© May 2019 | IJIRT | Volume 5 Issue 12 | ISSN: 2349-6002

Convolutional Neural Network in Medical Diagnosis

Vivek Paul[1,] Anshul Gupta[2]
[1,2]*Dept. of Information Technology, Maharaja Agrasen Institute of Technology, New Delhi, India*

Abstract- **Convolutional Neural Network (CNN) is transforming the field of medical diagnosis. CNN can help doctors make faster, more accurate diagnosis by providing automatic learning techniques for predicting the common patterns from the medical image data. Human expert provides limited interpretation of medical images due to its subjectivity, complexity and extensive variations across the image. CNN is able to provide state of the art solutions with good accuracy for medical imaging and is powered by the increasing availability of healthcare data. Major disease areas that use CNN includes cancer, dermatology, neurology and cardiology. This paper focuses on the use of different CNN architectures based on their performance for accurate medical diagnosis. We also discuss the current status of CNN applications in healthcare and its various limitations.**

Index Terms- Convolutional Neural Network, AlexNet, ZFNet, VGGNet, GoogleNet, ResNet and Medical Diagnosis

I. INTRODUCTION

Artificial Intelligence (AI) has sent waves across healthcare industry. It has the potential to revolutionize disease diagnosis and management by performing classification tasks which are difficult even for human experts and by rapidly analyzing large amounts of images. As health is a priority, medical experts are continually trying to find ways to implement new technologies and provide impactful results. Convolutional Neural Network (CNN) is making healthcare smarter. This powerful subset of artificial intelligence is most commonly applied to analyzing visual imagery. CNN has demonstrated its truly life impacting potential in healthcare through its various path breaking applications particularly in the field of medical diagnosis.

The classical algorithmic approach to image analysis for classification relied on (1) manual feature extraction (2) handcrafted object segmentation, followed by (3) identification of each segmented object using statistical classifiers or shallow neural computational machine-learning classifiers designed specifically for each class of objects, and finally (4) classification of the image [1]. Creating and refining multiple classifiers required many skilled people and was computationally expensive [2] [3].

In the recent years CNNs have garnered huge success in the field of image analysis. Specifically, the annual vision challenges from MS COCO [4] and ILSVRC [5] made it evident that humans are surpassed by contemporary CNN architectures in many vision related tasks. The ability to perform automatic feature engineering on a large image dataset has made the CNN very successful. Many studies [6] [7] [8] have reported that highly tuned non-CNN counterparts fare poorly when compared with the features extracted by contemporary CNNs which yield superior results in many tasks consistently.

The development of CNN layers has allowed for significant gains in the ability to classify images and detect objects in a picture. Image analysis filters, or convolutions, are applied on these multi-processing layers. By systematically convolving multiple filters across the image an abstracted representation of image is constructed within each layer, thus producing a feature map which acts as an input to the following layer. CNN makes it possible to process images in the form of pixels as input and to give the desired classification as output. CNNs can predict if a pixel belongs to the object of interest by using a patch or sub image centered on that pixel. This image to classification approach using a single classifier replaces the multiple steps required by previous image analysis methods. The multi-stream architecture of CNN can accommodate multiple sources of information or representations of the input in the form of channels presented to the input layer. CNNs can also be adapted to leverage intrinsic structure of medical images.

The profound and lasting impact made by CNN on healthcare is driven by four powerful trends: (1) Development of powerful and energy efficient

Graphics Processing Units (GPUs). Until recently, running AI algorithms was rarely cost-effective. (2) Besides the hardware development, the wide availability of open source packages and the GPU computing libraries such as OpenCL, CUDA has fueled the popularity of CNNs in medical imaging. (3) Development of sophisticated algorithms. We can now use CNN models at a fraction of past costs. (4) Availability of huge volumes of Healthcare data. Due to the digitization efforts, more and more healthcare data is available to train algorithms.

CNNs can assist physicians to make better clinical decisions or even replace human judgement in certain functional areas of healthcare in the future. The increasing availability of healthcare data and rapid development of image analysis techniques has made possible the recent successful applications of CNN in healthcare. Guided by relevant clinical questions, powerful CNN models can unlock clinically relevant information hidden in the massive amount of data, which in turn can assist clinical decision making.

The rest of this paper is organized as follows: section 2 covers the insights on the working of CNN. In section 3, we explain the various CNN architectures. Section 4 presents the applications of CNN in Medical Diagnosis. In section 5 the limitations are discussed and finally in section 6 we provide the conclusion.

II. CONVOLUTIONAL NEURAL NETWORK

Convolutional neural networks (CNN) are deep artificial neural networks that are used primarily for image classification, image clustering, and object detection within scenes. CNNs are inspired by biological processes [9, 10] in that the architecture of a CNN is analogous to that of the connectivity pattern between neurons in the Human Brain and resembles the organization of the animal visual cortex. Individual cortical neurons respond to stimuli only in a restricted region of the visual field known as the receptive field. The receptive fields of different neurons partially overlap to cover the entire visual field.

Compared to other image classification algorithms CNN requires relatively lower level of pre-processing. This means that with enough training the network can learn the filters that in traditional algorithms were hand-engineered. This independence from prior knowledge and human effort in feature design is a major advantage. CNN is able to successfully capture the Spatial and Temporal dependencies in complex images having pixel dependencies throughout using the application of relevant filters. The architecture performs a better fitting to the image dataset due to the reduction in the number of parameters involved and reusability of weights. In other words, the network can be trained to understand the sophistication of the image better.

As shown in Figure 1 a CNN arranges its neurons in three dimensions (width, height, depth), as visualized in one of the layers. Every layer of a CNN transforms the 3D input volume to a 3D output volume of neuron activations. In this example, the red input layer holds the image, so its width and height would be the dimensions of the image, and the depth would be 3 (Red, Green, Blue channels).

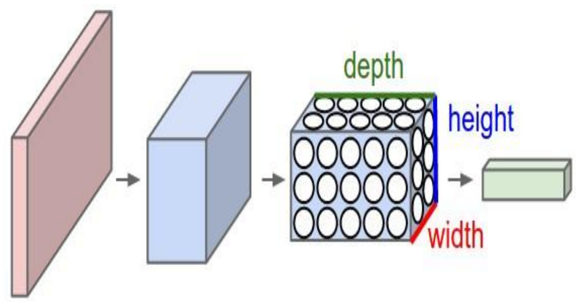

Figure 1: 3D arrangement of neurons by CNN

There are five different layers in a typical basic CNN, namely input layer, convolutional layers, rectified linear unit (ReLU) layers, pooling layers and fully-connected layers. Typical layered architecture of CNN is shown in Figure 2. Input layer holds the raw pixel values of the image. Convolutional layers are sets of filters that are learnable. When its input is scanned, a 2-dimensional activation map is generated, and it is used to describe which filter is active at each spatial position of the input. ReLU layer applies an elementwise activation function, such as the max (0, x) thresholding at zero. This leaves the size of the volume, i.e. height, width and depth, unchanged. Pooling layers are used to reduce the cost of computations and the venture of over-fitting by decreasing the dimension of the data representation. Fully connected layers can translate data to one dimensional data structure that is non-spatially dependent, and achieved by using the previous layer's nodes.

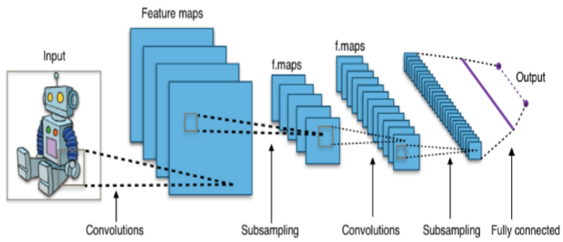

Figure 2: Typical layers of CNN

III. CNN ARCHITECTURES

The ImageNet project is a large visual database designed for use in visual object recognition software research. The ImageNet project runs an annual software contest, the ImageNet Large Scale Visual Recognition Challenge (ILSVRC), where software programs compete to correctly classify and detect objects and scenes.

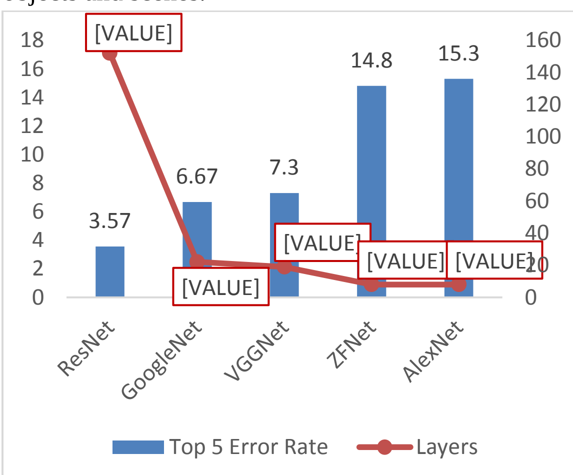

Figure 3: Top-5 Error Rate of various CNN architectures

ImageNet, is a dataset of over 15 million labeled high-resolution images with around 22,000 categories. ILSVRC uses a subset of ImageNet of around 1000 images in each of 1000 non-overlapping categories. In all, there are roughly 1.3 million training images, 50,000 validation images and 100,000 testing images. Top-5 error rate of various ILSVRC top competitors are given in Figure 3. There are various architectures of CNNs available which have been key in building algorithms which power and shall power AI as a whole in the foreseeable future. Some of them have been listed below:

A. AlexNet

AlexNet is a CNN developed by Alex Krizhevsky, Ilya Sutskever and Geoff Hinton [11]. The AlexNet won the ImageNet ILSVRC challenge in 2012 with a top 5 error rate of 15.3%. The Network has a very similar architecture to LeNet, but has deeper, with more filters per layer, bigger, and featured Convolutional Layers stacked on top of each other. It consists of 11x11, 5x5, 3x3, convolutions, max pooling, dropout, data augmentation, ReLU activations, SGD with momentum. It attached ReLU activations after every convolutional and fully-connected layer. The Architecture of the network is given in Figure 4. AlexNet popularized the use of CNNs in Computer Vision.

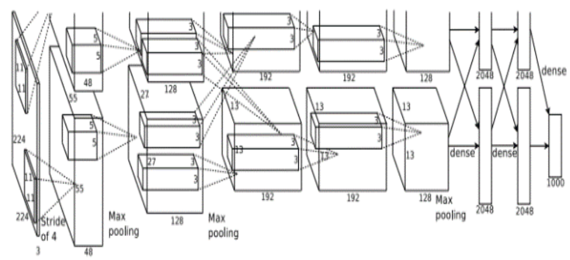

Figure 4: AlexNet Architecture

It uses ReLU activations instead of Tanh to add non-linearity. This accelerated the speed by 6 times at the same accuracy. To deal with overfitting it uses dropout instead of regularization. However this resulted in twice the training time required with the dropout rate of 0.5. To reduce the size of network AlexNet uses overlap pooling which also reduces the top-1 and top-5 error rates by 0.4% and 0.3%, respectively.

B. ZFNet

ZFNet is a CNN developed by Matthew Zeiler and Rob Fergus [12]. It won the ILSVRC 2013 challenge. The name of the CNN is based on the initials of their surnames Zeiler and Fergus. It is an improvement on AlexNet by tweaking the architecture hyperparameters, in particular by expanding the size of the middle convolutional layers and making the stride and filter size on the first layer smaller. It achieved a top-5 error rate of 14.8%. The architecture of ZFNet is shown in Figure 5. ZFNet offered essential insights into how CNNs are learning internal representations. To achieve this goal, the developers introduced a way to map learned features into input pixel space by using a specially designed Deconvolutional Network.

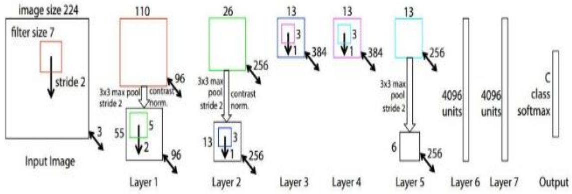

Figure 5: ZFNet Architecture

C. VGGNet

VGGNet is a CNN developed by Karen Simonyan and Andrew Zisserman [13]. It was the runner-up in ILSVRC 2014 challenge. It achieved a top-5 error rate of 7.3%. The final VGGNet consists of 16 Convolutional/Fully Connected layers and its uniform architecture makes it very appealing. Its main contribution was in showing that the depth of the network is a critical component for good performance. The homogeneous architecture performs 3x3 convolutions and 2x2 pooling from the beginning to the end which is similar to AlexNet but it uses lots of filters. The architecture of VGGNet is shown in Figure 6. The pretrained model is available for plug and play use in Caffe. The weight configuration of the VGGNet is publicly available and has been used in many other applications and challenges as a baseline feature extractor. A downside of the VGGNet is that it consists of 138 million parameters which uses a lot of memory and is therefore more expensive to evaluate. However most of these parameters are in the first fully connected layer, and can be removed with no performance downgrade, significantly reducing the number of necessary parameters.

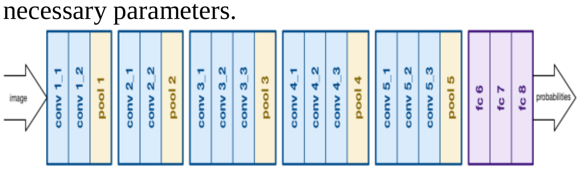

Figure 6: VGGNet Architetcure

D. GoogleNet

GoogleNet is a CNN developed by Szegedy et al. from Google [14]. It is the winner of the ILSVRC 2014 competition. It achieved a top-5 error rate of 6.67%. This is very close to human level performance. The network uses a CNN inspired by LeNet but implements a novel element which is dubbed an inception module. The idea of the inception layer is to cover a bigger area, but also keep a fine resolution for small information on the images. So the idea is to convolve in parallel different sizes from the most accurate detailing (1x1) to a bigger one (5x5). It uses batch normalization, image distortions and RMSprop. This module is based on several very small convolutions in order to drastically reduce the number of parameters. GoogleNet architecture consists of a 22 layer deep CNN but reduces the number of parameters from 60 million (AlexNet) to 4 million. Additionally, GoogleNet uses Global Average Pooling at the end of the CNN instead of Fully Connected layers, eliminating a large amount of parameters that do not seem to matter much. The architecture of GoogleNet is shown in Figure 7. There are also several follow-up versions to the GoogleNet, most recently Inception-v4.

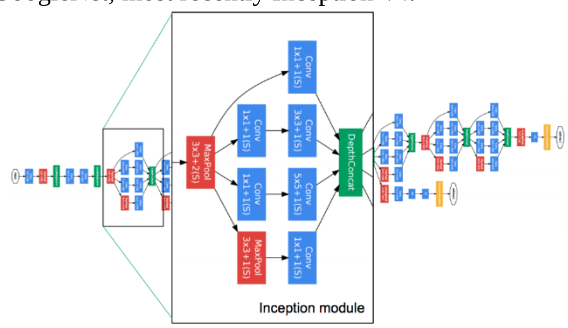

Figure 7: GoogleNet Architecture

E. ResNet

Residual Network is a CNN developed by Kaiming He et al. [15]. It is the winner of ILSVRC 2015 challenge. It utilizes special skip connections, or short-cuts to jump over some layers and features heavy batch normalization. The skip connections used are also known as gated recurrent units and share similarity with the elements recently applied in RNNs. The novel architecture is also missing fully connected layers at the end of the network. The architecture is shown in Figure 8. ResNets are currently by far state of the art Convolutional Neural Network models and are the default choice for using CNN in practice. ResNet can consist of 152 layers while still having lower complexity than VGGNet. It achieves a top-5 error rate of 3.57% which beats human-level performance.

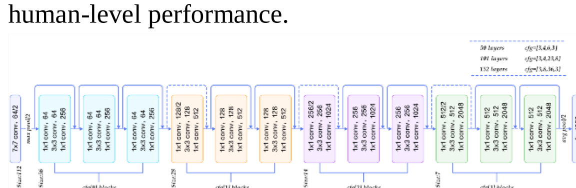

Figure 8: ResNet Architecture

The skip-connections used in ResNet help to address the Vanishing Gradient problem often experienced by very deep neural networks. They also make it easy for a ResNet block to learn an identity function. There are two main types of ResNets blocks: The identity block and the convolutional block. Very deep Residual Networks are built by stacking these blocks together.

IV. APPLICATION

Below we discuss some state of the art applications of CNN in medical diagnosis. Though, the list is by no means complete however it provides an indication of the long-ranging impact of CNN in the medical industry today.

A. Skin Cancer Diagnosis

Skin cancer diagnosis is a challenging task. To diagnose a skin disease, variety of visual clues are used such as the individual lesional morphology, the body site distribution, color, scaling and arrangement of lesions. Melanoma, has four major clinical diagnosis methods: ABCD rules, pattern analysis, Menzies method and 7-Point Checklist. To achieve a satisfactory diagnostic accuracy using these methods, a high level of expertise is required due to the variety of visual aspects of the skin lesions. Unlike the diagnosis by human experts, which depends essentially on subjective judgment and is not always reproducible, a computer aided diagnostic system is more objective and reliable.

Esteva et al. [16] used deep CNN which was trained end-to-end by taking the images directly, using only pixels and disease labels as inputs. The dataset used to train the network consisted of 129,450 clinical images of 2,032 different diseases. The network learned to distinguish keratinocyte carcinomas versus benign seborrheic keratoses, and malignant melanoma versus benign nevus. The researchers tested its performance against 21 board-certified dermatologists on biopsy-proven clinical images. Across both the tasks the CNN achieved performance on par with all certified experts. The network achieved a three way accuracy of 72.1% whereas the dermatologist achieved an accuracy of66%. It demonstrated that a CNN is capable of classifying skin cancer with a level of competence comparable to dermatologists.

B. Diabetic Retinopathy Diagnosis

Patients with severe diabetes are at risk of developing diabetic retinopathy (DR). Retinopathy leads to visual loss and blindness with the passage of time if not detected at an early stage. Regular eye examinations are the cornerstone of DR prevention and early detection. Manual detection of DR is difficult and time consuming process due to unavailability of equipment and the high level of expertise. Most of the current diagnostic programs employ 1- or 2-field retinal color fundus imaging due to its cost-effectiveness. With this in mind, researchers proposed automated DR detection [17] using a deep CNN architecture for diagnosing retinopathy and classifying its severity based on fundus images. They employed dataset consisting of about 80,000 digital fundus images from patients of various ethnicity and age, and then applied data augmentation on it. They also validated the network on 5,000 images. The network was able to identify micro-aneurysms, exudate and hemorrhages on the retina and automatically classified the images into 5 classes namely No retinopathy, Mild DR, Moderate DR, Severe DR and Proliferative DR. The final trained CNN achieved up to 95% specificity, 30% sensitivity and 75%. Researchers claim that a higher specificity is achieved at the expense of sensitivity.

C. Alzheimer's Disease Diagnosis

Alzheimer's disease (AD) is a chronic neurodegenerative disease that usually starts slowly and gradually worsens over time. It is the cause of 60–70% of cases of dementia. The most common early symptom is difficulty in remembering recent events. As the disease advances, symptoms can include problems with language, disorientation, mood swings, loss of motivation, not managing self-care, and behavioral issues. Alzheimer's disease is mainly diagnosed by studying the individual's behavior and medical history. Magnetic Resonance Imaging (MRI) is also used to analyze the brain's morphometric patterns for identifying disease-specific imaging biomarkers.

In this paper, we build a 3D Convolutional Neural Network (3D-CNN) and provide a simple method to interpret different regions of the brain and their association with the disease to identify AD biomarkers. Our method uses minimal preprocessing of MRIs (imposing minimum preprocessing artifacts) and utilizes a simple data augmentation strategy of down sampled MR images for training purposes. Unlike the vast majority of previous works, the proposed framework, thus, uses a voxel-based 3D-CNN to account for all voxels in the brain and capture the subtle local brain details in addition to better pronounced global specifics of MRIs.

Esmaeilzadeh et al. [18] proposed a simple 3D CNN and modified its parameters to tailor the end-to-end architecture for the diagnosis of Alzheimer's disease and its prodromal stage, MCI, using MRI images. The model achieved an accuracy of 94.1% on the popular ADNI dataset using only MRI data, which outperforms the previous state-of-the-art. The trained model was also able to identify the disease biomarkers. The researchers identified that the hippocampus region of the brain is critical in the diagnosis of AD. After extensive hyperparameter tuning the model was further trained using transfer learning to diagnose mild cognitive impairment (MCI), the prodromal stage of AD, which yield better results compared to other methods.

D. Tuberculosis Diagnosis

Tuberculosis (TB) is an infectious disease usually caused by Mycobacterium tuberculosis (MTB) bacteria. Tuberculosis generally affects the lungs, but can also affect other parts of the body. The classic symptoms of active TB includes a chronic cough with blood-containing sputum, fever, night sweats, and weight loss. Timely diagnosis and treatment is key to full patient recovery. The Microscopic Observed Drug Susceptibility (MODS) is a test to diagnose TB infection and drug susceptibility directly from a sputum sample in 7–10 days with a low cost and high sensitivity and specificity, based on the visual recognition of specific growth patterns of MTB in a broth culture. Despite that, MODS find limited usage in remote, low resource settings, because it requires permanent and trained technical staff for the image based diagnostics. Hence, an automated computer based alternative technique is required to analyze the MODS cultures.

Lopez et al. [19] designed an automated TB diagnosis system based on CNN. The system is trained to automatically evaluate and interpret MODS cultures digital images. The CNN was trained using a dataset of 12,510 images of 7-10 days MODS samples. In the dataset there were 4,849 positive and 7,661 negative images. Images were rescaled to 224 x 224 pixels, and converted to grayscale. The images were obtained from three different laboratories. The features extracted automatically by the CNN resembles visual cues used by expert diagnosticians to interpret MODS cultures, which suggests that the model has the ability to generalize and scale. The model 96.63 +/- 0.35% accuracy, 94.74 +/- 0.89% sensitivity and 97.83 +/- 1.07% specificity, when validated across held-out laboratory datasets. This CNN can assist laboratory personnel, in low resource settings, especially in developing countries.

V. LIMITATION

While CNNs exhibit high performance in image classification tasks, their capabilities aren't devoid of issues. For a problem as diverse and complex as medical imaging, CNN requires large datasets in order to reach the required levels of accuracy. In combination with long training durations, this presents a major drawback for any small to medium-scale institution or enterprise. Most of the medical images have poor signal-to-noise ratio when compared with images taken with a digital camera. This can hinder the successful computation as the contrast between anatomically distinct structures (e.g. lesion, edema, and healthy tissue) will be too low. Therefore adequate pre-processing is required to remove artifacts and reduce noise from the images. Next, there is always a possibility to over-train a CNN if the network is trained on the same dataset for too long. This results in the over-fitting of training images. To prevent this validation dataset is used to track a CNN's performance and stop training when it starts decreasing. CNNs work in a black-box way, not really offering explanations of their conclusions. With any input, a corresponding output can be produced, but it may be hard to explicate the decision and its reliability, which is certainly problematic for medical image processing.

VI. CONCLUSION

CNN methods have a wide application in the medical field for addressing problems such as diagnosing infectious diseases, different forms of cancer, ophthalmologic complications of diabetes and chronic neurological disorders. CNNs can perform a wide range of image processing operations along with increasing the speed of the procedure as compared to other classical techniques. In this paper, we highlighted the various state of the art CNN architectures as well as their applications in medical image analysis. Though, the list is by no means complete however it provides an indication of the long-ranging impact of CNN in the medical imaging industry today. The future development of CNN promises to diversify its uses in the field of medicine, particularly in the domain of medical diagnosis. However, we should not consider it as the only solution as there are several limitations that reduce its growth. It is evident CNN can't substitute the role of doctors/clinicians in medical diagnosis. So far CNN based applications provided positive feedback and we should continue to look into more sophisticated CNN methods that can deal with complex healthcare data efficiently. Lastly we conclude that CNNs will prevail in the near future, and that they will find many other uses in the healthcare industry.

REFERENCES

[1] Goldbaum, M., Moezzi, S., Taylor, A., Chatterjee, S., Boyd, J., Hunter, E., and Jain, R. (1996). Automated diagnosis and image understanding with object extraction, object classification, and inferencing in retinal images. Proceedings of 3rd IEEE International Conference on Image Processing 3, 695–698.

[2] Chaudhuri, S., Chatterjee, S., Katz, N., Nelson, M., and Goldbaum, M. (1989). Detection of blood vessels in retinal images using two-dimensional matched filters. IEEE Trans. Med. Imaging 8, 263–269.

[3] Hoover, A., and Goldbaum, M. (2003). Locating the optic nerve in a retinal image using the fuzzy convergence of the blood vessels. IEEE Trans. Med. Imaging 22, 951–958

[4] Lin, T.; Maire, M.; Belongie, S. J.; Hays, J.; Perona, P.; Ramanan, D.; Doll´ar, P.; and Zitnick, C. L. 2014. Microsoft COCO: common objects in context. In Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V, 740–755.

[5] Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M. S.; Berg, A. C.; and Li, F. 2015. Imagenet large scale visual recognition challenge. International Journal of Computer Vision 115(3):211–252.

[6] Razavian, A. S.; Azizpour, H.; Sullivan, J.; and Carlsson, S. 2014. CNN features off-the-shelf: An astounding baseline for recognition. In IEEE Conference on Computer Vision and Pattern Recognition, CVPR Workshops 2014, Columbus, OH, USA, June 23-28, 2014, 512–519.

[7] Donahue, J.; Jia, Y.; Vinyals, O.; Hoffman, J.; Zhang, N.; Tzeng, E.; and Darrell, T. 2014. Decaf: A deep convolutional activation feature for generic visual recognition. In Proceedings of the 31th International Conference on Machine Learning, ICML 2014, Beijing, China, 21-26 June 2014, 647–655.

[8] Zeiler, M. D., and Fergus, R. 2014. Visualizing and understanding convolutional networks. In Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part I, 818–833.

[9] Kunihiko Fukushima. Neocognitron: A Self-organizing Neural Network Model for a Mechanism of Pattern Recognition Unaffected by Shift in Position. Biological Cybernetics. 36 (4): 193–202. DOI:10.1007/BF00344251,

[10] D. H. Hubel, T. N. Wiesel. Receptive fields and functional architecture of monkey striate cortex. The Journal of Physiology. 195 (1): 215–243. DOI:10.1113/jphysiol.1968.sp008455.

[11] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. 2017. ImageNet classification with deep convolutional neural networks. Commun. ACM 60, 6 (May 2017), 84-90. DOI: https://doi.org/10.1145/3065386

[12] Zeiler M.D., Fergus R. (2014) Visualizing and Understanding Convolutional Networks. In: Fleet D., Pajdla T., Schiele B., Tuytelaars T. (eds) Computer Vision – ECCV 2014. ECCV 2014. Lecture Notes in Computer Science, vol 8689. Springer, Cham

[13] Karen Simonyan and Andrew Zisserman. Very Deep Convolutional Networks for Large-Scale Image Recognition. *CoRR* abs/1409.1556 (2015): n. pag.

[14] C. Szegedy et al., "Going deeper with convolutions," 2015 IEEE Conference on

Computer Vision and Pattern Recognition *(CVPR)*, Boston, MA, 2015, pp. 1-9. doi: 10.1109/CVPR.2015.7298594

[15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, Jian Sun. Deep Residual Learning for Image Recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016): 770-778.

[16] A. Esteva, B. Kuprel, R. A. Novoa, J. Ko, S. M. Swetter, H. M. Blau, et al., Dermatologist-level classification of skin cancer with deep neural networks, Nature, vol. 542, pp. 115-118, Feb 02 2017.

[17] Gulshan V, Peng L, Coram M, et al. Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs. JAMA. 2016; 316(22):2402–2410. doi:10.1001/jama.2016.17216

[18] Esmaeilzadeh, S., Belivanis, D.I., Pohl, K.M., & Adeli, E. (2018). End-To-End Alzheimer's disease Diagnosis and Biomarker Identification. *MLMI@MICCAI*.

[19] Lopez-Garnier S, Sheen P, Zimic M (2019) Automatic diagnostics of tuberculosis using convolutional neural networks analysis of MODS digital images. PLoS ONE 14(2): e0212094. https://doi.org/10.1371/journal.pone.0212094