CALIFORNIA STATE SCIENCE FAIR
2017 PROJECT SUMMARY

Name(s)

Anish R. Neervannan

Project Number

S1518

Project Title

A State-of-the-Art Approach to Detect Diabetic Retinopathy Using Convolutional Neural Networks

Abstract

Objectives/Goals

The purpose of this project was to determine if a state-of-the-art convolutional neural network image recognition algorithm could be developed to detect diabetic retinopathy in Fluorescein Angiography (fudus) photography to match the accuracy of a medical professional's mental model.

Methods/Materials

Materials included a Macbook with high computing power and internet access; a dataset of 35,000 retinal scans from Kaggle's diabetic retinopathy homepage; Inception v3 and ImageNet models; and machine learning libraries such as Tensorflow, Keras, scikit-learn, and matplotlib. A Java program was written to sort the scans based on their pre-classified labels. Python programs were written using the scikit-learn and matplotlib libraries to crop, reshape, denoise, visualize, and transpose the scans. The scans were split into a training set, with 29000 scans, a validation set, with 5000 scans, and a test set, with 200 scans. The Keras and Tensorflow libraries were used to build the Convolutional Neural Network architecture on top of the Inception base model and train the algorithm. The retinal scans in the test set were passed into the trained model and the algorithm's predictions were recorded for each image. The same set of scans was sent to three medical professionals.

Results

After the training process, a controlled test set of 200 scans was used to determine the accuracy, precision scores, and recall scores of the algorithm and the three medical professionals. The image recognition algorithm's accuracy was 79.50%, while the medical professionals' accuracies were 88.00%, 89.00%, and 88.00%.

Conclusions/Discussion

Although the computer algorithm's accuracy came close to the medical professionals' accuracies, it did not match them. The algorithm predicted scans without the disease more consistently than those with the disease. This algorithm breaks new grounds in this field by outperforming a prior model that was published in a Stanford PhD paper (trained on the same image dataset) by an 8% margin.

Summary Statement

I trained a convolutional neural network to detect diabetic retinopathy to be almost as accurate as a medical professional and better than a prior model published in a Stanford PhD paper.

Help Received

Dr. Kapil Kapoor (Vitreoretinal Surgeon), Dr. David Chia (Ophthalmologist), and Dr. Sanjay Kedhar (Gavin Herbert Eye Institute) classified the images in the test set. My father, Raj Neervannan, helped me understand the concepts required to complete the project.