CALIFORNIA SCIENCE & ENGINEERING FAIR
2018 PROJECT SUMMARY

Name(s)
Anish R. Neervannan

Project Number
38762

Project Title
A State-of-the-Art Approach to Detect Diabetic Retinopathy Using Convolutional Neural Networks and Adversarial Networks

Abstract

Objectives/Goals
The purpose of this project was to determine if a state-of-the-art generative adversarial network (GAN) could be used in conjunction with a convolutional neural network (CNN) image recognition algorithm to detect diabetic retinopathy in fudus photography more accurately than an independent CNN and match the accuracy of a medical professional's mental model.

Methods/Materials
The materials used included a System76 machine; a dataset of 35,000 retinal scans downloaded from Kaggle's diabetic retinopathy homepage (available for public use); the Inception Resnet v2 model; and Tensorflow, Keras, and other machine learning libraries. The image dataset was downloaded and retinal scans were randomly split into training & validation sets and test sets. Multiple programs were written in Python to crop, reshape, visualize, and transpose the scans. Keras and Tensorflow libraries were used to build the Convolutional Neural Network architecture on top of the Inception base model and train the algorithm. A separate Generative Adversarial Network architecture was built, the training data was fed through, and the generated images were added back to the original dataset for another round of training. 200 images from the test set were passed into both trained models and sent to three medical professionals. These predictions were used to compute precision, recall, and accuracy.

Results
The GAN+CNN algorithm's accuracy was 72.00%, the CNN-only algorithm's accuracy was 74.00%, and the three medical professionals' accuracies were 76.00%, 84.00%, and 88.00%. On scans without diabetic retinopathy, the CNN-only algorithm got a precision score of 75.00% and a recall score of 100.00%, the GAN+CNN algorithm got a precision score of 87.80% and a recall score of 100.00%, and the medical professionals got precision scores of 97.15%, 92.68%, and 88.50% and recall scores of 70.00%, 29.39%, and 71.24%.

Conclusions/Discussion
The hypothesis was not supported; even though both computer algorithms came close to the medical professionals' accuracies, it did not match them. In addition, all tests indicated that the GAN + CNN combination classifier came close to did not outperform the CNN-only classifier. However, this report broke new grounds in that the CNN-only classifier achieved a slightly higher accuracy than that of last year's project, which had outperformed a prior model that was published in a Stanford paper by an 8% margin.

Summary Statement
I developed a generative adversarial network to work in conjunction with a convolutional neural network to detect diabetic retinopathy nearly as accurately as a trained medical professional.

Help Received
Dr. Kapil Kapoor, Dr. David Chia, and Dr. Sanjay Kedhar were the three ophthalmologists that classified my test set data for comparison with my algorithms' results. My parents, Raj and Aruna Neervannan, and my science teachers, Mr. David Knight, Mr. Tim Smay, and Mr. Nicholas Brighton, reviewed my report