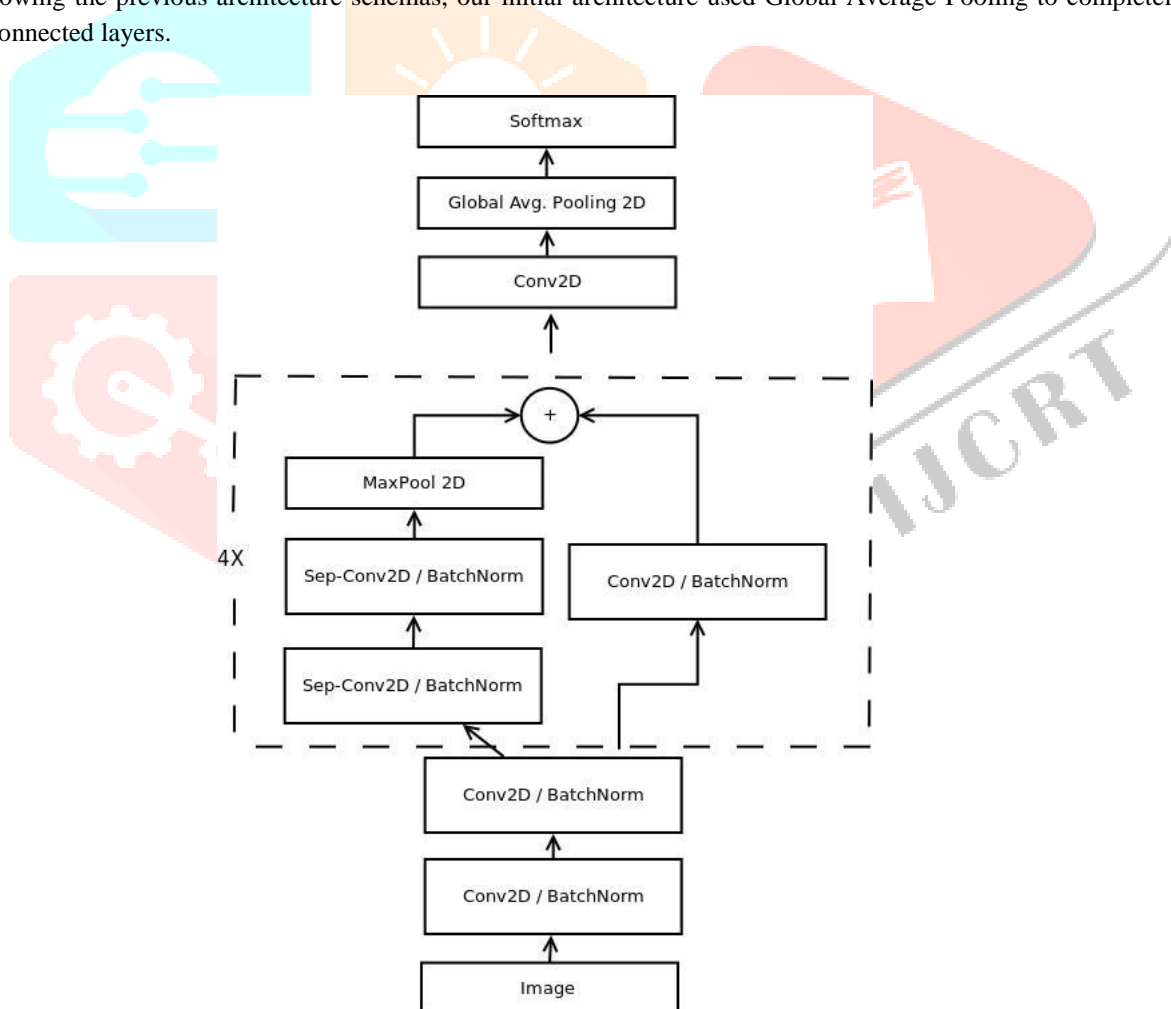


Fig. 3: Our proposed model for real-time classification.

each reduced feature map. Our initial proposed architecture is a standard fully-convolutional neural network composed of 9 convolution layers, ReLUs [5], Batch Normalization [7] and Global Average Pooling. This model contains approximately 600,000 parameters. It was trained on the IMDB gender dataset, which contains 460,723 RGB images where each image belongs to the class “woman” or “man”, and it achieved an accuracy of 96% in this dataset. We also validated this model in the FER-2013

dataset. This dataset contains 35,887 greyscale images where each image belongs to one of the following classes “angry”, “disgust”, “fear”, “happy”, “sad”, and “surprise, neutral”. Our initial model achieved an accuracy of 66% in this dataset. We will refer to this model as “sequential fully-CNN”.

Our second model is inspired by the Xception [1] architecture. This architecture combines the use of residual modules [6] and depth-wise separable convolutions [2]. Residual modules modify the desired mapping between two subsequent layers, so that the learned features become the difference of the original feature map and the desired features. Consequently, the desired features $H(x)$ are modified in order to solve an easier learning problem $F(X)$ such that:

$$H(x) = F(x) + x \quad (1)$$




Fig. 4:

Since our initial proposed architecture deleted the last fully-connected layer, we reduced further the amount of parameters by eliminating them now from the convolutional layers. This was done through the use of depth-wise separable convolutions. Depth-wise separable convolutions are composed of two different layers: depth-wise convolutions and point-wise convolutions. The main purpose of these layers is to separate the spatial cross-correlations from the channel cross-correlations [1]. They do this by first applying a $D \times D$ filter on every M input channels and then applying $N \times 1 \times 1$ convolution filters to combine the M input channels into N output channels. Applying $1 \times 1 \times M$ convolutions combines each value in the feature map without considering their spatial relation within the channel. Depth-wise separable convolutions reduces the computation with respect to the standard convolutions by a factor of $N^1 + D^1_2$ [2]. A visualization of the difference between a normal Convolution layer and a depth-wise separable convolution can be observed in Figure 4.

Our final architecture is a fully-convolutional neural net-work that contains 4 residual depth-wise separable convolutions where each convolution is followed by a batch normalization operation and a ReLU activation function. The last layer applies a global average pooling and a soft-max activation function to produce a prediction. This architecture has approximately 60; 000 parameters; which corresponds to a reduction of 10 when compared to our initial naive implementation, and 80 when compared to the original CNN. Figure 3 displays our complete final architecture which we refer to as mini-Xception. This architecture obtains an accuracy of 95% in gender classification task. Which corresponds to a reduction of one percent with respect to our initial implementation? Furthermore, we tested this architecture in the FER-2013 dataset and we obtained the same accuracy of 66% for the emotion classification task. Our final architecture weights can be stored in an 855 kilobytes file. By reducing our architectures computational cost we are now able to join both models and use them consecutively in the same image without any serious time reduction. Our complete pipeline including the openCV face detection module, the gender classification and the emotion classification takes 0:22 0:0003 ms on a i5-4210M CPU. This corresponds to a speedup of 1:5 when compared to the original architecture of Tang.

We also added to our implementation a real-time guided back-propagation visualization to observe which pixels in the image activate an element of a higher-level feature map. Given a CNN with only ReLUs as activation functions for the intermediate layers, guided-back propagation takes the derivative of every element $(x; y)$ of the input image I with respect to an element $(i; j)$ of the feature map f^L in layer L . The reconstructed image R filters all the negative gradients; consequently, the remaining gradients are chosen such that they only increase the value of the chosen element of the feature map. Following [11], a fully ReLU CNN reconstructed image in layer l is given by:

$$R_{i;j}^l = (R_{i;j}^{l+1} > 0) R_{i;j}^{l+1}$$

common misclassification such as persons with glasses being classified as “angry”. This happens since the label “angry” is highly activated when it believes a person is frowning and frowning features get confused with darker glass frames

Model	train acc	val acc	test acc
Baseline	0.25	0.25	0.24
Five-layer CNN	0.46	0.37	0.39
Deeper CNN	0.60	0.47	0.48
VGG16 fine-tuned CNN	-	0.29	-
VGGFace fine-tuned CNN	0.37	0.38	-

Table (1)

A comparison of the training, validation, and test accuracies of our different models when using the best parameters we found. Due to time constraints. We omitted some fine-tuning results that we were not able to run due to time constraints.

IV. RESULTS

Results of the real-time emotion classification task in un-seen faces can be observed in Figure 5. Our complete real-time pipeline including: face detection, emotion and gender classification.

An example of our complete pipeline can be seen in Figure 6 in which we provide emotion and gender classification. In Figure 7 we provide the confusion matrix results of our emotion classification mini-Xception model. We can observe several common misclassifications such as predicting “sad” instead of “fear” and predicting “angry” instead “disgust”.

A comparison of the learned features between several emotions’ and both of our proposed models can be observed in Figure 8. The white areas in figure 8b correspond to the pixel values that activate a selected neuron in our last convolution layer. The selected neuron was always selected in accordance to the highest activation. We can observe that the CNN learned to get activated by considering features such as the frown, the teeth, the eyebrows and the widening of one’s eyes, and that each feature remains constant within the same class. These results reassure that the CNN learned to interpret understand-able human-like features, that provide generalizable elements These interpretable results have helped us understand several

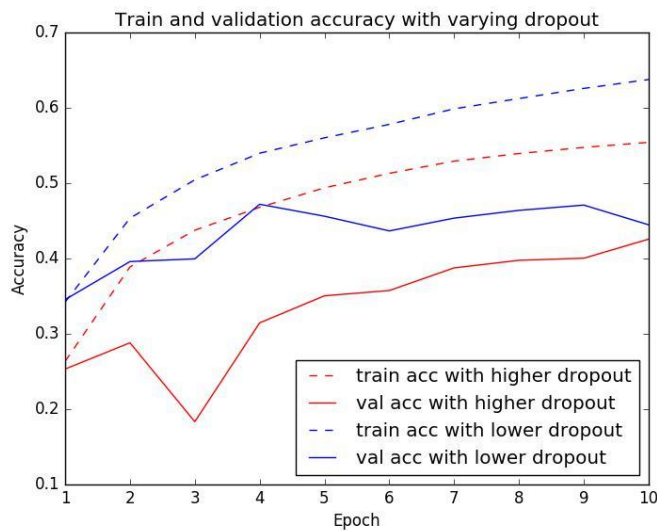


Figure 5: Higher dropout results in less over fitting but a longer time to reach accuracies achieved using the same parameters and lower dropout.




Figure 6. Visualizations of filters from VGG16. The visualizations represent filters from the first (top left), third (top right), fifth (bottom left), and seventh (bottom right) convolutional layers. These visualizations were generated by performing gradient ascent to generate an image that maximizes the activations at various layers in the network.


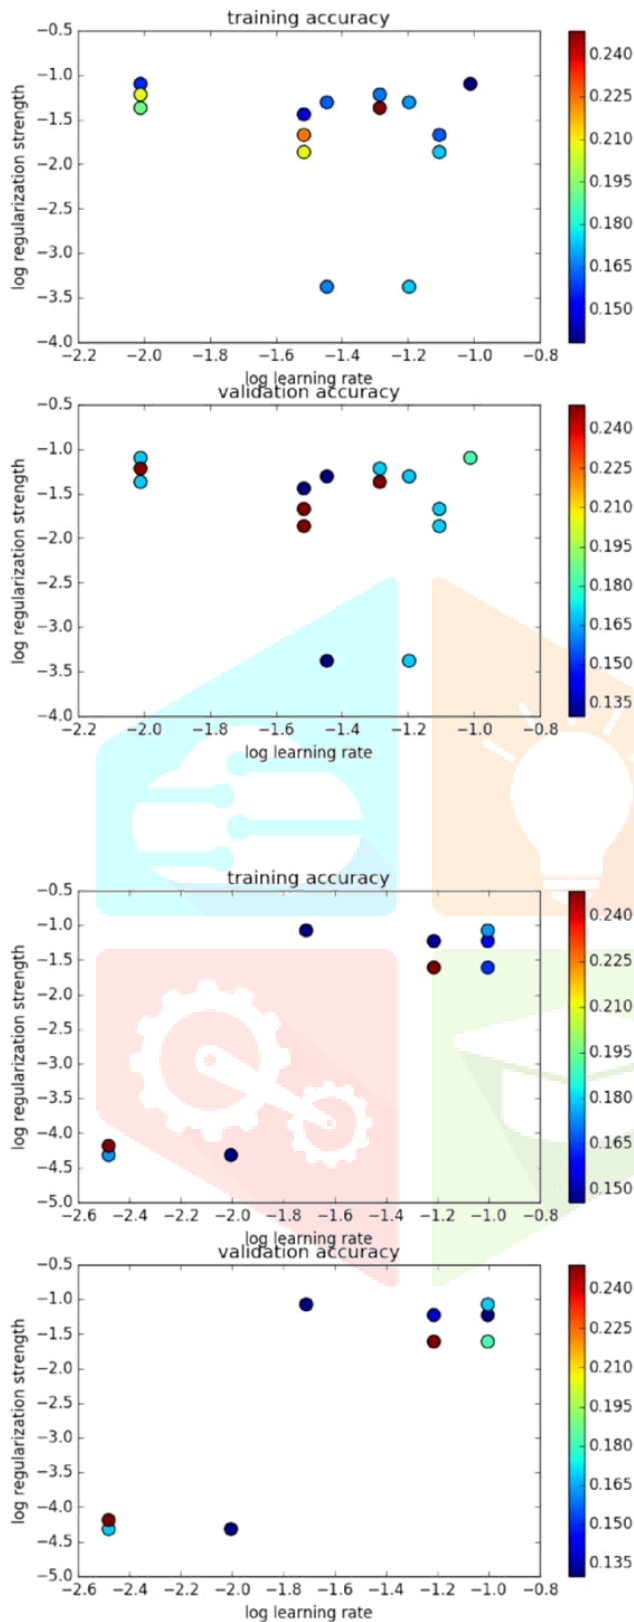



Figure 8. Visualizations of filters from the deeper CNN model. The visualizations represent filters from the first (top left), third (top right), fifth (bottom left), and seventh (bottom right) convolutional layers. These visualizations were generated by performing gradient ascent to generate an image that maximizes the activations at various layers in the network



Normalized confusion matrix of our mini-Xception network

Figure 7. The upper two graphs represent the training and validation accuracy for the the five-layer CNN during learning rate and regularization tuning. The accuracies are shown by the colour scale depicted to the right. The lower two graphs represent the five-layer CNN with fractional max-pooling.

Overall, however, training a model from scratch out-performed fine-tuning. This trend differs from what we have seen with other applications, but is understandable given that there has been relatively less work on face data sets compared to general image classification data sets, and given that state-of-the-art networks are less polished in this more specific task of FER. In addition, our application (FER) is not a sub-problem of the problem the network was trained to solve (e.g. classifying dog breeds vs. classifying animals); rather, FER is tangentially related to facial recognition, the task for which VGGFace was trained.




Fig. 9: Results of the provided real-time emotion classification provided in our public repository




Fig. 9: Results of the provided combined gender and emotion inferences demo. The colour blue represents the assigned class woman and red the class man


V. FUTURE WORK

Machine learning models are biased in accordance to their training data. In our specific application we have empirically found that our trained CNNs for gender classification are biased towards western facial features and facial accessories. We hypothesize that this misclassifications occurs since our training dataset consist of mostly western: actors, writers and cinematographers as observed in Figure 2.

Furthermore, as discussed previously, the use of glasses might affect the emotion classification by interfering with the features learned. However, the use of glasses can also interfere with the gender classification. This might be a result from the training data having most of the images of persons wearing glasses assigned with the label "man". We believe that uncovering such behaviours is of extreme importance when creating robust classifiers, and that the use of the visualization techniques such as guided back-propagation will become invaluable when uncovering model biases.



(a) Samples from the FER-2013 dataset



(b) Guided back-propagation visualization of our mini-Xception model

VI. CONCLUSION

We have proposed and tested general building designs for creating real-time CNNs. Our proposed architectures have been systematically built in order to reduce the amount of parameters. We began by eliminating completely the fully connected layers and by reducing the amount of parameters in the remaining convolutional layers via depth-wise separable convolutions. We have shown that our proposed models can be stacked for multi-class classifications while maintaining real-time inferences. Specifically, we have developed a vision system that performs face detection, gender classification and emotion classification in a single integrated module. We have achieved human-level performance in our classifications a task using a single CNN that leverages modern architecture constructs. Our architecture reduces the amount of parameters 80 while obtaining favourable results. Our complete pipeline has been successfully integrated in a Care-O-bot 3 robot. Finally we presented a visualization of the learned features in the CNN using the guided back-propagation visualization. This visualization technique is able to show us the high-level features learned by our models and discuss their interpretability.

REFERENCES

- [1]. SERGEY IOFFE AND CHRISTIAN SZEGEDY. "BATCH NORMALIZATION: ACCELERATING DEEP NETWORK TRAINING BY REDUCING INTERNAL COVARIATE SHIFT". IN INTERNATIONAL CONFERENCE ON MACHINE LEARNING, PAGES 448–456, 2015.
- [2]. S. G. Kong, j. Heo, f. Boughorbel, y. Zheng, b.r. abidi, a. Koschan, m. Yi, and m. A. Abidi: "multiscale fusion of visible and thermal ir images for illumination invariant face recognition", *ijcv* (2007)
- [3]. N. Kar, m. Kanti debbarma, a. Saha, and d. Rudra pa: "study of implementing automated attendance system using face recognition technique", volume 1, no. 2, *international journal of computer and communication engineering* (2012)
- [4]. S. Rani k, t.c. sharma : "face recognition office security system using lab view 8.6", *ijecierd*, volume 3, pp: 195-200 (2013)
- [5]. Buddharaju, p. Et al.: "physiology-based fr in the thermal irs", *ieee transactions on pami*, volume 29, no. 4, pp: 613-626 (2007).
- [6]. S. Kong, j. Heo, b. Abidi, j. Paik, and m. Abidi: "recent advances in visual and infrared face recognition: a review", *computer vision and image understanding*, volume 97, pp: 103-135(2005)
- [7]. R. Gross, i. Matthews, and s. Baker: "active appearance models with occlusion", *image and vision computing*, pp: 593-604(2006)
- [8]. L. Wolf and a. Shashua: "learning over sets using kernel principal angles", *jmlr*, volume 4, no.10, pp: 913-931(2003).
- [9]. S. Q. Wu, l.z. wei, z.j. fang, r.w. li, and x. Q. Ye: "infrared face recognition based on blood perfusion and sub-block dct in wavelet domain", *conference on wavelet analysis and pattern recognition*, volume 3, pp: 1252-1256(2007).
- [10]. D. A. Socolinsky and a. Selinger: "thermal face recognition over time, pattern recognition", volume 4, pp: 187-190(2004)
- [11]. XAVIER GLOROT, ANTOINE BORDES, AND YOSHUA BENGIO. "DEEP SPARSE RECTIFIER NEURAL NETWORKS". IN PROCEEDINGS OF THE FOURTEENTH INTERNATIONAL CONFERENCE ON ARTIFICIAL INTELLIGENCE AND STATISTICS, PAGES 315–323, 2011
- [12]. CHRISTIAN SZEGEDY, VINCENT VANHOUCKE, SERGEY IOFFE, JON SHLENS, AND ZBIGNIEW WOJNA. "RETHINKING THE INCEPTION ARCHITECTURE FOR
- [13]. COMPUTER VISION." IN PROCEEDINGS OF THE IEEE CONFERENCE ON COMPUTER VISION AND PATTERN RECOGNITION, PAGES 2818–2826, 2016.
- [14]. RASMUS ROTHE, RADU TIMOFTE, AND LUC VAN GOOL. "DEEP EXPECTATION OF REAL AND APPARENT AGE FROM A SINGLE IMAGE WITHOUT FACIAL LANDMARKS". *INTERNATIONAL JOURNAL OF COMPUTER VISION (IJCV)*, JULY 2016.
- [15]. RUSSAKOVSKY, J. DENG AND L. FEI-FEI. "IMAGENET LARGE SCALE VISUAL RECOGNITION CHALLENGE". *INTERNATIONAL JOURNAL OF COMPUTER VISION IJCV*, 2015
- [16]. C.W. PABLO BARROS AND S. WERMTER. "EMOTIONAL EXPRESSION RECOGNITION WITH A CROSS-CHANNEL CONVOLUTIONAL NEURAL NET-WORK FOR HUMAN-ROBOT" INTERACTION. *IEEE 15TH INTERNATIONAL CONFERENCE ON HUMANOID ROBOTS*, 2015
- [17]. KAIMING HE, XIANGYU ZHANG, SHAOQING REN, AND JIAN SUN. "DEEP RESIDUAL LEARNING FOR IMAGE RECOGNITION". IN PROCEEDINGS OF THE IEEE CONFERENCE ON COMPUTER VISION AND PATTERN RECOGNITION, PAGES 770–778, 2016.
- [18]. T. Elguebaly and nizar bouguila: "a bayesian method for infrared faces recognition, machine vision beyond visible spectrum" (2011).

- [19]. Y. Yoshitomi: "face identification using thermal image processing, workshop on robot & human communication", pp: 374-379(1997)
- [20]. R. Siddiqui: "face identification based on biological trait using infrared images after cold effect enhancement and sunglasses filtering", in proceedings of international conference in central european computer graphics, visualization and computer vision (2004).
- [21]. X. Chen, k. Bowyer: "visible light and infrared fr", procd. In workshop on multimodal user authentication, pp: 48-55(2003)
- [22]. Oliver rockinger and thomas fechner: "pixel-level image fusion:" the case of image sequences, spie (1998).
- [23]. IAN GOOD FELLOW ET AL. CHALLENGES IN "REPRESENTATION LEARNING": A REPORT ON THREE MACHINE LEARNING CONTESTS, 2013.
- [24]. ANDREW G. HOWARD ET AL. MOBILENETS: "EFFICIENT CONVOLUTIONAL NEURAL NETWORKS FOR MOBILE VISION APPLICATIONS". CORR, ABS/1704.04861, 2017.
- [25]. DARIO AMODEI ET AL. "DEEP SPEECH 2: END-TO-END SPEECH RECOGNITION IN ENGLISH". CORR, ABS/1512.02595, 2015

