Retinal Blood Vessel Segmentation and Exudates Detection Using Deep Neural Networks

Shruti Sharan [1], Anamika Awasthi [2], Satya [3], A.Aruna [4]
[1, 2, 3] Student, SRMIST, Chennai, Tamil Nadu, India.
[4] Assistant Professor, SRMIST, Chennai, Tamil Nadu, India.

Abstract – The STARE (STructured Analysis of the Retina) Project was conceived and initiated in 1975 by Michael Goldbaum, M.D., at the University of California, San Diego. It was funded by the U.S. National Institutes of Health. Image segmentation is a major part of this project and in our attempt to contribute further to this endeavor, we have implemented the project using deep neural networks. The Generative Adversarial Networks have proved to be a great alternative for the previous algorithms employed for image classification and have superseded performance by several measures. In this paper we have used GANs to train our model and predict whether the images are healthy or not.

Index Terms –Retinal, Segmentation, Generative Adversarial Networks (GAN), detection, Neural Networks, Ophthalmology.

1. INTRODUCTION

An ophthalmologist is a medical doctor that specializes in the structure, function, and diseases of the human eye. During a clinical examination, an ophthalmologist notes findings that are visible in the eyes of the subject. The ophthalmologist then uses these findings to reason about the health of the subject. For instance, a patient may exhibit discoloration of the optic nerve, or a narrowing of the blood vessels in the retina. An ophthalmologist uses this information to diagnose the patient, as having for instance Coats' disease or a central retinal artery occlusion.

A common procedure during an examination is retinal imaging. An optical camera is used to see through the pupil of the eye to the rear inner surface of the eyeball. A picture is taken showing the optic nerve, fovea, surrounding vessels, and the retinal layer. The ophthalmologist can then reference this image while considering any observed findings.

This research concerns a system to automatically diagnose diseases of the human eye. The system takes as input information observable in a retinal image. This information is formulated to mimic the findings that an ophthalmologist would note during a clinical examination. The main output of the system is a diagnosis formulated to mimic the conclusion that an ophthalmologist would reach about the health of the subject.

Our approach breaks the problem into two components. The first component concerns automatically processing a retinal image to denote the important findings. The second component concerns automatically reasoning about the findings to determine a diagnosis. Additional outputs include detailed measurements of the anatomical structures and lesions visible in the retinal image. These measurements are useful for tracking disease severity and the evaluation of treatment progress over time. By collecting a database of measurements for a large number of people, the STARE project could support clinical population studies and intern training.

2. RELATED WORK

Although Generative Adversarial Network is a recent breakthrough but the application for image processing are already numerous bearing a testimony to the efficacy of the adversarial training paradigm.

Initially GAN was introduced as a Multi-Layer Perceptron network, but these networks suffered from high rate of non-convergence. GAN shortcomings were further investigated and improvements were suggested to meet them. LAPGAN presented the idea of employing convolutional neural networks with the adversarial training. Another version of GAN called Deep Convolutional GAN or DCGAN was also presented exploring unsupervised learning with Convolutional Neural Networks for discriminator and generator.

DCGAN model was explored in detail and architectural guidelines like using strided convolutions instead of up sampling, de-convolution for avoiding pooling layers were presented. Batch normalization was suggested and fully connected layers were completely removed from the networks except the first layer in Generator and the last in the Discriminator. They also advocated Relu layers as most suitable activation functions for Generators and leaky Relu for Discriminators. Image generation was also conditioned to a label in a number of ways. An MLP implementation incorporated class labels as one hot vectors mapped to additional hidden dense layer before being combined with hidden layer generated by latent noise.

Another application of GAN employed Convolutional Neural Networks while both the generator and discriminator were conditioned on the class label implemented as embedding layer. SGAN demonstrated semi-supervised approach by training the discriminator to predict labels for images and then the generator utilizing these labels for generation. Image generation and classification through GAN was further improved through techniques like feature matching, mini-batch discrimination, historical averaging, label smoothing and batch normalization. ACGAN adopts the semi-supervised approach and demonstrates that introduction of a label results in faster convergence and stable performance.

3. PORPOSED MODELLING

Generative Adversarial Networks consist of a Generator and a Discriminator competing against each other. The Discriminator is constantly trying to catch the fake output produced by a Generator while the Generator tries to fool it. The Discriminator should be updated in such a manner so as to be able to differentiate between the true data and fake output produced by generator while the Generator should be updated so as to minimize the discrimination. Goodfellow et.al proved that there was a unique solution to the problem with Generator covering the entire data distribution and discriminator getting totally confused. The Generator network G tries to map a noise vector z with distribution p(z) to a data distribution p(x). While the Discriminator network D tries to differentiates between data distribution and the generated distribution.

The Discriminator tries to maximize the value function while Generator tries to minimize it. The first term on the right hand side in the above equation denotes the loss for real data while the second term is the loss for generated data. When the Discriminator is totally confused then D(x) = 1/2 and generative distribution p_g is equal to data distribution p_{data}.

In the current effort a three dimensional imaged reconstruction is implemented through GAN incorporating image class labels. The discriminator as well as generator consisted of four convolution layers as shown in figure. The leaky relu activation function was used. The discriminator input was a batch of images to be discriminated. The discriminator also had batch normalization and dropout layers to improve performance. There were two categorical outputs obtained sigmoidal activation functions on the final flattened layer. The first output was binary presenting whether the images were actual or generated while the second output denoted the image class of weather it is a healthy eye or a defective one.

The generator required a latent noise vector from a uniform probability distribution as well as the desired class of the image to generate. The embedding of image class vector was applied to the latent noise thus implementing class conditioning on generated images. Apart from four convolution layers an initial dense layer and two up-sampling layer were also added to get the desired image dimensions. The generated image was of the same dimensions as the actual images.

4. RESULTS AND DISCUSSIONS

Training:

Adversarial training trains both discriminator and generator against each other. The training strategy was same as. The Figure 1 describes the training process for the GAN.

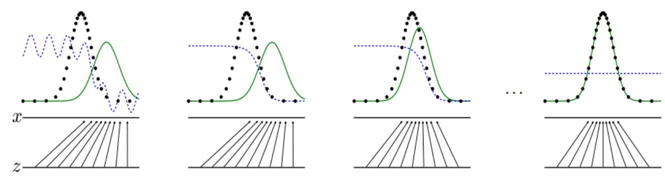

Initially generator was used to generate images using random class labels. The discriminator was then used to predict whether the images were real or fake as well as the class of the image. The first batch of images provided to the discriminator were actual images from the data set. The discriminator was trained to make correct prediction. Then it was trained once more on fake images. It has been accessed from past implementations that training on a batch of all real images and then a batch of all fake images provided better performance than a batch of mixed images. As the discriminator was trained twice for each batch, the generator was also trained twice. The loss of the generator was taken from the discriminator loss function if all the generated images were actual images. Thus training the generator so as to reduce this loss and generate images that have very close resemblance with actual eye images. For the present work 30 epochs were used for training.

Both generator and discriminator losses are calculated at each step and the mean value of the test and train losses are computed. The weights from each epoch are saved. Once the training is done, the generator model is built from latent space and the weights are loaded into it. The discriminator is also loaded with the weights and the model is built using the generated images from the generator.

Ophthalmoscopy:

An ophthalmoscopy is a very important part of diabetic retinopathy diagnosis because it allows the ophthalmologist to see the entire back portion of the eyeball, which includes the optic disc, choroid, retina, and blood vessels. The test is a common part of routine eye exams and takes just a few minutes to complete. During the examination, the doctor will beam a bright light through your pupil using an, which has a series of rotating lenses through which the back of the eye can be viewed. Because the human eye is a natural magnifier, the ophthalmologist is able to easily, the most popular being the “air puff” test. During this test, the ophthalmologist uses a special instrument to calculate the IOP by measuring changes in the light reflected off the corneas when the air puff is blown into the eyes. A more specialized form of this tomography is called optical coherence tomography (OCT). It produces an image of the inner structures of the eye to detect macular edema, which is one of the telltale signs of diabetic retinopathy.

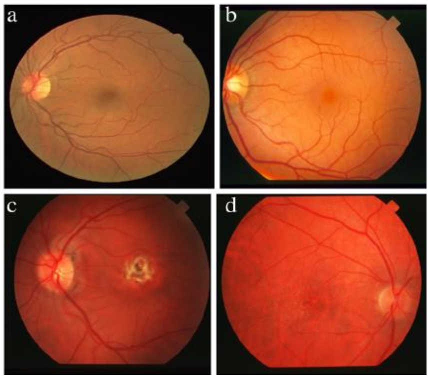

Figure 2: Some of the retinal images in DRIVE and STARE dataset (a,b,c,d)

Segmentation:

Human eye is a very sensitive organ to deal with in Medical circumstances. Its segmentation in bodily structure imaging could be a nontrivial task because of variable size of vessels, comparatively low distinction, and potential presence of pathologies like small aneurysms and hemorrhages. Several algorithms, each unsupervised and supervised, have been planned for this purpose within the past. We tend to propose a supervised segmentation technique that uses a deep neural network trained on an outsized (up to four hundred, 000) sample of examples, preprocessed with international distinction standardization, zero-phase lightening, and increased mistreatment geometric transformations and gamma corrections. Many variants of the strategy are thought of, as well as structured prediction, wherever a network classifies multiple pixels at the same time. Once applied to plain benchmarks of bodily structure imaging, the DRIVE, STARE, and CHASE databases, the networks considerably crush the previous algorithms on the world below mythical monster curve live (up to > zero.99) and accuracy of classification (up to > zero.97). The strategy is additionally proof against the development of central vessel reflex, sensitive in detection of fine vessels (sensitivity >zero.87), and fares well on pathological cases and the ones with Ilasitik.

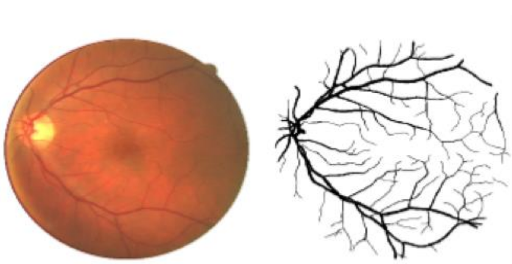

Figure 3: A training image from the DRIVE database (left) and the corresponding manual segmentation (right).

Diabetic Retinopathy Diagnosis:

Diabetic retinopathy occurs when the tiny blood vessels, known as capillaries, within the retina are damaged. In patients with non-proliferative diabetic retinopathy (NPDR), the walls of the capillaries weaken and develop micro aneurysms, or tiny bulges protruding from the blood vessels. Eventually these micro aneurysms begin to leak blood and fluid into the retina, causing vision loss.

In patients with proliferative diabetic retinopathy (PDR), not only are there progressively more micro aneurysms, but new, abnormal capillaries begin to develop within the retina. As these blood vessels spread throughout the retina, they often begin to grow into the jelly-like substance (vitreous) that fills the center of the eye. Ultimately, this abnormal growth causes the capillaries to shut down, leading to vision loss and, in some cases, retinal detachment.

In addition to diabetic retinopathy, there are two other eye diseases associated with diabetes: cataracts and glaucoma. These conditions are also treatable and preventable, but require comprehensive eye care.

Both type I and II diabetes patients are at a very high risk of developing diabetic retinopathy. How much the disease progresses and spreads is in almost direct correlation to how long the patient has had diabetes and how long they have gone without consistent eye examinations.

Diabetes is currently the number one cause of new cases of blindness in the United States; serious complications from diabetic retinopathy affect approximately 24,000 new people each year. However, studies also show that given adequate preventative measures and the right diabetic retinopathy treatment plan, severe vision loss can be reduced by as much as 94 percent. So while all diabetes patients are at risk of developing diabetic retinopathy, not all of them are destined for blindness. Undergoing yearly eye exams and tests for diabetic retinopathy diagnosis are crucial steps to preventing total vision loss.

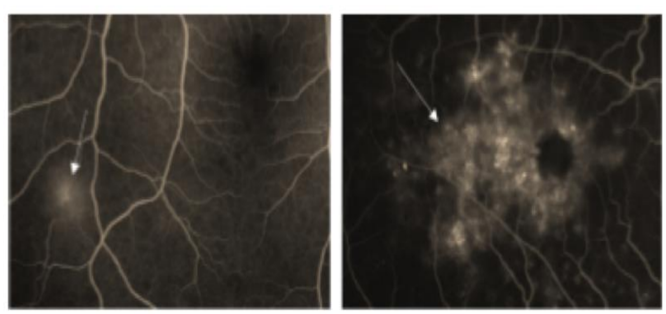

Figure 4: Illustration of focal leakages in two retinal diseases. Left: malarial retinopathy. Right: diabetic retinopathy. There is a large increase in brightness in leaking regions (white arrows) compared to surrounding non-leaking regions.

5. CONCLUSION

In this Paper they enforced on a completely unique efficient thanks to enhance outflow regions by exploitation the thought of prominence. Prominence indicates the relative importance of visual options, and is closely associated with the characteristics of human perception and process of visual stimuli. Prominence emerges from such characteristics in options of the image as visual individuation, unpredictability, or rarity, and is usually attributed to variations in specific image attributes like color, gradient, edges, and bounds. Such attributes are characteristics of retinal outflow in sulfa syllable pictures. For instance, outflow of fluorescent dye causes an oversized increase in brightness in unseaworthy regions in comparison to close non-leaking regions. For this application, unseaworthy regions are often defined as those of high strikingness.

In consequence, we tend to area unit driven to firstly determine the unseaworthy regions in sulfa syllable pictures through a prominence detection technique, then estimate their areas from the obtained prominence map. And, they improved on multi scale prominence maps with integration of the intensity and compactness cues of super pixels for this specific application. a lot of specifically, ancient prominence extraction ways typically cypher the strikingness of a picture in an exceedingly pixel-by-pixel manner, and ignore the neighborhood and edge data of the objects of interest. Conclusion part depicts the main points as the constructive finds obtained from the proposed system. Conclusion should not be the same as abstract. Conclusion should be modelled efficiently.

REFERENCES

[1] J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. WardeFarley, S. Ozair, A. Courville, and Y. Bengio. Generative Adversarial Networks. ArXiv e-prints, June 2014.

[2] I. J. Goodfellow. On distinguishability criteria for estimating generative models. ArXiv e-prints, December 2014.

[3] Emily L. Denton, Soumith Chintala, Arthur Szlam, and Robert Fergus. Deep generative image models using a laplacian pyramid of adversarial networks. CoRR, abs/1506.05751, 2015.

[4] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. CoRR, abs/1511.06434, 2015.

[5] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. CoRR, abs/1411.1784, 2014.

[6] Jon Gauthier. Conditional generative adversarial nets for convolutional face generation. Class Project for Stanford CS231N: Convolutional Neural Networks for Visual Recognition, Winter semester, 2014(5):2, 2014.

[7] A. Odena. Semi-Supervised Learning with Generative Adversarial Networks. ArXiv e-prints, June 2016.

[8] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In Advances in Neural Information Processing Systems, pages 2234–2242, 2016.

[9] A. Odena, C. Olah, and J. Shlens. Conditional Image Synthesis With Auxiliary Classifier GANs. ArXiv e-prints, October 2016.