*i*Journals: International Journal of Software & Hardware Research in Engineering (IJSHRE)
ISSN-2347-4890
Volume 10 Issue 11 November 2022

Developing a Machine-Learning Algorithm to Diagnose Age-Related Macular Degeneration

Authors: Shikhar Gupta
Monta Vista High School
DOI: 10.26821/IJSHRE.10.10.2022.101012

ABSTRACT

Age related macular degeneration (AMD) is an ocular disease that affects the central retina. The disease most commonly leads to blindness, and it involves a gradual deterioration of the retina cells or a fluid buildup in the retina that progressively hinders the sight. Ocular diseases including age-related macular degeneration are found in more than 196 million people aged 40 and up world wide, and these diseases are expected to grow to 240 million by 2050. The cost of the special machinery used to detect AMD can be a heavy burden for small eye clinics that are sporadically found in the regions outside of a city. As a result, we can find these advanced types of diagnostic machinery in inner city large hospitals and large eye clinics. Access to these types of machinery is a challenge that with affordable and efficient options will drastically lower the inequities regarding access to eye treatment. As a result of the issues regarding the diagnostic processes of age-related macular degeneration, a type of technology known as a convolutional neural network (CNN), a deep learning algorithm designed specially for images and pixel processing, can be especially used effectively in this scenario. A CNN completes this task in a fraction of the time it would take with conventional methods, and is functional in use for other contexts as well, especially in areas where a comprehensive exam may not be feasible. Strategies that are needed to train CNN models for the assessment and diagnosis of ocular diseases have not yet been well characterized. By these ends, we trained a CNN model consisting of Convolution, Max Pooling, and ReLU layers on 5000 images of retinas affected by age-related macular degeneration, to conclude a diagnostic F1 score of 89%.

SUMMARY

The human eye, scientifically termed *oculus*, is essential to our functioning – without it, we would not be able to observe the world around us. Ocular diseases however, are frequent among elderly populations (elderly being defined as 40+). Today, more than 12 million people over the age of 40 suffer from ocular diseases. Most commonly, older patients are susceptible to age related macular degeneration, an eye disease that causes blurring of the central vision due to the deterioration of the macula which is the central part of the retina. The former can only be detected through complex and expensive imaging software, markedly a visual field test; this leaves a significant population with untreated eye disease and holds them at risk for complete vision loss. The use of machine learning algorithms has been proposed for treating eye disease. However, the development of these models is limited by a lack of understanding regarding appropriate model and training parameters to maximize model performance. In our study, we address these points by generating multiple machine learning models, each with a learning rate with our highest F1 score being 0.89.
Our analysis shows that sample imbalance is a key challenge in training of machine learning models and can result in deceptive improvements in training cost which does not translate to true improvements in model predictive performance.

Considering the wide ranging impact of ocular eye disease in elderly populations and its repercussions on the audience, we developed a machine learning algorithm to treat the same. We trained our model on varying ocular disease datasets consisting of over 5,000 patients, and the laser ophthalmology of their infected eyes. In the future, we hope this model is used extensively, especially in areas of the world that are under-resourced, to better diagnose eye disease, improve overall efficiency and increase access to global healthcare.

INTRODUCTION

Age related macular degeneration is an ocular disease that affects the central retina.. The disease most commonly leads to blindness, and it involves a gradual deterioration of the retina cells or a fluid buildup in the retina that progressively hinders the sight (1). Ocular diseases including age-related macular degeneration are found in more than 196 million people aged 40 and up world wide, and these diseases are expected to grow to 240 million by 2050 (2). The above is characterized by abnormalities in vision and retinal pigment. The disease is divided into three distinctions based on the age of the individual because of the prospect of development of advanced AMD. An evaluation for the diagnosis of age-related macular degeneration involves a comprehensive eye exam of the macula that hones in on aspects that are unique qualities of AMD (3). Factors that may be assessed are decreased vision, difficulty in adjusting to darkness, and a history of ocular issues. The process of gathering this information can be time consuming and requires advanced machinery to produce a clear image of the macula (4). One example of the type of machinery required to diagnose age-related macular degeneration successfully is an optical coherence tomography (OCT). The issue is not only on the premise of efficiency but costs too. The cost of purchasing an OCT can range from $40,000 to $150,000 (5). This cost can be a heavy burden for small eye clinics that are sporadically found in the regions outside of a city. As a result, we can find these advanced types of diagnostic machinery in inner city large hospitals and large eye clinics. Access to these types of machinery is a challenge that with affordable and efficient options will drastically lower the inequities regarding access to eye treatment.

As a result of the issues regarding the diagnostic processes of age-related macular degeneration a type of technology known as a convolutional neural network (CNN), a deep learning algorithm designed specially for images and pixel processing, can be especially used effectively in this scenario. This network is made up of three layers: a convolutional layer, a pooling layer, and finally the fully connected/output layer. They work in unison to take in image inputs and output a 4D array, which houses predictions and conclusions regarding the data (6). A CNN completes this task in a fraction of the time it would take with conventional methods, and it is functional in use for other contexts as well, especially in areas where a comprehensive exam may not be feasible. CNN models can be trained up to 99% accuracy; this process has high yielding results. Strategies that are needed to train CNN models for the assessment and diagnosis of ocular diseases have not been well characterized (7). By these ends, we trained a CNN model consisting of Convolution, Max Pooling, and ReLU layers on 5000 images of retinas affected by age-related macular degeneration, to conclude a diagnostic F1 score of 89%.

MATERIALS AND METHODS

Datasets

The source where our dataset is originally from is. The health-related images in the dataset are anonymised and de-identified. The dataset is empirical patient’s health information gathered by Shanggong Medical Technology Co., Ltd. from different hospitals/medical centers in China. The data of the patient's’ health is categorized and labeled as follows:

The original dataset is made up of patients whose ages range from 1 to 94, with the majority falling into the 55-64 age group. The distribution of sex is relatively equal, with 54% are male, and 46% are female.

From this dataset, we narrow our research topic down to the Age-Related Macular Degeneration (A) and Normal (N). There are no missing datas in the final dataset that we use for training our machine learning model.

Convolutional Neural Network

In our research, we utilize a ResNet 18 convolutional neural network for classification [1].
The Convolutional Neural Network consists of several different layers that are used in conjunction to process the image and output a result.

The first layer in the convolutional neural network was the Convolutional Layer. The convolutional layer is a layer in a neural network which uses weighted filters of a specific size to bring out the feature maps in an image. The kernel is an *n x n* square matrix which moves across the whole image, multiplying each pixel of the image by a constant value, which allows specific features of the image to come out. The stride of the convolutional layer is how many pixels the convolutional layer crosses over each time it brings out the features of a feature map. We use several different kernels with different strides in each layer to bring out many different features of the image. In the end, we combine the feature maps of each kernel in order to bring out a convoluted image. For this mode, we used a variety of filters and kernel sizes: we used 64, 128, 256, and 1x1, 3x3, 5x5 respectively.

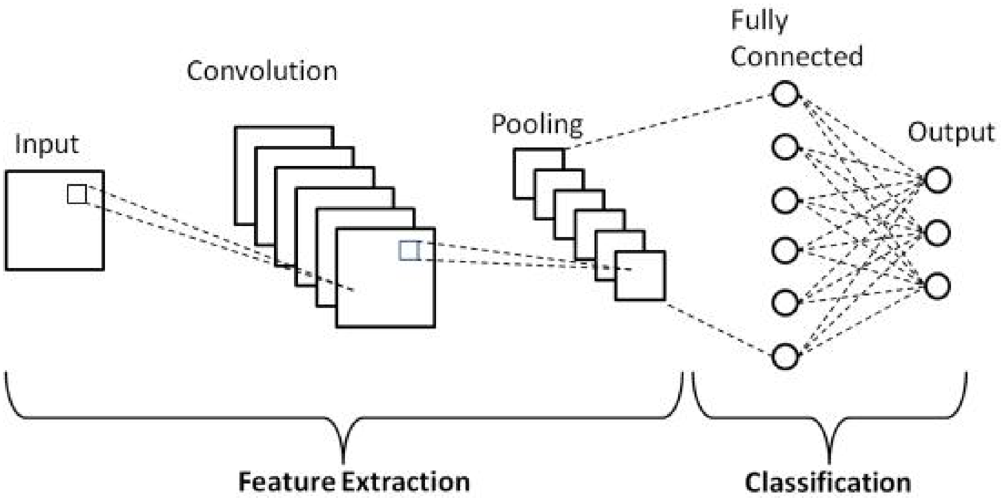

Figure 1. The structure of a Convolutional Neural Network

Another important layer that we used for this model was the Max Pooling layer. In order to make the convolutional network become more flexible with spatial variance of the inputs, the pooling layer is implemented through leveling down the resolution of the inputs. During the process, the number of parameters is reduced, which also helps prevent overfitting in the model. The pooling layer with pooled layer whose size is 2*2 and stride of two is commonly used, and each pooled feature map matches with the corresponding feature map of the convolutional layer. The popular procedure used in pooling is max pooling and average pooling. For the former method, the max element is picked out from the neighborhood and for the latter, the mean element is selected.

$$b_c^{(x,y)} = \max_{\delta_x, \delta_y \in \{-R, \cdots, R\}} \{ m_c^{\mathbf{z}_{\delta_x, \delta_y}} - \sum_{n=1}^{N} a_{c,n} d_{c,n}^{\delta_x, \delta_y} \}$$

$$\text{where } \mathbf{z}_{\delta_x, \delta_y} = (s_x \cdot x + \delta_x, s_y \cdot y + \delta_y).$$

Figure 2. The Max Pooling Layer Function

In addition, we used a RELU Layer as part of our Convolutional Neural Network as a way to enhance the nonlinearity that our model outputted. In real life, there is nothing really linear about images: the pixels are arbitrary, and it is not often that individual pixel values can be directly proportional to the output that the model will have. To combat this, we can make use of a RELU layer, which is a simple layer: it allows the value to pass as it is if it is greater than 0; however, if it is less than 0, the value is set to 0. This is highly effective in introducing non linearity in the model.

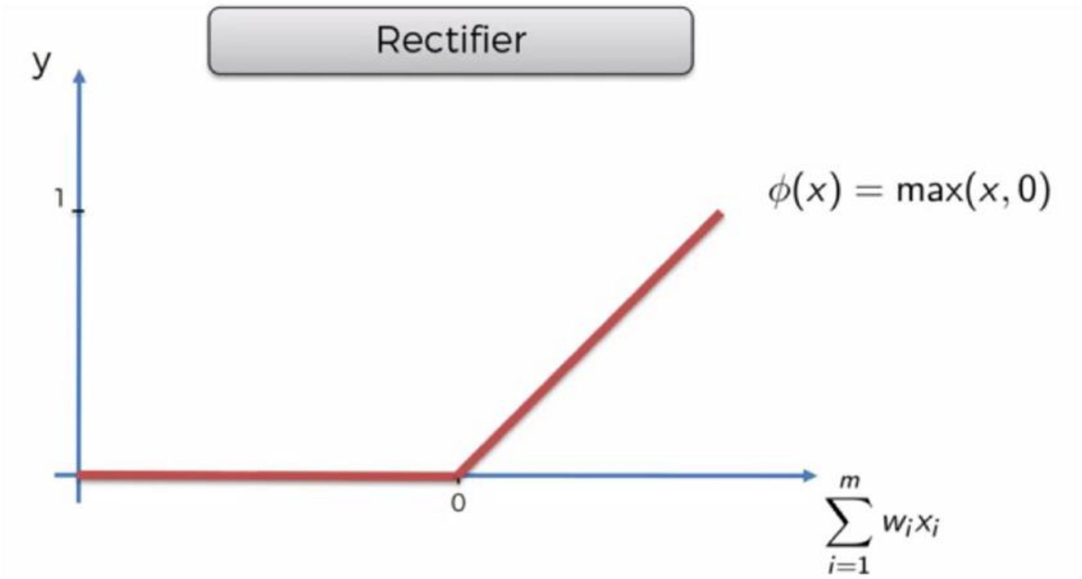

Figure 3. The RELU Layer Function

For our performance metric, we used F-Score, which is an evaluation metric useful for unevenly distributed datasets, as it gives a true result that isn't biased to class sizes. It's really useful in the medical field, as it takes into account recall and precision. This is helpful for the medical field because both of these values take into account false negatives - in a medical model, you don't want any false negatives, as it gives the patient false hope for treatment when they should have really gotten it. Because of this, F-Score was the best evaluation metric. To calculate F-Score, we used precision and recall, given that they are different from 0:

$$\frac{2}{F1} = \frac{1}{precision} + \frac{1}{recall} \text{ or } F1 = 2\frac{1}{\frac{1}{precision}+\frac{1}{recall}} = 2\frac{precision \times recall}{precision + recall}$$

The F-score value falls into the interval of (0,1]

.

RESULTS

The main objective of the research was to see if there is a correlation between different learning rates on our pre-compiled convolutional neural network model, Resnet18, offered by the PyTorch library. We ran several tests on different learning rates, and looked at different metrics to gauge the performance of the model on real world macular degeneration imagery.

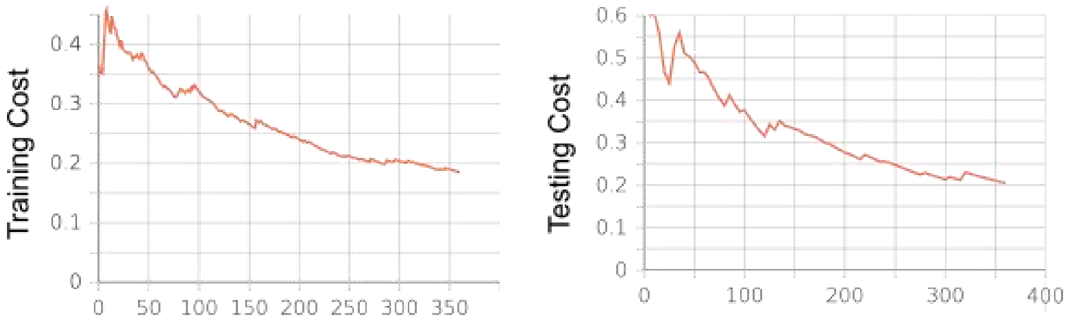

Figure 4
Unweighted base model

Initially, the tests were run on unweighted batches of images. An important remark is that our batches were unweighted. Our training set consisted of a significantly greater number of regular, non-macular degeneration disease images, compared to the images with the macular degeneration. It can be seen that the model has been very well trained, as can be seen by the training and testing costs graphs, which all indicate to be moving down as the model goes through more epochs of training.

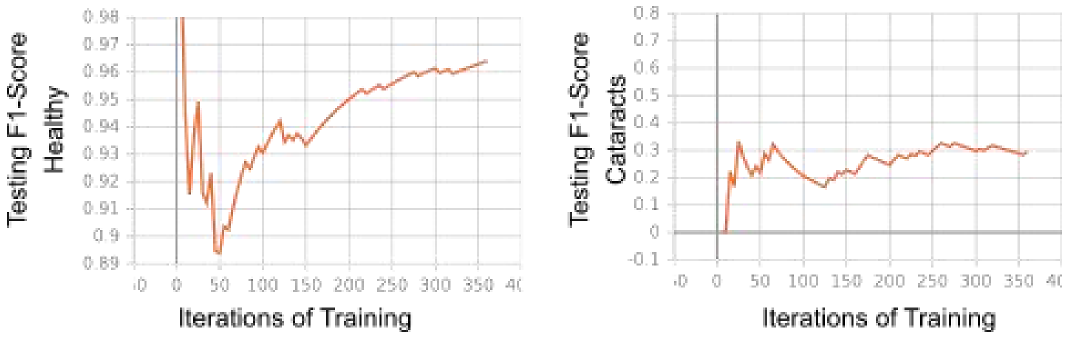

Figure 5
Weighted base model

However, as can be seen by Figure 2, this conclusion by the testing and training loss curve is deceptive. When using the F1 score to assess our model, we saw that the F1 score remained stagnant for the testing set, unlike what we would have assumed when solely looking at the training and testing costs. This is due to the model simply predicting no macular degeneration, and being correct most of the time, because of the imbalance in the dataset classes.

In our next two experiments, we made two main changes: the first one was to weight the batches of the images passed into the neural network, and the second one was to weight the cost function so that it is penalized heavier for misclassifying the macular degeneration image, so that it does not simply predicted no macular degeneration for all images and have a decreasing cost. This increased the F1 score by 47%.

DISCUSSION

We ran several tests on different learning rates, and looked at different metrics to detect age-related macular degeneration using a machine learning model. After analyzing over 5000 images, we were able to develop a model with an F1-Score of 0.89, indicating that the model was trained successfully. We saw that higher image size improves model performance, but there is a compromise on its learning rate. This would not be applicable in the real world where there would be millions of images. Additionally, we observed how a higher batch size meant that there would be more images of each category in every batch, so the model wouldn't be predicting the same result every time. Finally, we saw that weighting the dataset balances out the difference in number of images in each category.

It is also important to note certains sources of bias in our current model. We utilized a lower batch size due to constraints with our data collections. Additionally, a lower batch size lowers cost function. However, it can sometimes increase the training time for some machines that would support parallelism, such those with Graphical Processing Units or Tensor Processing Units. The best batch size depends on the use case scenario in combination with the training time and accuracy necessary.

ACKNOWLEDGMENTS

We would like to acknowledge Parsa Akbari for assisting [1] with generation of results and write-up of the report. [1] Downing College, Cambridge University, United Kingdom, CB2 1DQ

REFERENCES

1.“Common Eye Disorders and Diseases.” *Centers for Disease Control and Prevention*, Centers for Disease Control and Prevention, 3 June 2020, www.cdc.gov/visionhealth/basics/ced/index.html.

2. “Age-Related Macular Degeneration: Facts & Figures.” *Age-Related Macular Degeneration: Facts & Figures | BrightFocus Foundation*, www.brightfocus.org/macular/article/age-related-macular-facts-figures.

3.“Age-Related Macular Degeneration.” *National Eye Institute*, U.S. Department of Health and Human Services, www.nei.nih.gov/learn-about-eye-health/eye-conditions-and-diseases/age-related-macular-degeneration.

4.Li, Ji-Peng Olivia, et al. “Digital Technology, Tele-Medicine and Artificial Intelligence in Ophthalmology: A Global Perspective.” *Progress in Retinal and Eye Research*, Elsevier Ltd., May 2021, www.ncbi.nlm.nih.gov/pmc/articles/PMC7474840/.

5.Turbert, David. “What Is Optical Coherence Tomography?” *American Academy of Ophthalmology*, 8 Mar. 2022, www.aao.org/eye-health/treatments/what-is-optical-coherence-tomography.

6.Quetscher, Felizia. “A Comprehensible Explanation of the Dimensions in Cnns.” *Medium*, Towards Data Science, 24 June 2021, towardsdatascience.com/a-comprehensible-explanation-of-the-dimensions-in-cnns-841dba49df5e?gi=69def2f384b5.

7.Yamashita, Rikiya, et al. “Convolutional Neural Networks: An Overview and Application in Radiology - Insights into Imaging.” *SpringerOpen*, Springer Berlin Heidelberg, 22 June 2018, insightsimaging.springeropen.com/articles/10.1007/s13244-018-0639-9.

8. He, K., Zhang, X., Ren, S. & Sun, J. Deep Residual Learning for Image Recognition. *arXiv [cs.CV]* (2015).

9. Larxene. (2020, September 24). *Ocular disease recognition*. Kaggle. Retrieved October 15, 2021,

from https://www.kaggle.com/andrewmvd/ocular-disease-recognition-odir5k

10. Transitive transfer learning. (2020). *Transfer Learning, 151–167*. https://doi.org/10.1017/9781139061773.013

11. "Odir-2019 - Grand Challenge." *Grand*, odir2019.grand-challenge.org/.