IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

Detecting Diabetic Retinopathy using Deep Learning

Ramakrushna Swain, Sovit Ranjan Rath, Jogesh Agrawal, Harshit Verma, Rashmiranjan Pradhan

Department of Computer Science & Engineering, Silicon Institute of Technology, Bhubaneswar

Abstract**—***Diabetic Retinopathy is one of the major causes of blindness around the world. This disease affects those people who suffer from diabetes and mainly aged diabetes patients. Many hospitals around the world try their best to conduct research and prevent blindness caused due to diabetic retinopathy. Being able to detect diabetic retinopathy before it is too advanced is one of the major aims of this project. In this work, we will try to detect whether a patient has reached the stage of diabetic retinopathy or not. In accordance with this, we will use deep learning neural networks to classify the images of eyes of patients. Medical imaging analysis using deep learning and neural networks is gaining popularity because of its effectiveness and positive results. In this project, we will train a neural network on the retina images taken using fundus photography. After the training the neural network will be able to classify whether a patient has diabetic retinopathy or not and to what extent. The successful completion of this project may help many medical practitioners carry out the eye examination of diabetes patients in an easy and effective manner. The detection of diabetic retinopathy will also become more accurate and less time consuming. This will lead to better treatment and also a reduction in cost as physical examination of the retina images by technicians will be minimal.*

Keywords—Deep learning, neural networks, image classification, medical imaging, diabetic retinopathy, medical imaging analysis, machine learning.

INTRODUCTION

Diabetic retinopathy is a leading cause of blindness among working-age adults. Early detection of this condition is critical for good prognosis. Many adults who are suffering from diabetes have a high probability of being diagnosed with diabetic retinopathy at some stage of their life. In general, diabetic retinopathy is curable if detected at the early stages. But if the patient is diagnosed at a very severe stage, then it may not be curable at all. Therefore, in this project, we try to detect whether a person suffering from diabetes has diabetic retinopathy or not. For that we are employing the power of deep neural networks from the field of deep learning. In specific, we will be using the CNN (Convolutional Neural Network) architecture in this work.

We will be using a large set of retina images taken using fundus photography under a variety of imaging conditions. An example of how the images look like is given below. Our main aim will be to use different types of images, different learning algorithms, and different neural network architectures and compare which combination gives the best results for detection of diabetic retinopathy.

Figure 1 : A retina image taken using fundus photography

A. Fundus Imaging

Fundus photography involves photographing the rear of an eye; also known as the fundus. Specialized fundus cameras consisting of an intricate microscope attached to a flash enabled camera are used in fundus photography. The main structures that can be visualized in a fundus photo are the central and peripheral retina, optic disc and macula. Fundus photography can be performed with colored filters, or with specialized dyes including fluorescein and indocyanine green.

DATA COLLECTION

Our current data has been collected from the Kaggle website. The dataset contains retina scan images taken using fundus photography under a variety of imaging conditions. The dataset contains 3,662 images for training 1,928 images for testing.

A. Type of Data

The default images provided on the website are colored images. They are very high-resolution images some ranging to around 4k dimensionality. We have resized the images for practicability and more description can be found in the section where we have defined our approaches.

B. Division in the Dataset

The dataset has been divided into a train directory, test directory, train.csv file, and test.csv file.

The data is divided into five classes based on the severity of diabetic retinopathy. The following are the five classes.

0 - No DR

1 - Mild

2 - Moderate

3 - Severe

4 - Proliferative DR

The numbering is done based on increasing severity where 0 means no diabetic retinopathy and 4 means the highest stage, that is, proliferative diabetic retinopathy.

In the dataset we have 1,805 images with no diabetic retinopathy, 370 images with mild, 999 images with moderate, 193 images with severe, and 295 images with proliferative diabetic retinopathy.

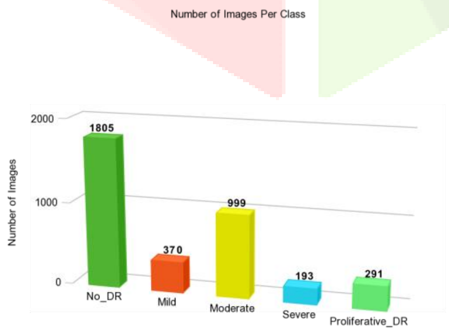

Figure 2 : Number of Images Per Class

Train directory contains all the images of retina that we train and validate our neural network model on. Train CSV file contains the labels corresponding to the images in the train directory. Similar is the structure for the test data as well. But the test CSV file does not contain the labels.

APPROACH

We are classifying the severity of diabolic retinopathy by the use of neural networks.

A. Neural Network Model

The classification is done by ResNet34[4]. This model contains 34 CNN layers. The first layer uses 7*7 filter and the next layers use 3*3. It uses Global average pooling layer and a 5-way fully-connected layer with Softmax in the end.

ResNet34 solves the degradation problem i.e. the accuracy gets saturated and degrades rapidly on the increase in network depth.

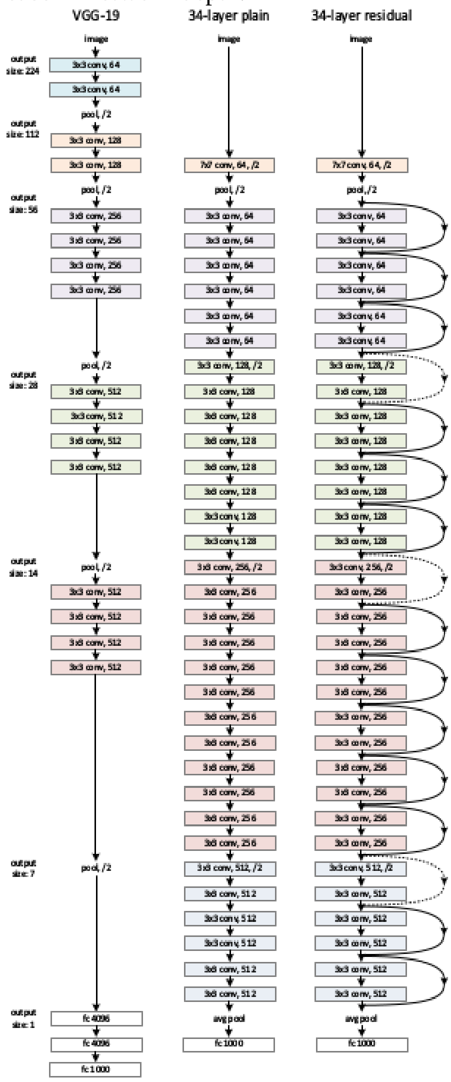

Figure 3 : ResNet34 blocks

EXPERIMENTAL RESULTS

For the first stage of our experimentation we have trained the ResNet34 for 20 epochs and for 3 epochs after finding the optimal learning rate of the neural network. This process is repeated for all the three different colored datasets (colored, grayscale, and gaussian).

A. Colored Images Results

During the initial training of colored images, for the first 20 epochs the best error rate is 0.195355. After learning rate optimization we achieved the best error rate of 0.178962. The lowest training loss if 0.441782 and validation loss is 0.541847.

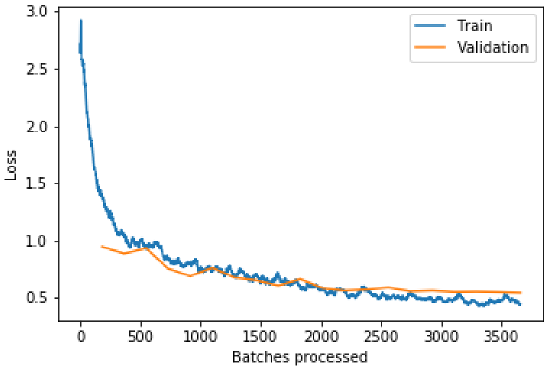

Figure 4 : Colored Images Traning Loss After 20 Epochs

B. Grayscale Images Results

For grayscale images, the lowest error rate after 20 epochs of training is 0.196721 with best training loss of 0.480602 and best validation loss of 0.520497.

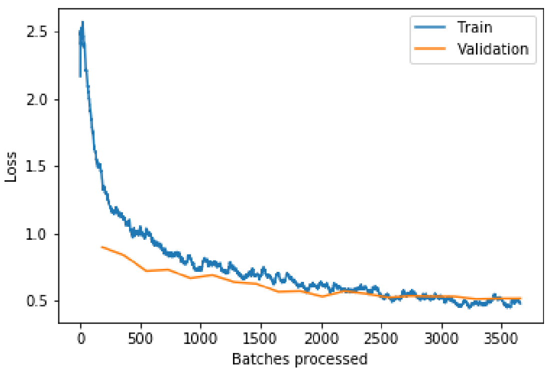

Figure 5 : Grayscale Images Training Loss After 20 Epochs

C. Gaussian Filtered Images Results

For grayscale images, the lowest error rate after 20 epochs of training is 0.221311 with best training loss of 0.478799 and best validation loss of 0.578836.

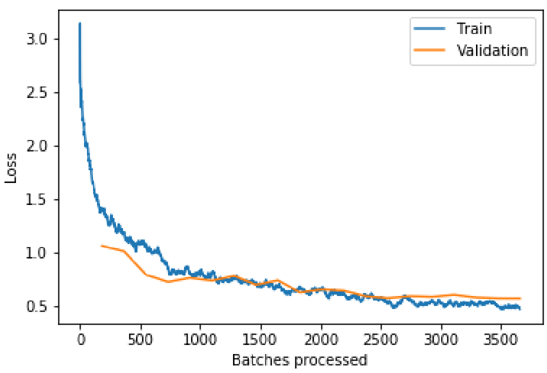

Figure 6 : Gaussian Filtered Images Training Loss After 20 Epochs

D. Training Result Tables

The following tables show the results of training on colored, grayscale, and gaussian filtered retina images. We have trained each of the image sets for 20 epochs. In the following tables we are showing the results of the last three epochs only.

epoch	**train_loss**	**valid_loss**	**error_rate**
18	0.456303	0.551894	0.207650
19	0.483539	0.548007	0.214481
20	0.441782	0.541847	0.211749

Table 1. Last three epoch results for colored retina images

epoch	**train_loss**	**valid_loss**	**error_rate**
18	0.492623	0.516096	0.200820
19	0.512907	0.520497	0.200820
20	0.480602	0.521209	0.196721

Table 2. Last three epoch results for grayscale retina images

epoch	**train_loss**	**valid_loss**	**error_rate**
17	0.519309	0.587234	0.229508
18	0.510745	0.579334	0.224044
19	0.478799	0.578836	0.221311

Table 3. Last three epoch results for gaussian filtered retina images

From the above tables we can see that the training loss after 20 epochs is the least for colored retina images (Table 1). For the error rate, we get the least error rate in the case of training grayscale images. But the training loss is more in case of grayscale images than in colored images.

The neural network seems to be performing the worst in the case of gaussian filtered images when taking the error rate as a result metric. Also, the validation loss is the highest in this case.

The results show that perhaps, training a bigger neural network model on a mixture of all images may provide much better results than we are currently getting.

E. Validation Results

The validation results provide a good insight of how good/bad our neural network model is performing while validating on the dataset.

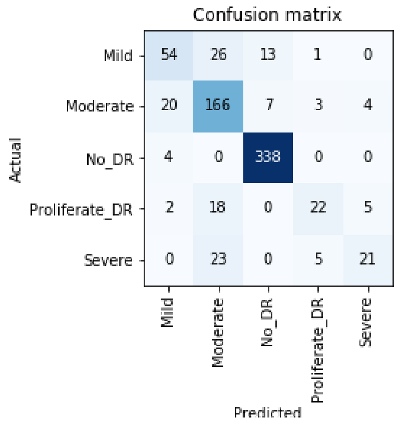

Figure 7 : Confusion Matrix for Colored Images Validation

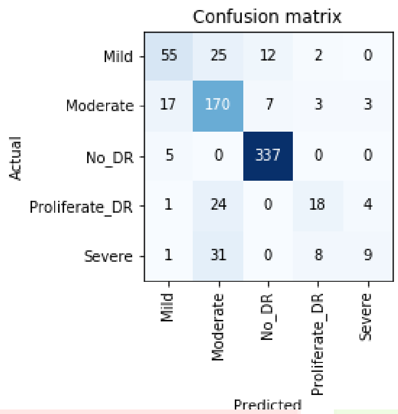

Figure 8 : Confusion Matrix for Grayscale Images Validation

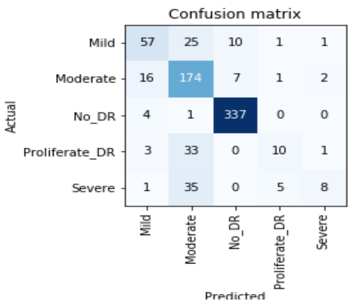

Figure 9. Confusion Matrix for Gaussian Filtered Images Validation

Figures 7, 8, and 9 shows the confusion matrix interpolation plot for the validation of the retina images. The figures correspond to colored images validation, grayscale images validation, and gaussian filtered images validation results respectively.

We can clearly observe that in all cases, the ResNet34 neural network model is perfectly classifying the *No_DR* class. This is because, this class has the most number of images (almost 1800 images) and it is getting enough images for that class to learn the patterns.

Whereas, for the other classes, the number images are quite less, and below 1000 images per class as well.

The validation mistakes in some cases are quite profound as well. For example, in the case of colored retina images validation, the neural network model is classifying 18 *Moderate* class as *Proliferate_DR* class. And in total, it is ***misclassifying 131 images out of 732 validation images***.

In the case of grayscale validation images, the neural network model is ***misclassifying 143 images out of 732 validation images.*** Here, validation results are worse than the colored images validation results.

In the case of gaussian filtered validation images, the neural network model is ***misclassifying 146 images out of 732 validation images.*** The validation results are the worst here. Mostly because, after applying gaussian filtering to the retina images, they are losing many of the pixel and color features. Therefore, the neural network model is finding it difficult to learn all the important features.

CONCLUSION AND FUTURE WORK

During the next stage of our experimentation, we intend to carry out the following: Train a ResNet34 model on a mixture of all images. Collect more data for training to achieve better results. Train bigger ResNet models like ResNet50.

The detection of diabetic retinopathy become more accurate, less time consuming, reduction in cost and lead to better treatment.

VI. REFERENCES

[1] Carson Lam, Darvin Yi, Margaret Guo, Tony Lindsay, "Automated Detection of Diabetic Retinopathy using Deep Learning". AMIA Jt Summits Transl Sci Proc. 2018; 2018: 147–155.

[2] Jaakko Sahlsten, Joel Jaskari, Jyri Kivinen, Lauri Turunen, Esa janio, Kustaa Hietala, Kimmo Kaski, "Deep Learning Fundus Image Analysis for Diabetic Retinopathy and Macular Edema Grading," Science Reports, 2019.

[3] Anonymous authors, "Fixing Weight Decay Regularization in Adam,". Under review as a conference paper at ICLR 2018

[4] Kaiming He, Xiangyu Zhang, Shaoqing Ren, Jian Sun, "Deep Residual Learning for Image Revognition," Microsoft Research, 2015

[5] Goodfellow et al., "Convolutional Networks", Deep Learning, p.326 - p.366., 2016