Deep CNNs for Detecting Diabetic Retinopathy

Iain Usiri (iusiri@), Chala Fufa (cfufa@), Alex Tamkin (atamkin@)
CS229, Stanford University

What is Diabetic Retinopathy?

- DR is an eye disease caused by diabetes that can lead to loss of vision or even complete blindness.
- DR is the leading cause of blindness for people aged 20 to 64 years. It is estimated to affect about 93 million people globally, though only half are aware of it.
- If caught early enough, vision degeneration can be slowed if not stopped, but this is often difficult because symptoms may appear too late to provide effective treatment.

DR is characterized by lesions. These lesions identify blood vessels that have burst/leaked and formed scar tissue.

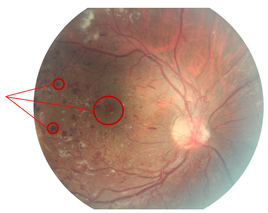

Retina Image - Stage 4 DR

Goal

Our objective was to create a model that would take a preprocessed retina image as input and output whether or not it had stage 4 diabetic retinopathy.

Dataset

We are using a dataset of retina images from a Kaggle competition. These are a set of high resolution color retina images taken with a variety of cameras in different lighting conditions. Each image is tagged with a clinician's DR classification.

Data Preprocessing

1. Resized each image to 256 by 256 pixels.
2. Rescale each image to have the same radius (the eyeball).
3. Clip the edges of the images to minimize the variation on the boundaries or edges of the images.
4. Subtract the local average color from each pixel, mapping the average to 50% gray.

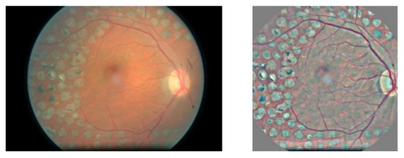

Retina image before and after pre-processing

Models

- We used deep convolutional neural networks (CNNs) and transfer learning, where one takes a pre-trained model and adapts it to a new domain.
- The pre-trained model we used was the Inception V3 model, with ImageNet weights.
- We then added a fully-connected layer of size 1024, and a fully-connected output layer and trained both with the Adam optimizer. We used dropout and maxnorm regularization.
- We also trained the top two blocks (sections) of the Inception V3 model in addition to the two dense layers, using SGD with a low learning rate to avoid overfitting.

Results

	Accuracy on Training Set (600 samples)	Accuracy on Validation Set (600 samples)
Training last two layers	0.8950	0.8517
Training last two layers and top two Inception V3 blocks	0.9417	0.8733

Discussion

- Our results were better than expected, with our best model correctly classifying around seven in eight examples correctly in the validation set (87.33%).
- The Inception VC model likely has good low-level feature extractors towards the bottom of the network, which allowed it to be repurposed for this fairly different task.
- Since the top layers in a deep network are speculated to contain higher-level, more domain-specific features, training the top blocks probably produced better performance by adapting these blocks to the diabetic retinopathy task.
- These results give hope that computational models could help catch DR early in populations across the world without access to trained ophthalmologists.

Future Work

- Implement multiclass classification, attempting to diagnose not just the presence of DR but also the stage.
- Experiment with different architectures and regularization schemes.
- Perform data augmentation to reduce overfitting.

References

- Szegedy, Christian, et al. "Rethinking the inception architecture for computer vision." *arXiv preprint arXiv:1512.00567* (2015).
- Oquab, Maxime, et al. "Learning and transferring mid-level image representations using convolutional neural networks." *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2014.