M.Sc. Engg. Thesis

Automatic Detection of Diabetic Retinopathy from Fundus Images Using Image Processing and Artificial Neural Network

by
Ishtiyaque Ahmad

Submitted to

Department of Computer Science and Engineering

in partial fulfilment of the requirements for the degree of
Master of Science in Computer Science and Engineering

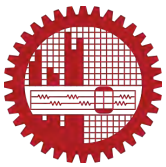

Department of Computer Science and Engineering

Bangladesh University of Engineering and Technology (BUET)

Dhaka 1000

January 2019

Dedicated to my loving parents

Author's Contact

Ishtiyaque Ahmad

Lecturer

Department of Computer Science and Engineering

Bangladesh University of Engineering and Technology (BUET).

Email: ishtiyaque_ahmad@cse.buet.ac.bd

The thesis titled "Automatic Detection of Diabetic Retinopathy from Fundus Images Using Image Processing and Artificial Neural Network", submitted by Ishtiyaque Ahmad, Roll No. **1015052042 P**, Session October 2015, to the Department of Computer Science and Engineering, Bangladesh University of Engineering and Technology, has been accepted as satisfactory in partial fulfillment of the requirements for the degree of Master of Science in Computer Science and Engineering and approved as to its style and contents. Examination held on January 26, 2019.

Board of Examiners

1. _____
Dr. A.K.M. Ashikur Rahman
Professor
Department of CSE, BUET, Dhaka.
Chairman
(Supervisor)

2. _____
Dr. Md. Mostofa Akbar
Head and Professor
Department of CSE, BUET, Dhaka.
Member
(Ex-Officio)

3. _____
Dr. Mohammed Eunus Ali
Professor
Department of CSE, BUET, Dhaka.
Member

4. _____
Dr. Rifat Shahriyar
Associate Professor
Department of CSE, BUET, Dhaka.
Member

5. _____
Dr. Mohammad Rashedur Rahman
Professor
Department of Electrical and Computer Engineering,
North South University, Dhaka-1229.
Member
(External)

Candidate's Declaration

This is hereby declared that the work titled "Automatic Detection of Diabetic Retinopathy from Fundus Images Using Image Processing and Artificial Neural Network" is the outcome of research carried out by me under the supervision of Dr. A.K.M. Ashikur Rahman, in the Department of Computer Science and Engineering, Bangladesh University of Engineering and Technology, Dhaka 1000. It is also declared that this thesis or any part of it has not been submitted elsewhere for the award of any degree or diploma.

Ishtiyaque Ahmad

Candidate

Acknowledgment

I express my heart-felt gratitude to my supervisor, Dr. A.K.M. Ashikur Rahman for his constant supervision of this work. He helped me a lot in every aspect of this work and guided me with proper directions whenever I sought one. His patient hearing of my ideas, critical analysis of my observations and detecting flaws (and amending thereby) in my thinking and writing have made this thesis a success.

I would also want to thank the members of my thesis committee for their valuable suggestions. I thank Dr. Md. Mostofa Akbar, Dr. Mohammed Eunus Ali, Dr. Rifat Shahriyar, and specially the external member Dr. Mohammad Rashedur Rahman.

In this regard, I remain ever grateful to my beloved parents, who always exist as sources of inspiration behind every success of mine I have ever made.

Abstract

Diabetic Retinopathy (DR), a disease that affects the retina, is a complication of diabetes and the fifth most common reason causing blindness and loss of vision over the world. The risk of severe damage to vision can be reduced by early screening and treatment. The absence of any visual symptom, lack of an adequate number of ophthalmologists, difficulties in the manual detection process, etc. are the key challenges for early detection of DR. An automated detection method would address these issues and facilitate the detection process. DR can be diagnosed by analyzing pictures of retina (known as fundus images) captured by special cameras. The severity of DR depends on the properties of some anatomical features present in a fundus image namely - optic disc, blood vessel, exudate, microaneurysm, and hemorrhage. This work has first developed techniques to identify each of the components separately from a fundus image. In the course of that, it has been able to overcome many limitations prevailing in state-of-the-art techniques. After that, a technique for detecting the severity of DR has been developed based on the identified components.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Diabetic Retinopathy (DR) is a complication of diabetes, caused by high blood sugar levels damaging the back of the eye (retina). If left untreated for several years, Diabetic Retinopathy may reach a stage causing blindness. In the year 2010, 0.8 million people worldwide were found to be blind and another 3.7 million visually impaired due to Diabetic Retinopathy [1]. Currently, at least 366 million people worldwide have diabetes and an estimated 41.6% of them are affected by some level of DR. Globally, the number of people with DR will reach up to 191 million by 2030 [2]. Early screening of DR and treatment may reduce the risk of vision loss. However, there is no significant visual symptom of Diabetic Retinopathy in the earlier stages. For this reason, potential patients are often reluctant to visit ophthalmologist at that stage. Ophthalmologists diagnose DR by taking pictures of retina (known as fundus images) using specialized cameras. There are mainly five anatomical components in a fundus image, identification of which are crucial for the detection of Diabetic Retinopathy - namely, optic disc, blood vessel, exudate, microaneurysm, and hemorrhage. A typical fundus image with these components labeled is shown in Figure 1.1. Three of these components - exudate, microaneurysm, and hemorrhage are actually considered as lesions and indicate abnormalities if present in a fundus image. Diagnosis is done based on different properties of these anatomical components.

There are several conventions prevailing in the medical community for grading Diabetic Retinopathy. Among them, the severity scale proposed by Wilkinson et. al. [3] is the most

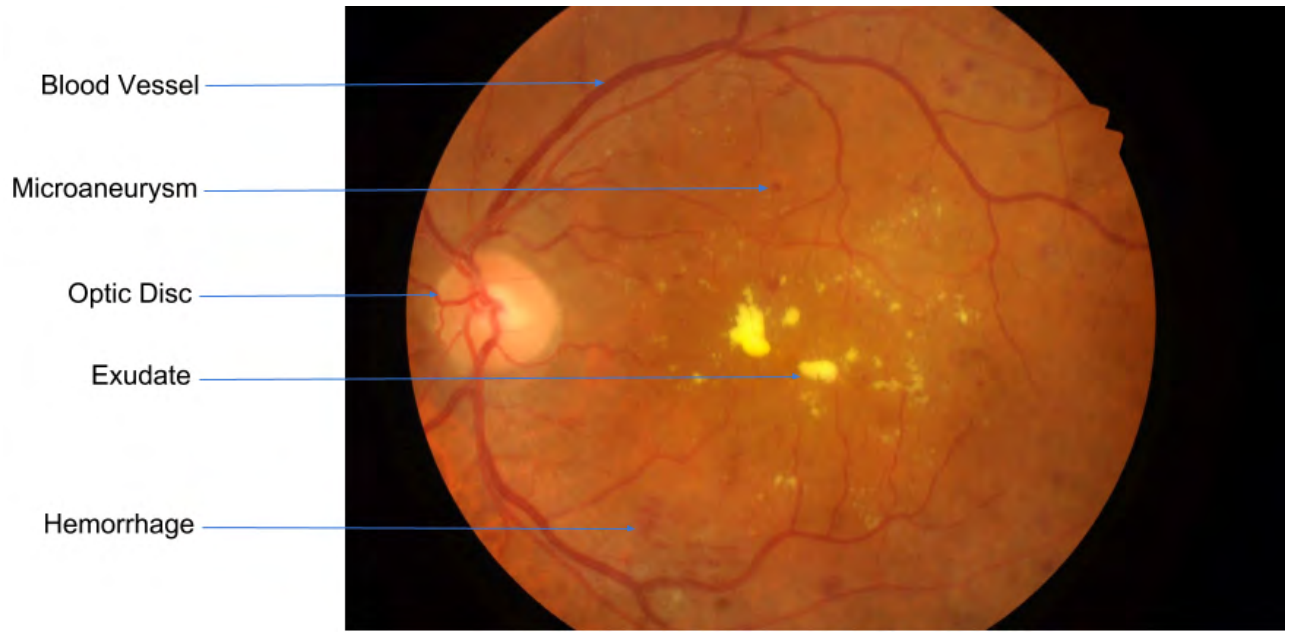

Figure 1.1: A typical fundus image labeled with different anatomical components.

widely used and therefore followed in this work. They have proposed five severity levels of DR. In ascending order of severity they are:

i No apparent retinopathy

ii Mild nonproliferative DR

iii Moderate nonproliferative DR

iv Severe nonproliferative DR

v Proliferative DR

This work first presents techniques developed for the identification of each of the five components from a fundus image. After that, a technique for determining the severity of DR according to the severity scale mentioned above is described. No existing work has been found in the literature that determines DR severity from fundus image according to Wilkinson's scale based on identification of components.

1.1 Motivation

During the period of 1990-2010, DR has been found to be the fifth most common cause of preventable or treatable blindness [2]. The treatment of DR needs to be started at an early stage to prevent severe damage to vision or blindness. An early detection is a prerequisite for that. Currently, there are too few ophthalmologists compared to the huge number of images that need to be diagnosed per year. Diagnosis is conducted by manual observation of images and identification of lesions. Unfortunately, this task is extremely labor intensive, time consuming, and prone to human error. An automated detection method would come to the aid of an ophthalmologist and thus increase the accuracy of detection. Moreover, it is found that large fraction of screened images bear no sign of retinopathy [4]. An automated detection method would screen these images out and allow the ophthalmologists to focus on more complicated cases. It has been found that, the prevalence of Diabetic Retinopathy is higher in people of South Asian, African, Latin American, and indigenous tribal descent [5]. A major fraction of these people are underprivileged and often unwilling to visit ophthalmologists due to ignorance and financial reasons. An automated detection process will make the procedure less cumbersome and thus contribute in the elevation of health care facilities for these people.

1.2 Contribution

[label=] The major contributions of this work can be enumerated as:

1. Techniques for identification of different anatomical components in a fundus image. While developing these techniques, care was taken to make them independent of image dataset and robust to variation in illumination, color, texture, and shape of the images. Although several works are available that describe techniques for identification of one or more of these components, those works are reported to operate on small datasets without significant diversity in the images.

2. A tool to provide comprehensive details regarding lesions present in an image. This tool provides information about the number of different lesion pixels present in a fundus image

as well as their location. It also calculates the diameter of the optic disc and the relative proximity of different lesions with respect to the optic disc. This set of information would aid an ophthalmologist for the detection of DR severity.

3. A novel technique for the detection of DR severity from the identified components. This would help building a full-fledged system that can identify DR severity from raw fundus images. Moreover, it would be possible to identify DR severity by manual inputs also.

1.3 Thesis Organization

This thesis is organized as follows. Chapter 1 gives a brief introduction to the problem under investigation. It also discusses the motivation to pursue the problem as well as the contributions of this work. Chapter 2 discusses the relevant researches on the problem. Chapter 3 first elaborately discusses the methodologies followed in this work to identify different components from a fundus image. After that, it describes the steps followed to detect DR severity from those identified components. Chapter 4 presents the results obtained and contributions made in this work. Finally, Chapter 5 provides some concluding remarks and possible future researches along this direction.

Chapter 2

Background

The approaches made for determining the severity of DR from fundus images can broadly be categorized into two types:

i Classification after component extraction

ii Clasification without component extraction

The former approach aims the detection by identifying different anatomical components in a fundus image. In contrast, the later approach performs the detection directly from raw fundus images without explicit identification of any component. This approach usually adopts deep learning techniques and requires a huge dataset of images. One key limitation of this approach is, it cannot provide insights for the detection made. As a result, further analysis of the condition as well as identification of misclassifications get difficult. This work detects DR following the first approach. A brief study of the existing works in the literature categorized according to the approach made be them is discussed as followed.

2.1 Classification After Component Extraction

There is no complete work in the literature that identifies all of the components and then makes the detection from fundus images. However, a good number of works for extracting one or more

of the components exist in the literature. Sinthanayothin et. al. [6] used automated image analysis systems to detect optic disc, blood vessels, and fovea. They proposed a new technique named Ḿoat operatorẃhich was used along with recursive region growing segmentation algorithms to detect hemorrhages and microaneurysms as one group, and hard exudates as another group. They applied the exudate recognition algorithm on 30 retinal images and attained a sensivity of 88.5% and a specifity of 99.7%. Applying on 14 retinal images, they achieved a sensitivity of 77.5% and specificity of 88.7% for detection of hemorrhages and microaneurysms. Yun et. al. [7] used binarization and morphological operations to detect perimeter and area of veins, hemorrhages and microaneurysms for each image channel. However, their work does not include identification of the anatomical components from the image. They trained a feed forward neural network with 77 images to classify them into 4 groups. They tested the model with 37 images and reported a specifity of more than 90%. [8] used region growing to determine the shape of the abnormalities and then classified the abnormalities using different classifiers. The best result was found with Mahalanobis classifier where microaneurysms, haemorrhages, exudates, and cotton wool spots were detected with a sensitivity of 69%, 83%, 99%, and 80%, respectively. [9] used trace transforms to extract features from fundus image.

Different techniques are found in existing literature for the detection and localization of optic disc. The most frequently used techniques include adaptive thresholding, identification the point of convergence of the main retinal blood vessels, active contour models, snakes, principal components analysis (PCA), and the watershed transform [10]. Sinthanayothin et al. [11] located the optic disc by identifying the area with the highest variation in intensity of adjacent pixels and obtained a sensitivity and specificity of 99.1%. However, this algorithm was reported to fail for images with a large number of white lesions, light artifacts or strongly visible choroidal vessels in [12]. Walter et. al. [13] detected optic disc by applying area threshold and then morphological operations. Their technique based on the assumptions that optic disc is the brightest part of a fundus image and exudates are never big enough to be compared with the size of optic disc. Ravishankar et. al. [14] found out the intersection point of the major blood vessels to approximate the location of optic disc. Their approach yield a 97.1% success rate for optic disc localization [15] used Hausdorff-based template matching and

pyramidal decomposition to detect optic disc in low resolution fundus image and reported an average error of 7%. [16] used three vessel distribution characteristics to find possible horizontal coordinate and global vessel direction characteristic to identify the vertical coordinate of optic disc. They reported to have a localization accuracy from 93.8% to 99.7%. [17] used firefly algorithm to detect optic disc.

Individual works for identification of blood vessels are also found in the literature. Use of both image processing and machine learning techniques are prevailing for this purpose. Welikela et. al. [18] divided the new vessels into two categories : new vessels at the optic disc and new vessels elsewhere. Then they used a genetic algorithm based approach for feature selection. Two dimensional match filter was used in [19], [20], [21]. Hoover et. al. [22] proposed an algorithm using two dimensional match filter and threshold probing and obtained 75% true positive rate over 20 fundus images. [23] used vessel segmentation method based on fully connected conditional random fields. Staal et. al. [24] extracted image ridges which coincide with vessel centerlines. Xu et. al. [25] detected large blood vessels by identifying large connected components in a binary image formed by thresholding. The remaining thin vessels were detected by an SVM classifier. They reported an average sensitivity of 77%. [26] constructed a 41-D feature vector and trained a feature-based AdaBoost classifier for detection of vessel and attained an accuracy of 95.97% over a test set of 20 images. [27] proposed a new clustering algorithm using a partial supervision strategy. [28], [29] used contour models. A hybrid approach was proposed by Yang et. al. [30] which consisted of mathematical morphology and a fuzzy clustering algorithm. A purification procedure was then used to reduce weak edge and noises. [31], [32] used Convolutional Neural Network (CNN) for vessel detection. Wang et. al.[33] extracted trainable features with the help of CNN and then used Random Forest for classification based on the features.

In most of the cases, first the optic disc is detected and masked for exudate detection. Earlier works on the detection of exudate mostly used thresholding and morphological image processing techniques. Walter et. al. [34] first detected the optic disc using area threshold in the green channel and then identified exudates by using their high grey level variation. They determined exudate contours using morphological reconstruction techniques. [35] used domain

knowledge to differentiate exudates from other bright regions generated due to reflection of light in the image. [36] extracted bright candidate lesions from fundus image and then applied a curvelet-based algorithm for detection of optic disc and exudates. [37] applied average filtering and contrast adjustment for preprocessing fundus images and then used marker controlled watershed algorithm for extracting contours of optic disc and exudates. A model based approach by combining region growing and edge detection was used in [38]. Some approaches have adopted techniques related to artificial intelligence to differentiate exudates from other bright regions in fundus image. Osareh et. al. [39] segmented color retinal images using fuzzy c-means clustering. Features were extracted from the segments and a multilayer neural network classifier was used to classify each segment as exudate or non-exudate based on those features. They used 300 manually labeled retinal images for their algorithm and attained a sensitivity of 93.5%. [40] used contextual clustering algorithm to segment the retinal image. The segment with the greatest area was considered as optic disc and then eliminated. Three features : convex area, solidity, and orientation were extracted from each segment and a modular neural network was trained with the feature vector to detect exudates. Some recent works have used deep learning for the detection of exudates from fundus images. [41] first identified optic disc and blood vessels using existing methods and then used deep convolutional neural network for classifying each pixel as exudate and non-exudate. They reported an F-score of 0.78. [42] divided the image into patches and classified each patch into exudate or non-exudate using deep learning.

Due to similarity in color and texture, microaneurysms and hemorrhages are detected using similar techniques in many of the methods described in existing literature. Both are usually dark red in color and circular or elliptical in shape. Due to this consistency in shape, image processing methods popular for detecting shapes can bes used to detect them. Fleming et. al. [43] used linear structuring element from eight different directions after local contrast normalization. Regions only having circular or near circular regions got enhanced by all the structuring elements from different directions. Dots within vessels initially wrongly detected as microaneurysms were removed using local vessel detectionn. They reported a sensitivity of 85.4% and specificity of 83.1% for microaneurysm detection. Mizutani et. al. [44] selected

candidates for microaneurysm using double ring filter and then classified them using rule based classifier and artificial neural network. Circular Hough transform was used to detect microaneurysm candidates in [45] after blood vessels were removed from the image. Filter banks were used in [46]. Sopharak et. al. [47] detected microaneurysm using mathematical morphology methods. They considered circular patterns having diameter less than 125 μm, a constant intensity, and boundary pixels with higher value in the green plane. Some recent works have used artificial intelligence techniques to detect microaneurysms and hemorrhages. Tang et. al. [48] partitioned the fundus image into non-overlapping segments containing similar color and spatial location. Features were extracted from each of the segments and then a classifier was used for hemorrhage detection. [49] used a hybrid method consisting of both mathematical morphology and pixel classification for selecting red lesion candidates. They extracted 21 different features from the candidate objects. A kNN classifier was used for classification of the candidate objects into red lesion or not. Grinsven et. al. [50] used Convolutional Neural Network (CNN) to identify micoaneurysm and hemorrhages; they achieved area under the receiver operating characteristics (ROC) curve of 0.894 and 0.972 on two data sets.

2.2 Clasification without component extraction

Gulshan et. al.[51] used deep learning to detect Diabetic Retinopathy directly from fundus images without explicitly detecting any lesion. They used the Inception-v3 model[52] - consisting of 42 layers - which is known to perform very well for image classification. An ensemble of 10 networks each containing 22 million parameters were trained with 128,175 images and a linear average of the predictions of the ensembles was considered as the final prediction. They validated the prediction model with two data sets containing 9963 and 1748 images respectively. They achieved 90.3% and 87.0% sensitivity and 98.1% and 98.5% specificity respectively at the operating point selected for high specificity. On the other hand, they achieved sensitivity of 97.5% and 96.1%; and specificity of 93.4% and 93.9% at the operating point selected for high sensitivity.

Chapter 3

Methodology

In this work, the severity of DR is detected from different properties of the anatomical components present in a fundus image. Therefore, identification of all those components is a pre-requisite for the DR detection task. For this reason, five anatomical components as shown in Figure 1.1 were first identified. They are:

i Optic disc

ii Blood vessel

iii Exudates

iv Microaneurysms and Hemorrhages (Dark Lesions)

The identification procedures for each of the components followed in this work are not absolutely independent of each other. For instance, in order to identify the exudates, the optic disc needs to be identified first and then masked out. Similarly, due to a similarity in color, detection of dark lesions require that the blood vessels are identified and masked out. However, there is no dependency between the identification of optic disc and blood vessels. As a result, the entire component identification part can be thought of as two parallel sequences of tasks. In one sequence, the optic disc is identified followed by the identification of exudates whereas the other sequence consists of identification of blood vessels followed by dark lesions.

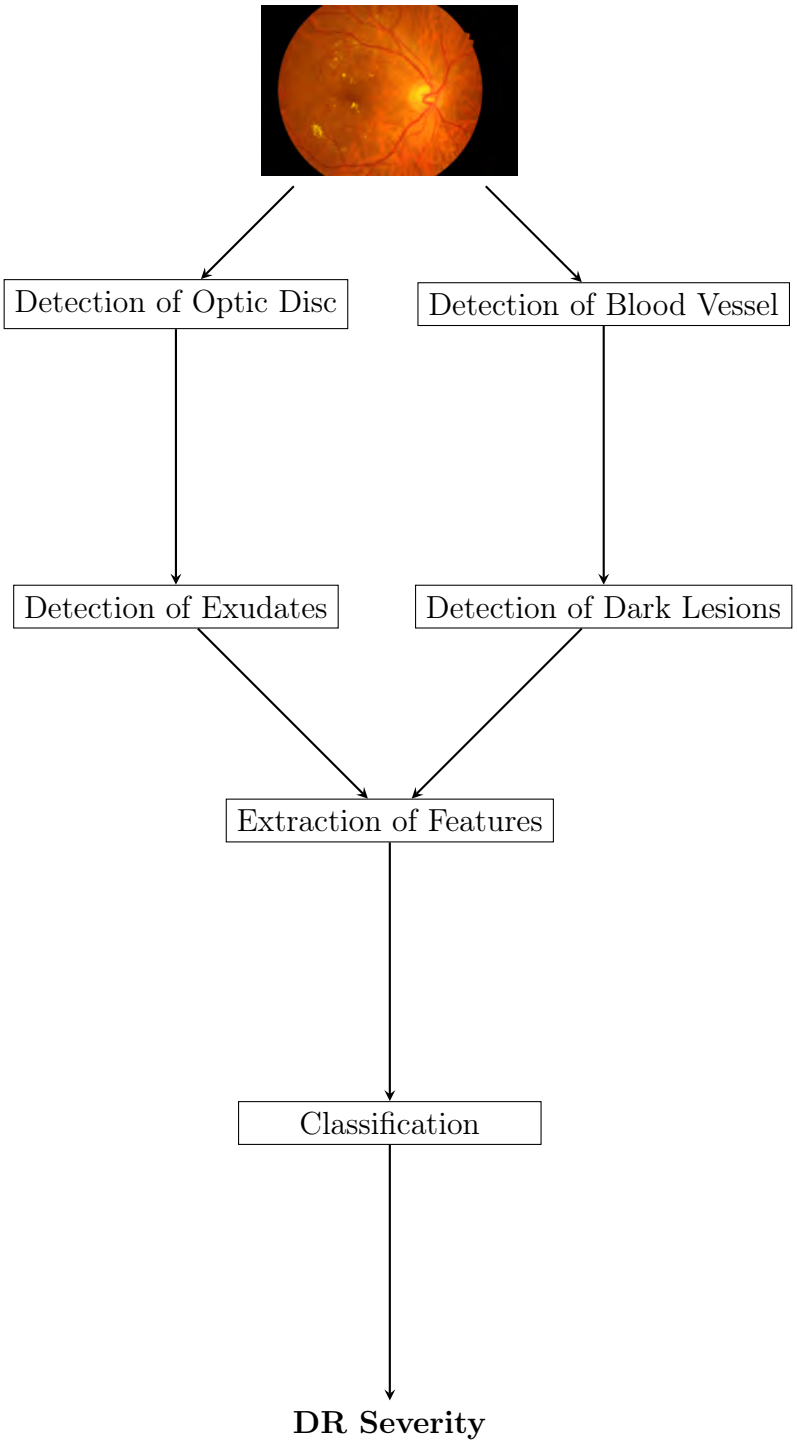

Figure 3.1: Flow Diagram for DR Severity Detection

After having all the components identified, a total of 14 features related to their presence and location were extracted. These features were then used to determine the DR severity level from a particular fundus image. A flow diagram of the steps followed in this work for the detection of DR severity from a fundus image is shown in Figure 3.1. These steps are elaborately discussed below.

3.1 Detection of Optic Disc

Optic disc is the oval shaped bright region in a fundus image where all the blood vessels join together. It can be positioned anywhere in a fundus image. Moreover, the color and texture of an optic disc may vary based on the illumination of an image. This work first discusses some possible approaches made for optic disc detection which are also extant in the literature. The limitations and difficulties associated with each approach are also discussed along with. After that, the technique for optic disc detection proposed by this work as well as its advantages over other techniques are discussed.

3.1.1 Possible Approaches and Their Limitations

Due to its visually prominent characteristics from rest of the fundus image, an approach was made to identify the optic disc by using image segmentation based on intensity values of the pixels. In stead of applying segmentation on a color fundus image, the green channel of an RGB image was chosen to apply segmentation. One advantage of this choice was the cheaper computational cost of applying the same operations on a single channel image rather than on a three channel image. The rationale for choosing green channel over red and blue channel is that, the optic disc as well as other image components are visually more prominent in the green channel than blue or red channels. A comparison of the different channels of a fundus image is shown in figure 3.2.

The green channel was first preprocessed by applying mean filtering and histogram equalization. The resultant image was then applied a piecewise linear transformation so that the pixel value range becomes the maximum and same for all images. Then, thresholding and

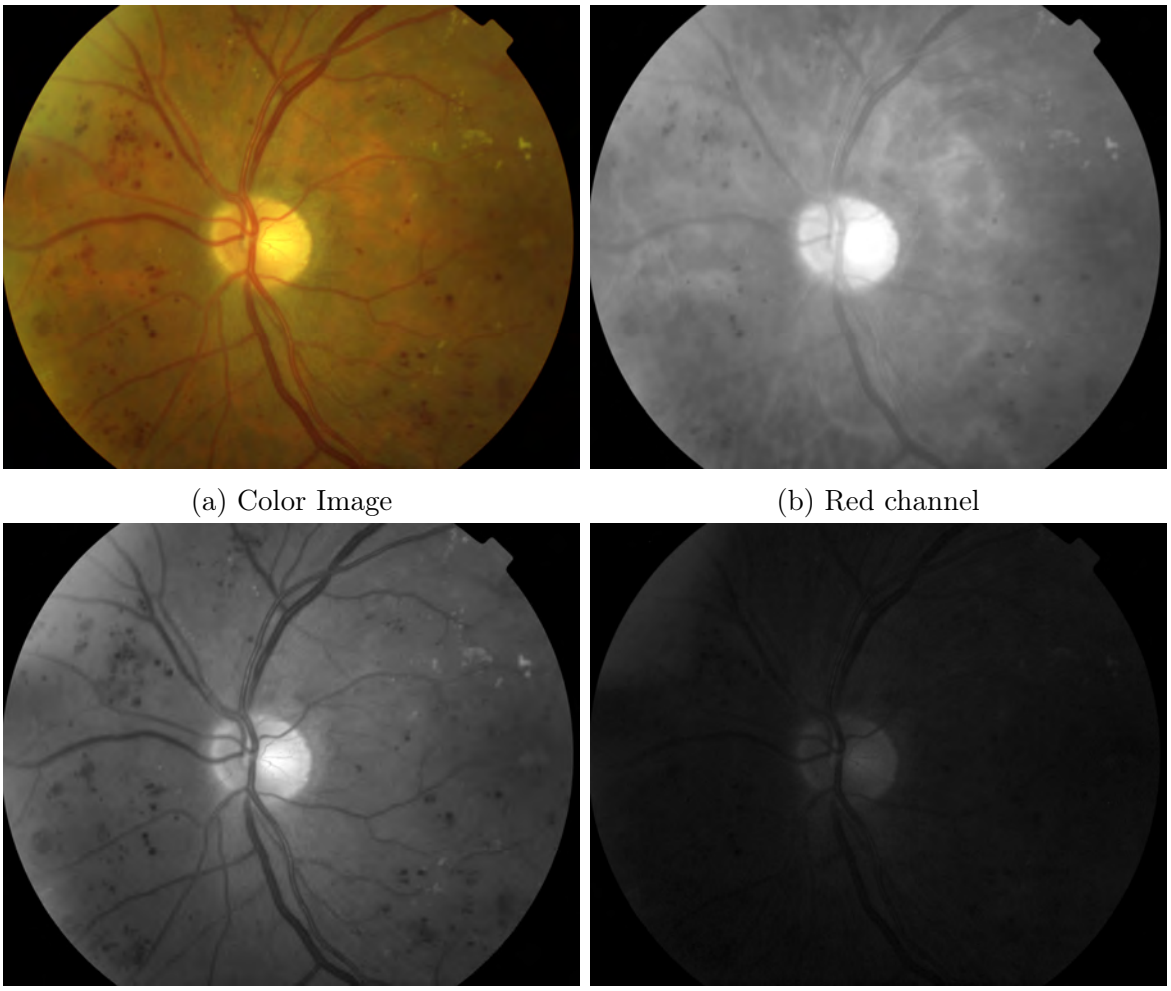

(a) Color Image (b) Red channel

(c) Green channel (d) Blue channel

Figure 3.2: A color fundus image and its red, blue, and green channels.

segmentation were performed to separate the optic disc from the green image. However, this method came out to have limitations as shown in figure 3.3. In figure 3.3a another bright circular region apart from the optic disc is present in the image which is in fact a lesion rather than optic disc. Nonetheless, this region was also detected as optic disc by the segmentation method. In addition to that, in figure 3.3b, the optic disc is obscured by lesions and are not as visually distinguishable like the ones in most fundus images. As a result, the segmentation method could not identify the optic disc for this image as well. Notwithstanding the reasonably good performance of image segmentation method to identify the optic disc, this method was found not to be robust enough to the diversities prevailing in fundus images.

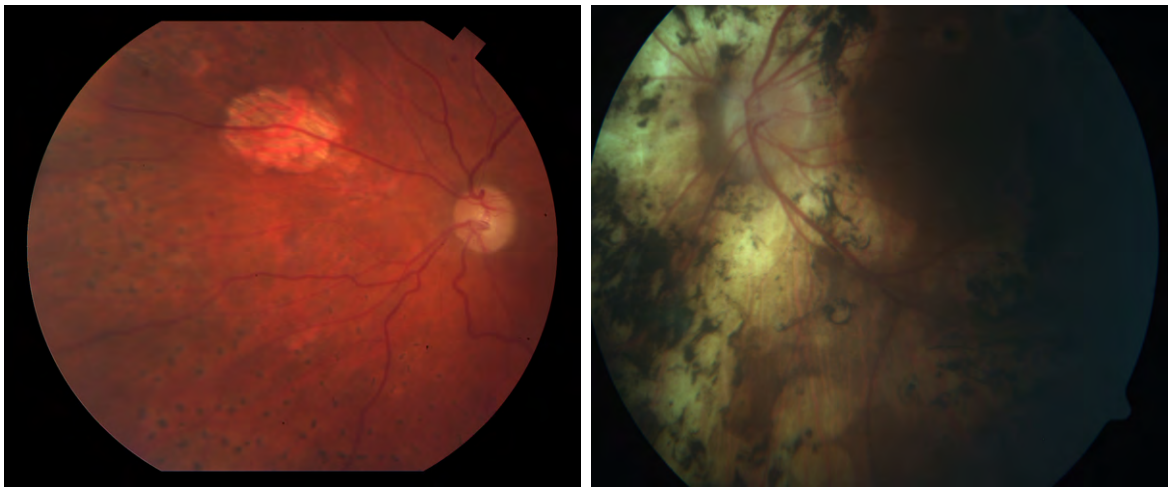

(a) incorrect detection (b) unable to detect

Figure 3.3: Optic disc identified incorrectly by segmentation method

The next approach was to divide the fundus image into small rectangular segments and then classify that segment based on whether it contains an optic disc or not. For this approach, image segments were cropped from sample fundus images. From 1200 sample images, one image segment was constructed by cropping the optic disc only. These were considered as positive samples. In addition, 1200 image segments were cropped which were not optic discs. These segments were considered as negative samples. Many of the negative samples were chosen to have lesions which look similar to the optic disc. Some sample positive and negative image segments are shown in figure 3.4. The obtained image segment set was then divided into two sets: a *training set* containing 1000 positive images and 1000 negative images; and a *testing set* containing 200 positive images and 200 negative images. At first, a deep convolutional neural network (CNN) was trained using the training dataset and then tested using the testing set. Training the network for several times after changing network hyperparameters, the best training accuracy obtained was 62.5% and test accuracy was 56.3%. It could be inferred that the training dataset was too small for the network. There were two possible solutions to this probelm: to increase the size of the dataset, and to use a technique named transfer learning which uses a network trained previously with another larger datset. Since a a larger dataset was not available, the transfer learning approach was adopted to improve the performance. Simonyan et. al. developed a convolutional neural network based image classifier named

VGG16 [53] which was trained using more than 1 million images. This trained model was used for transfer learning.

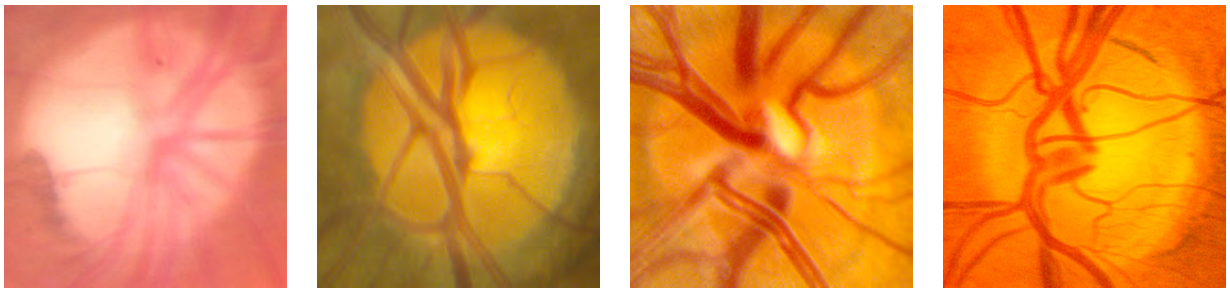

(a) Positive samples

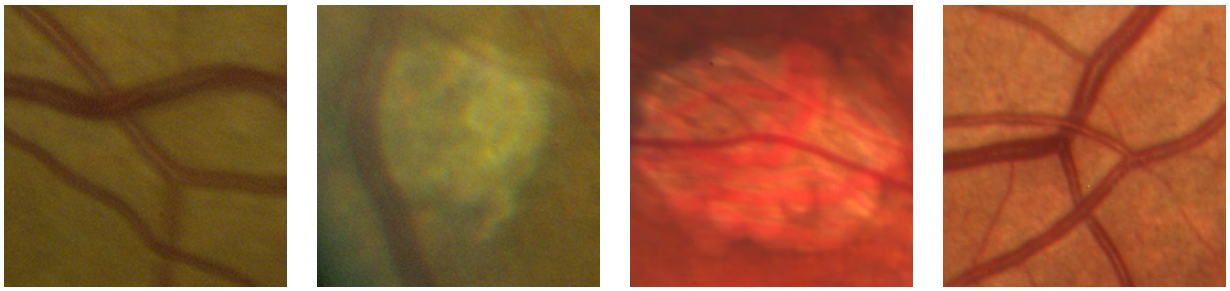

(b) Negative samples

Figure 3.4: Sample images from the dataset created by cropping fundus images for detection of optic disc.

The main limitation of the classification method was its inability to identify the location of an optic disc within reasonable time frame. Scanning the entire image sequentially with different window sizes required a long time. So, identifying the location of the optic disc faster in a fundus image became necessary.

3.1.2 Proposed Approach

The proposed method in this work identifies optic disc following two sequential steps: i) Localization, ii) Segmentation. The localization step actually identifies the region where the optic disc is located by drawing a bounding box around it. Then, the segmentation step identifies the exact boundary of the optic disc. These two steps are discussed as followed.

Localization

The localization step considered the optic disc as an object on the fundus image background and an approach was made to localize the object. Several works have been done in computer vision domain for solving the object localization problem. However, application of object localization technique in order to detect optic disc in a fundus image has not been tried yet and this work is the first one to do so. For this purpose, a state-of-the art object localization system named "You Only Look Once" (YOLO)[56] has been used. YOLO can draw a bounding box around different objects in an input image given that it was trained with those objects. YOLO uses a deep neural network consisting of 24 convolutional layers followed by 2 fully connected layers for object localization. The architecture of YOLO is depicted in Figure 3.5.

In order to make YOLO capable of localizing optic disc, the deep neural network was required to be trained for that. For the purpose of this training, a dataset containing 1250 fundus images obtained from Ubicomp Lab [54] was annotated with the location of optic disc in each image. The annotation contained the location of the top left corner point, width, and height of a bounding rectangle for the optic disc in an image. The constructed dataset has been made public [55] to facilitate reproduction and further research. However, the size of the dataset is not large enough to train YOLO from scratch. Hence, transfer learning was used where only fully connected layers of YOLO were re-trained. The training was conducted using Darknet [58] framework for 3000 epochs obtaining a training accuracy of 99.6% which took around seven days on an Nvidia GeForce 940MX GPU. The trained model could localize the optic disc in a fundus image by drawing a rectangular bounding box around it as shown in Figure 3.6b.

Segmentation

After identifying a bounding box around the optic disc, the next goal was to determine its exact boundary. This was done by applying local thresholding inside the bounding box obtained in the localization step. In contrast to thresholding the entire image, which was found to have limitations discussed earlier, local thresholding attempted to identify the brightest region inside

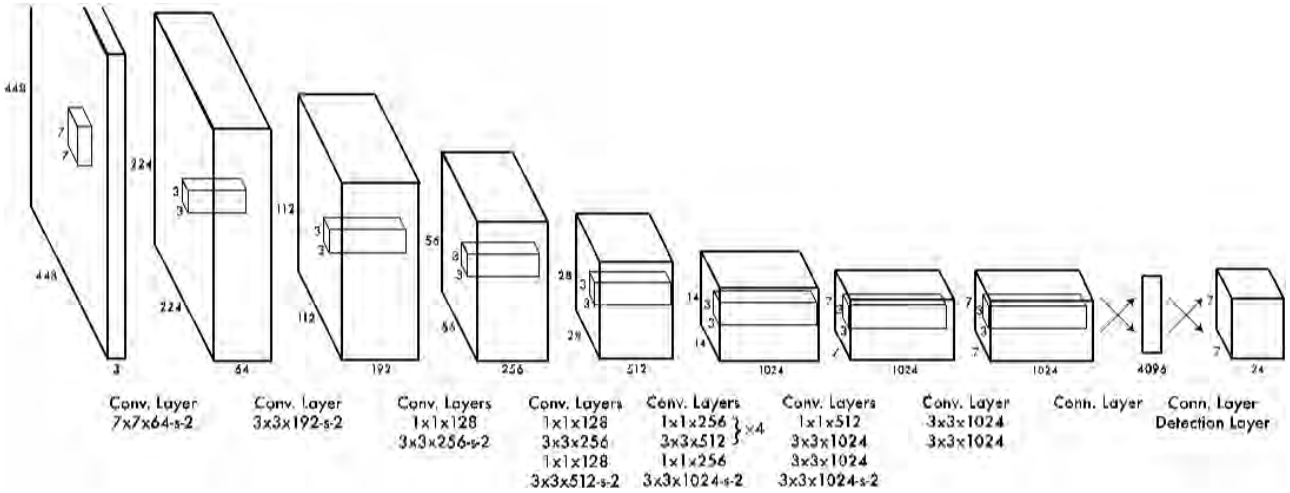

Figure 3.5: The architecture of Deep Neural Network YOLO used for optic disc localization.

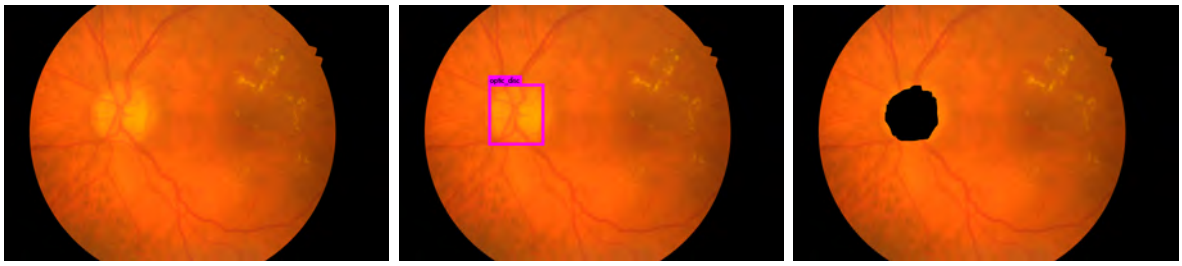

(a) A sample fundus image from IDRiD data set for optic disc detection

(b) Output of optic disc localization model, a bounding box is drawn around the optic disc.

(c) After segmentation, optic disc pixels are marked with black.

Figure 3.6: Output of different stages of optic disc detection.

the bounding box. The threshold value for each image was determined using Otsu's method [59] which is a well known algorithm for calculating threshold value for bimodal image. Pixels inside the bounding box having value greater than the threshold were considered belonging to the optic disc. The collection of all such pixels for a particular image was identified as the optic disc. Figure 3.6c shows the output of segmentation operation. The optic disc pixels are masked out by assigning a zero value for better visualization. The performance of the optic disc detection technique on unobserved test images is discussed in Section 4.

3.2 Detection of Blood Vessel

Blood vessels form a connected network like architecture in a fundus image which is often called as vessel network. The main vessels originate from the optic disc and are thicker in size. Branches and sub-branches originate from different points of the main vessel and thus construct a well-spread network. This work used Recurrent Neural Network (RNN) for the detection of blood vessels. Different steps of the detection methodology are described as followed.

3.2.1 Data Preparation

Two publicly available datasets: STARE [60], and DRIVE [61] were used for the detection of blood vessel. The STARE dataset contains 20 color fundus images. Each fundus image is accompanied by a binary grayscale label image of the same width and height. The label image contains 1 at pixel positions where there were blood vessels in the fundus image, and 0 otherwise. Similarly, the DRIVE dataset contains 40 color fundus images and 40 label images. These two datasets were combined to get a total of 60 labeled images which were then used for blood vessel detection. A sample image from the dataset along with its label is shown in Figure 3.7.

At first, the green channel of the images were extracted to form a single channel gray-scale image. Then, the obtained gray-scale image was resized to have a width of 500 pixels keeping the aspect ratio the same as original image. Bi-cubic interpolation was applied while resizing the images. It was observed that, in most of the images, the illumination all over the image was not uniform. As a result, the pixel values for similar components varied largely. Contrast enhancement was performed on the images using Speeded-Up Adaptive Contrast Enhancement (SUACE)[62]. The enhanced images were then randomly divided into two sets: the training set consisting of 48 labeled images and the test set consisting of the remaining 12 labeled images.

Each image was split into overlapping patches of size 50 pixels by 50 pixels starting from the top left corner of the image with a stride length of 10 pixels both along width and height of the image. The image was padded with zero values for boundary patches. The labels for each image was cropped in a similar way. Figure 3.8 shows non-overlapping patches in an

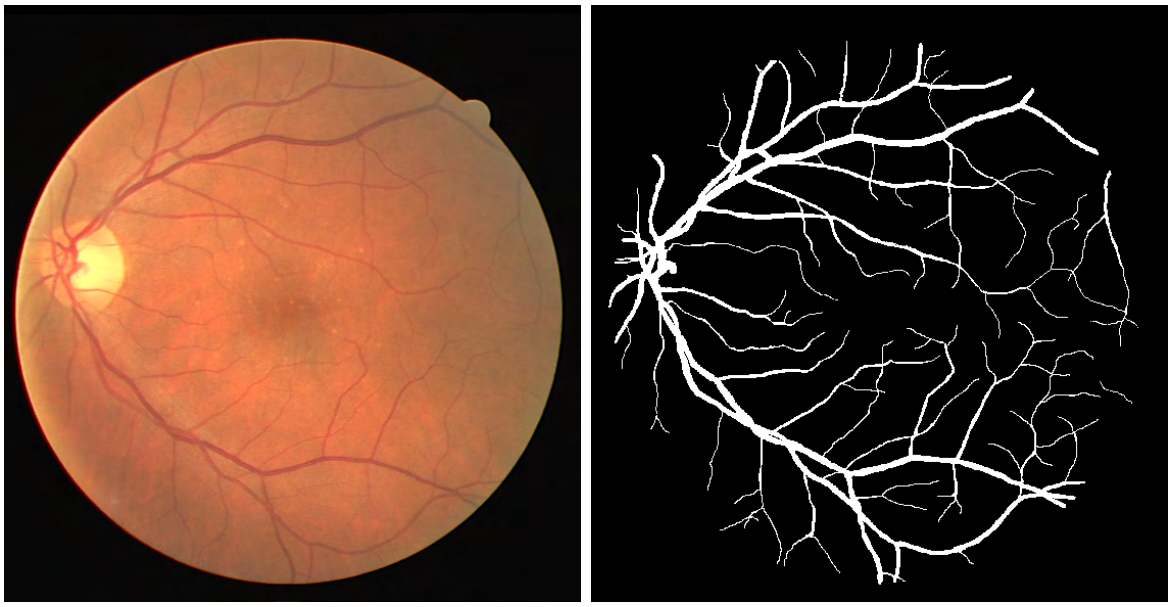

(a) Sample fundus image (b) Blood vessel labeling

Figure 3.7: A sample image along with corresponding label image used for blood vessel detection.

enhanced image. Overlapping patches can be obtained in a similar way starting from a stride away from the starting point an so on. Each patch was then reshaped to a 1 dimensional array of length 2500. Each image patch was paired with its corresponding label patch and a dataset was constructed containing all such pairs obtained from the images. These arrays were then normalized and used as inputs for our training model.

3.2.2 Model Construction and Training

Recurrent Neural Networks are best known for text sequence prediction. This work models the vessel detection problem as a sequence detection problem and addresses it by applying RNN. It was observed that, whether a pixel in a fundus image belongs to blood vessel or not depends not only on the value of that particular pixel, but also on the fact whether it is a part of a sequence of vessel pixels. Often pixels belonging to microaneurysms, and hemorrhages have similar pixel value as that of a pixel in a blood vessel. The distinction between such pixels and blood vessels can be made by looking at the nearby pixels. If the pixel under consideration is the part of an elongated tubular structure, it should be a vessel pixel. In contrast, if the

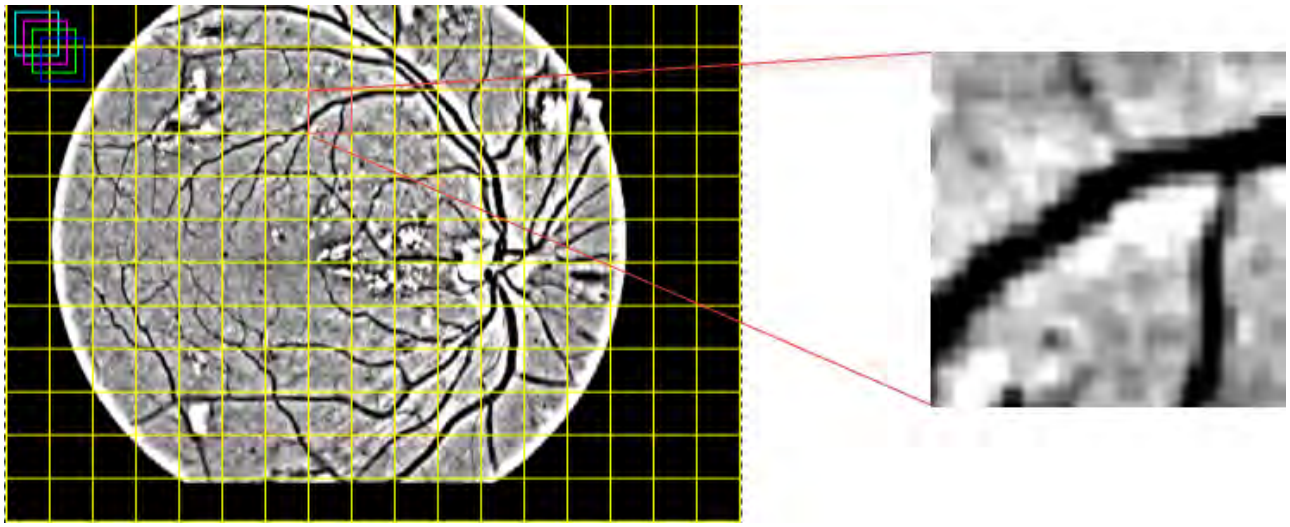

Figure 3.8: A contrast enhanced image is divided into patches of size 50 by 50. The yellow grid represents the patches obtained by starting from the top left corner of the image. Changing the starting point will generate overlapping patches as shown near the top left corner. A scaled up view of one such patch is shown alongside.

pixel belongs to an isolated dark region, it is more likely to be a part of microaneurysm or hemorrhage. Therefore, each patch constructed from the fundus images was considered as a sequence of pixels. The corresponding label patch would then also be a sequence of pixel denoting the label sequence. An RNN was constructed with two bidirectional LSTM layers using the deep learning framework Keras[63] with theano[64] as backend. The loss function used was binary crossentropy and optimization algorithm used was rmsprop. The model was trained with 80% of the training dataset created as mentioned earlier and the remaining 20% was used for cross-validation. Number of nodes in each layer was fine-tuned experimentally. A model with 100 nodes in the first LSTM layer and 50 nodes in the second LSTM layer produced the minimum loss after experimentation with different values. The architecture of the RNN model used for blood vessel detection is depicted in Figure 3.9.

The model was then trained for 38 epochs on a Tesla K80 GPU in Floydhub which took around 120 hours to complete. The loss curve demonstrating the change of training loss and validation loss with respect to training epochs is shown in Figure 3.10. The obtained minimum validation loss after 38 epochs was 0.12. The model weights generating this minimum validation loss was used for further detection from unobserved images.

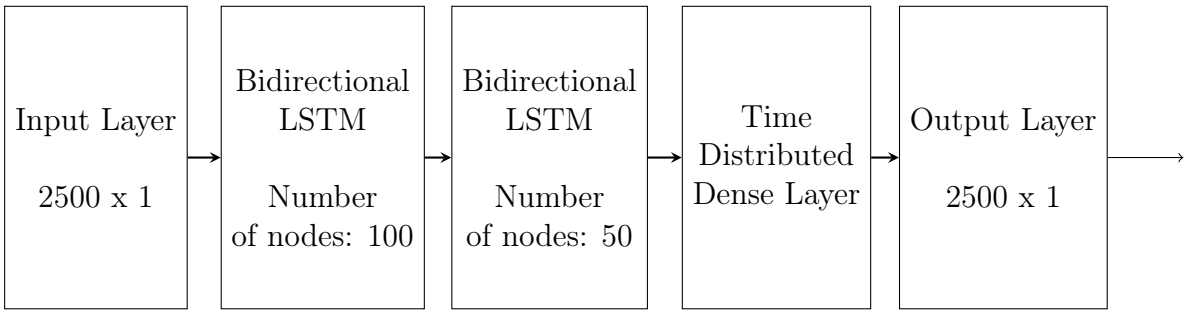

Figure 3.9: Architecture of the recurrent neural network used for detection of blood vessels.

3.2.3 Detection using Model

The trained RNN model was used for the detection of blood vessels in a test fundus image. For this, the test image was first resized to have a width of 500 pixels keeping the aspect ratio unchanged. The green channel was then extracted and contrast enhancement was performed. Next goal was to divide the image into non-overlapping patches of size 50 pixels by 50 pixels. However, there was a chance that the height of the image might not be divisible by 50. For that, rows padded with zeros were augmented with the image. If h was the original height of the image, then $\{(\lceil \frac{h}{50} \rceil * 50) - h\}$ rows were augmented with the image. That enabled dividing the image into non-overlapping patches of size 50 pixels by 50 pixels. Each patch was then reshaped to a 1 dimensional array of length 2500. The single dimensional arrays were then normalized and used as input for prediction. The constructed model would give an output of size 2500 containing the probability of each corresponding pixel in the input to be a vessel pixel. The output array was then reshaped to a two dimensional array of size 50 by 50. This output patch was considered as the output label predicted by the model for that particular input patch. In this way, output label patches were predicted by the model for all the input patches generated from a particular test image. The output label patches were then combined to construct a single output label image. After that, the number of extra rows augmented as mentioned earlier were removed from the label image to make its shape equal to that of the test image before augmentation. A sample test image and its corresponding label image formed in this way is shown in Figure 3.11b

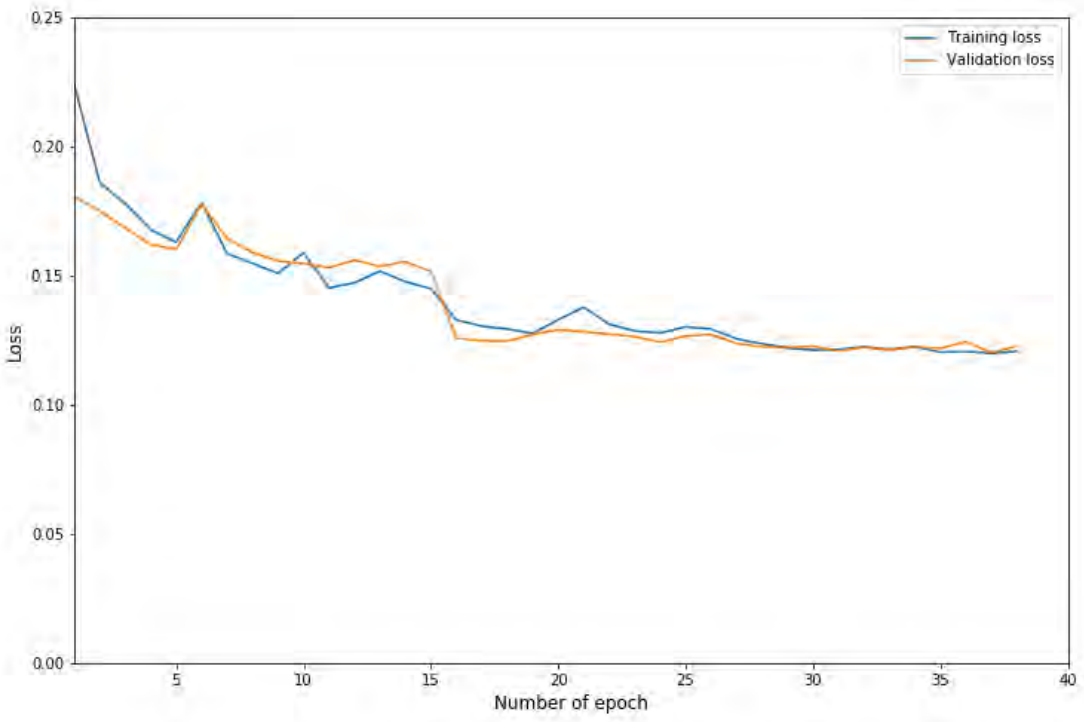

Figure 3.10: Loss curve for blood vessel detection model

3.2.4 Postprocessing

All the pixels in the output probability image I generated by the RNN model contained a double value within zero and one. This value denoted the probability of a pixel being a vessel pixel. However, in order to obtain a binary label image identifying the blood vessel in the fundus image, some post-processing on the probability image were performed. The post-processing steps are briefly discussed below in the same order of execution.

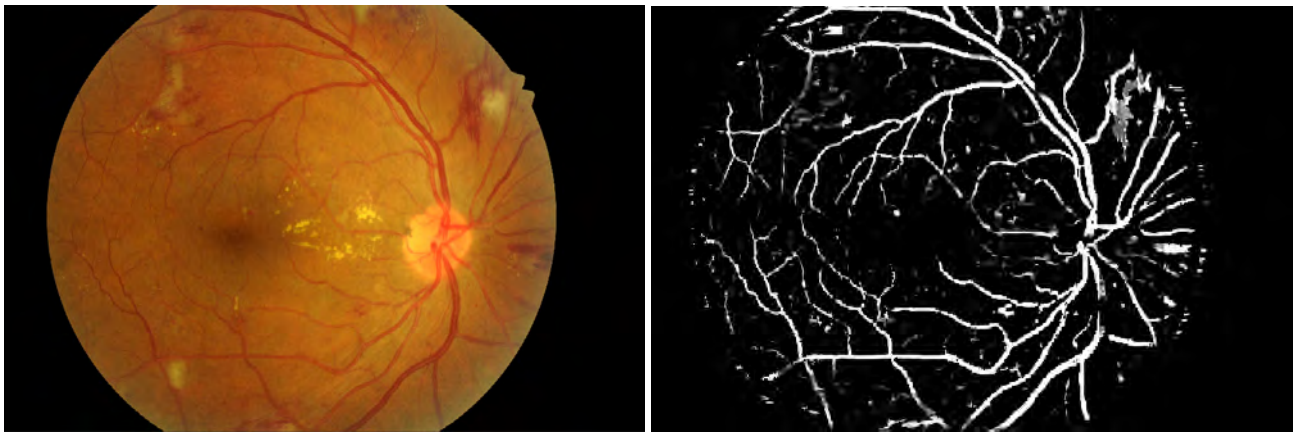

(a) A test fundus image. (b) Probability label image.

Figure 3.11: A test fundus image and the corresponding probability label of blood vessel generated by RNN model.

Thresholding

The probability image was thresholded at a value t to retain the predictions with a high probability only. A binary grayscale image I_t was constructed where,

$$I_t(x,y) = \begin{cases} 1, & \text{if } I(x,y) \geq t \\ 0, & \text{otherwise} \end{cases}$$

Experiments were conducted with different values of t for doing the threshold. The impact of t on the performance of vessel detection is discussed in Section 4. Figure 3.12b shows the output generated after thresholding the probability image in 3.12a at a value t=0.6. Pixels with value 1 in I_t represents presence of blood vessel at that corresponding pixel of the test image.

Noise Removal

Though the thresholding step could remove almost all of the non-vessel pixels from the probability image, some isolated noises were still remaining. These noises were in the form of small regions in the thresholded image which were incorrectly assigned a probability higher than t by the RNN model. A morphological opening operation with an elliptical structuring element was

performed. As a result, isolated noises with elliptical or circular shape were removed. Noises with other shapes were identified by calculating the connected components in the image. The components having size less than 100 pixels were identified as non-vessel noise regions and therefore were discarded. A noise reduced label image generated by this process is shown in Figure 3.12c.

Bridging Gaps

Besides removing non-vessel pixels, the thresholding and noise removal steps also removed some vessel pixels. As a result, the blood vessels got discontinued at some points as shown by red circles in Figure 3.12c. In order to bridge these gaps, a flood-filling operation was performed on the test fundus image. Flood filling operation takes some seed regions and include any adjacent pixel having similar color into the region. This process goes on until no neighboring pixel that can be included is available. For bridging the gaps between identified vessels, seeds were identified from the fundus image. Those pixels in the fundus image for which the corresponding pixels in the obtained binary label image are marked as vessel were considered as seeds. Then by flood filling, the nearby similar colored pixels were also included in the vessel network and thus bridging the undesired gaps. The resultant image produced after applying flood-filling operation on Figure 3.12c is shown in Figure 3.12d. This represents the final blood vessel detection for a test image by the described method in this work.

The performance of the vessel detection method over all of the 20 test images is discussed elaborately in Section 4.

3.3 Detection of Exudates

Exudates are bright colored lesions in a fundus image. They were detected using three steps. First, regions in a fundus image which are comparatively brighter than their surroundings and have some other distinguishing properties were marked as Regions of Interest (ROI). Second, 14 different features were extracted from each region. Finally, the regions of interest were classified as to be exudate or non-exudate based on the features using a machine learning

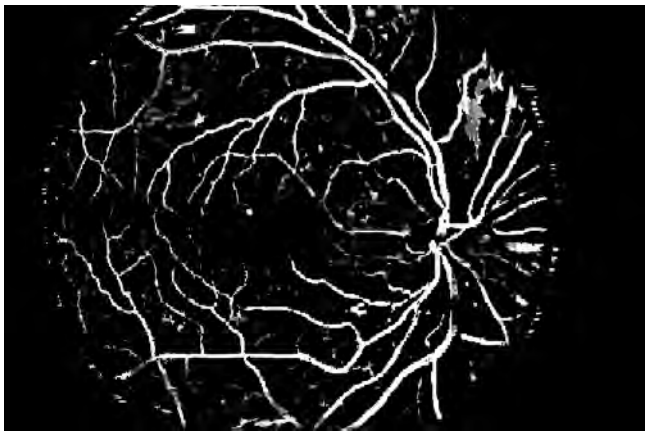

(a) Output probability image by RNN.

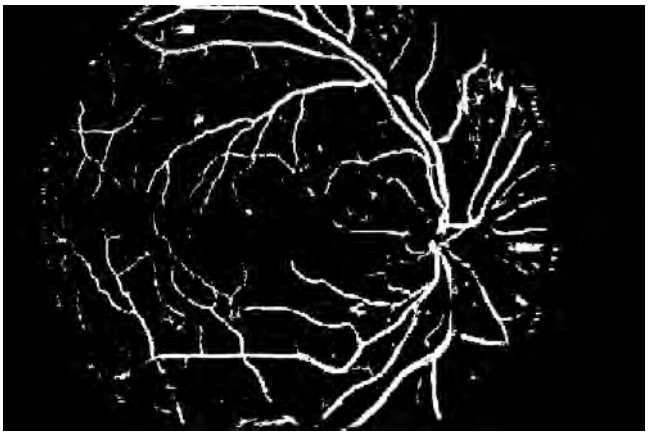

(b) Probability image thresholded at value $t = 0.6$

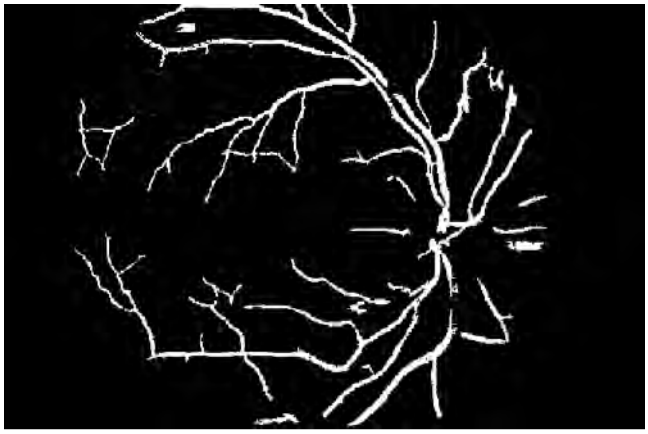

(c) Label image after reduction of noise. Discontinuities within vessels are marked with red circles.

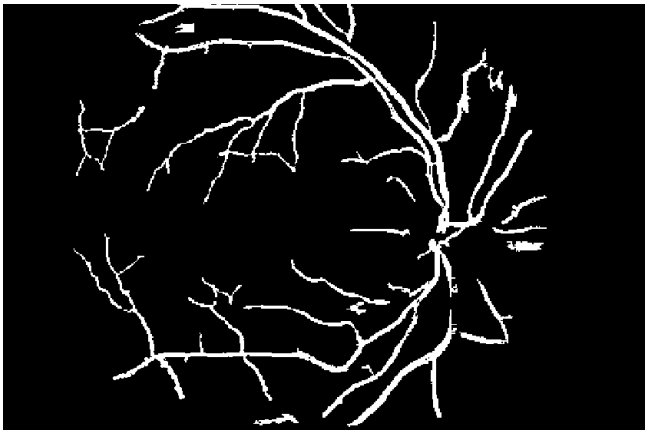

(d) Label image after bridging discontinuities within vessels by flood-filling operation.

Figure 3.12: Output produced after the completion of each intermediate step during blood vessel detection.

classifier. These three steps are elaborately discussed below.

3.3.1 Extraction of Regions of Interest

An exudate is always brighter than its surrounding region in a fundus image. However, due to non-uniformity of illumination over the image, some non-exudate bright colored patches also arise. Therefore, all regions that are brighter than its surroundings were considered as regions of interest (ROI) for exudate detection. The non-exudate regions identified as ROI were later eliminated in the classification step. The optic disc was first detected using methods mentioned in Section 3.1 and masked out with zero values to avoid it to be falsely identified as a region

of interest. It was observed that, the range of values spanned by all the pixels in a fundus image depends on the illumination of the image taken. However, the deviation of values among neighboring pixels is not the same everywhere. For example, the difference of pixel values on the same exudate region is significantly smaller than the difference of any particular pixel on the exudate and any other pixel in the neighboring non-exudate background region. To reduce the number of distinct colors and yet maintain the color difference between different regions, the fundus image was first represented with a smaller number of colors by applying quantization. This quantization was performed by dividing the image pixels into a number of clusters by applying k-means clustering algorithm on the image pixels. This algorithm distributed all the pixles in a color fundus image into k distinct clusters. Pixels having similar color values were grouped into the same cluster while dissimilar pixels were grouped into different clusters. Color similarity between two pixels was calculated by converting the image into CIELAB color space and then taking Euclidean distance between the pixels. CIELAB color space is known to have better similarity with human visual perception than RGB color space for such type of color similarity measured by Euclidean distance. The less the Euclidean distance between two pixels is, the more similar the pixels are.

Different values of k were tried starting from k=4 with an increment by one up to k=10. By empirical analysis, it was found that k=8 produces the best result for clustering. Lower values of k fail to distinguish between exudates and other non-exudate bright regions. On the other hand, higher values of k result in multiple clusters in the same region. The image was converted back to RGB color space after quantization. A sample fundus image and its quantized form for k=8 is demonstrated in figure 3.13.

After quantization, the green channel of the quantized image was considered for extraction of ROI. This is because the difference of pixel values is more prominent in the green channel than others. Since the image was quantized to contain exactly 8 distinct values, all the pixels in the green channel belonged to a set of values $\{v_1, v_2, v_3, \ldots, v_8\}$ sorted in ascending order. It was observed that, exudate pixels must have value greater than or equal to v_3 since exudates are brighter than blood vessels and the background without exception. All the pixels having the same color value and having a connected path between themselves were considered as a

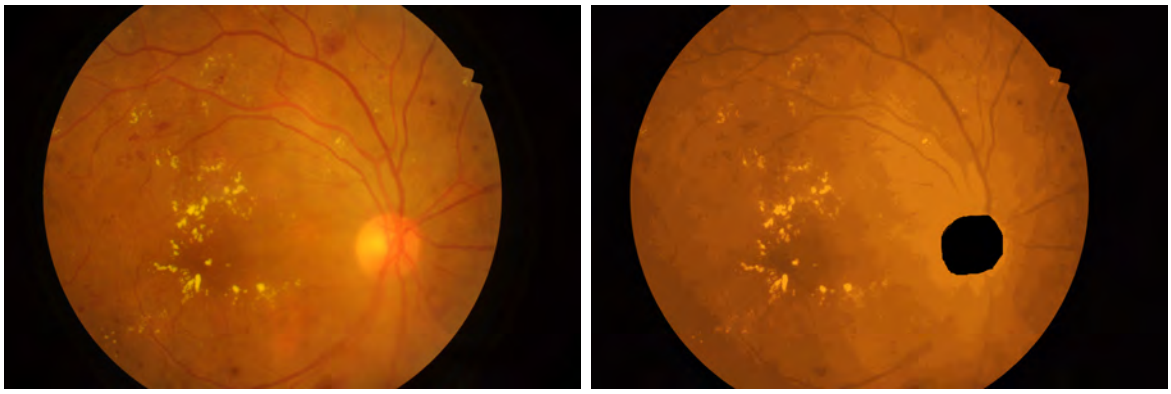

(a) Sample fundus image (b) After quantization

Figure 3.13: A sample fundus image and the corresponding resultant image after masking out the optic disc and applying quantization into 8 levels.

region. As a result, the image was split into many regions of different sizes. The common pixel value of all the pixels of a region was considered as the pixel value of the region. Two regions were considered adjacent if they shared a common boundary. Any region R with pixel value v_R was considered as a region of interest for exudate detection if it fulfilled two conditions :

i All the regions adjacent to the external boundary of R have pixel values less than v_R.

ii There is no region totally contained in R that has pixel value greater than v_R.

3.3.2 Feature Extraction for Exudate Detection

A total of 14 features were extracted corresponding to each region of interest. Among them, eight features were related to the properties of pixels belonging to the region. In addition, six more features were extracted from the surrounding region of each region of interest. The surrounding of each region of interest was defined as the region bounded by a rectangle having width and height 10 pixels larger than the minimum bounding rectangle of the candidate region. The total 14 features extracted from each candidate region and its surrounding are listed below:

i Mean value of red component of pixels in region

ii Mean value of green component of pixels in region

iii Mean value of blue component of pixels in region

iv Standard deviation of red component of pixels in region

v Standard deviation of green component of pixels in region

vi Standard deviation of blue component of pixels in region

vii Mean value of red component of pixels in surrounding area

viii Mean value of green component of pixels in surrounding area

ix Mean value of blue component of pixels in surrounding area

x Standard deviation of red component of pixels in surrounding area

xi Standard deviation of green component of pixels in surrounding area

xii Standard deviation of blue component of pixels in surrounding area

xiii Ratio of blue component to green component of pixels in region

xiv Ratio of green component to red component of pixels in region

3.3.3 Classification for Exudate Detection

Four different classifiers, namely - Naïve Bayes, Support Vector Machine (SVM), k-Nearest Neighbor (kNN), and Artificial Neural Network (ANN) were used for classifying each extracted region of interest as exudate or non-exudate. For this purpose, a dataset named IDRiD[65] was used. This dataset consists of 81 fundus images and each image is accompanied with a label image. For each exudate pixel in the fundus image, the corresponding pixel in the label image is marked with a value of one. All other pixels in the label image are marked with zero.

A total of $10,786$ regions of interest were extracted from the IDRiD image dataset. For any region, if the number of exudate pixels in the corresponding pixels in the label image was more than half of the number of pixels in the region, then that region was labeled as a positive region for exudate. Otherwise, the region was labeled as negative region. In the constructed region

dataset, 6,272 regions were found to be positive and the remaining 4,514 regions were found to be negative. The dataset was split into two parts for training and testing purpose. 2,157 regions were selected randomly which is in fact almost 20% of the total dataset and included in the test dataset. The remaining 8,629 regions were used for training purpose.

For Naïve Bayes classifier, the dataset was assumed to be distributed according to a Gaussian distribution. kNN was tried with different values of k and experimentally k=5 was found to provide the best results. The dataset was normalized to have zero mean and unit standard deviation before training the SVM and ANN classifiers. A radial basis kernel function was used for performing SVM classification. Hyperparameters for ANN classifier were tuned by extensive trial and error experimentations. The best results were found with a network consisting of two hidden layers with 15 and 12 neurons respectively. The network was trained for 200 epochs with an Adam optimization function. The results of each classifier are discussed in details in Section 4.

3.4 Detection of Hemorrhages and Microaneurysms

Both hemorrhages and microaneurysms are darker than the background of a fundus image and therefore they were together considered as dark lesions. Detection of dark lesions was conducted following three major steps. First, some preprocessing techniques were applied to enhance the dark lesion regions in a fundus image. Second, 16 distinguishing features were extracted from each region. Finally, the regions were classified as to be dark lesions or not based on the features using a machine learning classifier. These three steps are elaborately discussed below.

3.4.1 Preprocessing

Hemorrhages and microaneurysms are circular or nearly circular dark lesions present in a fundus image. However, blood vessels are similar in color with dark lesions. So, merely identifying dark pixels as dark lesions could result in misclassifying the blood vessels into dark lesions. In order to minimize this misclassification, blood vessels were first detected following procedures

mentioned in Section 3.2 and masked out. Yet, some isolated vessel pixels could remain in the image due to existing inaccuracy in the vessel detection process. Distinguishing these isolated vessel segments from dark lesions was a major challenge in this detection process. First, green channel of the masked image was extracted. Then, in order to reduce the effect of variable illumination, piecewise linear contrasst enhancement was applied on the green channel. It was observed that, pixel values of dark lesions were always situated within the first quartile of all the values of a contrast enhanced green channel. Therefore, the contrast enhanced green channel was thresholded at the first quartile of the pixel values. This resulted in a binary image where pixels darker than the threshold value were marked zero and the other pixels were marked with one. The resultant image was then inverted to obtain a binary image with black background and white regions on it. The regions were identified by detecting connected components in the binary image.

3.4.2 Feature Extraction for Dark Lesion Detection

A total of 16 features were extracted from each connected component identified in the pre-processing step. Six features - consisting of mean and standard deviation of each of the three channels were extracted from each region. Since the values of each channel may vary depending on the illumination of the image taken, ratio of the blue channel to green channel and green channel to red channel were also considered to reduce the effect of variable illumination. In addition, six more features were extracted from the surrounding area of each region. The surrounding of each candidate region was defined in the same way as used in detection of exudates i.e. the region bounded by a rectangle having width and height 10 pixels larger than the minimum bounding rectangle of the candidate region. The mean and standard deviation of the pixel values in the surrounding region were also considered as features for dark lesion detection from the candidate regions. Two more features related to shape of the regions were considered for differentiating dark lesions from isolated blood vessel segments as mentioned earlier. Dark lesions are usually circular or elliptical in shape whereas blood vessel segments are linear or cylindrical in shape.The circularity of each region of interest was measured by fitting an ellipse around the circumference of the region and then taking the ratio of the two axes of the ellipse.

The closer the value to 1.0, the more circular the region was. The ratio of the total area of the region to the area of the circle circumscribing the region was also calculated for checking circularity of the region. All the features identified for the classification of the candidate regions into dark lesions or blood vessel segments are listed below.

i Mean value of red component of pixels in region

ii Mean value of green component of pixels in region

iii Mean value of blue component of pixels in region

iv Standard deviation of red component of pixels in region

v Standard deviation of green component of pixels in region

vi Standard deviation of blue component of pixels in region

vii Mean value of red component of pixels in surrounding area

viii Mean value of green component of pixels in surrounding area

ix Mean value of blue component of pixels in surrounding area

x Standard deviation of red component of pixels in surrounding area

xi Standard deviation of green component of pixels in surrounding area

xii Standard deviation of blue component of pixels in surrounding area

xiii Ratio of blue component to green component of pixels in region

xiv Ratio of green component to red component of pixels in region

xv Ratio of area of the region to the area of a circle circumscribing the region

xvi Ratio of the two axes of an ellipse circumscribing the region

3.4.3 Classification for Dark Lesion Detection

Four different classifiers, namely - Naïve Bayes, Support Vector Machine (SVM), k-Nearest Neighbor (kNN), and Artificial Neural Network (ANN) were used for the classification purpose. The IDRiD dataset mentioned earlier was also used for the classification of dark lesions from isolated vessel pixels. A total of $7,789$ regions were extracted from 81 images in the dataset. Each region was labeled as positive or negative based on the presence of dark lesions in the corresponding pixels in the label image. In the constructed region dataset, $4,172$ regions were found to be positive and the remaining $3,617$ regions were found to be negative. The dataset was split into two parts for training and testing purpose. $1,558$ regions were selected randomly which is in fact almost 20% of the total dataset and included in the test dataset. The remaining $6,231$ regions were used for training purpose.

For Naïve Bayes classifier, the dataset was assumed to be distributed according to a Gaussian distribution. kNN was tried with different values of k and experimentally k=7 was found to provide the best results. The dataset was normalized to have zero mean and unit standard deviation before training the SVM and ANN classifiers. A radial basis kernel function was used for performing SVM classification. Hyperparameters for ANN classifier were tuned by extensive trial and error experimentations. The best results were found with a network consisting of two hidden layers with 20 and 15 neurons respectively. The network was trained for 200 epochs with an Adam optimization function. The results of each classifier are discussed in details in Section 4.

3.5 Extraction of Features for DR Detection

The severity of Diabetic Retinopathy can be detected from a fundus image based on the number of different lesion pixels and their location. The severity detection method proposed by Wilkinson et. al. divides a fundus image into four quadrants with respect to the optic disk. Different severity levels are determined based on the existence of exudates, microaneurysms, and hemorrhages in each quadrant and their distance from the centre of the optic disc. The detailed description of the grading process can be found in the paper [3]. Based on this method,

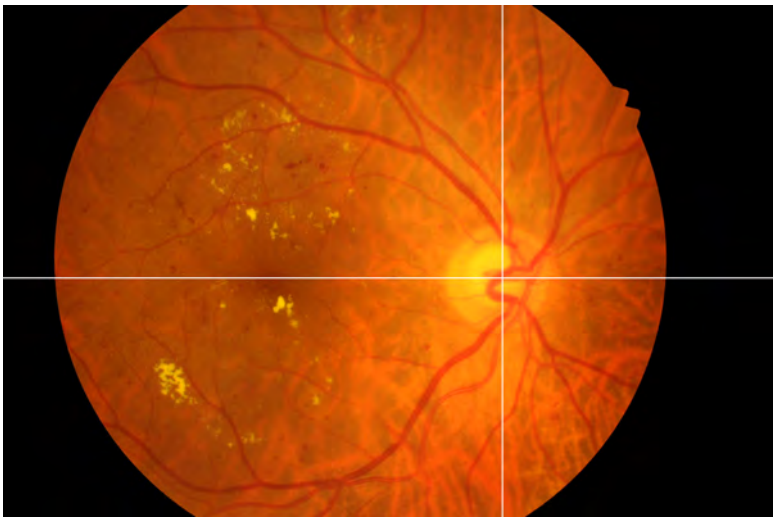

Figure 3.14: A typical fundus image divided into quadrants

14 features were extracted from the components identified in the previous sections.

At first, it was required to divide a fundus image into four quadrants. A rectangular bounding box surrounding the optic disc was constructed as described in Section 3.1. The mean of the breadth and height of the rectangular region was considered as the diameter of the optic disc D. The intersecting point of two diagonals of the bounding box was considered as the centre of the optic disc. The image was divided into four quadrants with respect to this centre of optic disc as origin. A sample fundus image divided into four quadrants is illustrated in figure 3.14. Exudates and dark lesions were then identified from each image. The distance of the lesion pixels from optic disc centre were determined by calculating Euclidean distance between the pixel locations in the image.

Following is the list of features extracted from a fundus image which was used for DR detection.

i Total number of exudate pixels in the image

ii Total number of dark lesion pixels in the image

iii Exudate pixels within distance D from the Optic Disk centre

iv Dark lesion pixels within distance D from the optic disk centre

v Exudate pixels within distance $2D$ from the Optic Disk centre

vi Dark lesion pixels within distance $2D$ from the optic disk centre

vii Number of exudate pixels in the first quadrant

viii Number of exudate pixels in the second quadrant

ix Number of exudate pixels in the third quadrant

x Number of exudate pixels in the fourth quadrant

xi Number of dark lesion pixels in the first quadrant

xii Number of dark lesion pixels in the second quadrant

xiii Number of dark lesion pixels in the third quadrant

xiv Number of dark lesion pixels in the fourth quadrant

3.6 Classification of DR Severity

The severity of Diabetic Retinopathy in a fundus image was detected by classifying it into different DR levels basing on 14 features. Two different classifiers - SVM and Artificial Neural Network were trained for this detection task. A publicly available labeled data set by IDRiD[65] containing 413 training images and 103 test images were used for training and testing the classifiers respectively. The images were labeled from 0 to 4 denoting the severity of Diabetic Retinopathy as followed:

0 - No apparent retinopathy

1 - Mild nonproliferative DR

2 - Moderate nonproliferative DR

3 - Severe nonproliferative DR

4 - Proliferative DR

A training dataset was constructed with features extracted from 413 training images. The dataset was normalized to have zero mean and unit variance. An SVM classifer with radial basis kernel function was trained with the training dataset. The number of layers and number of nodes in each layer for the artificial neural network classifier was determined by performing a grid search. A network consisting of two hidden layers with four neurons in each of them was found to provide the best performance. For each of the classifiers, training was performed with cross-validation over 20% of the training data. Performance of each classifier was measured by evaluating them on the test images. The performance obtained by each classifier is discussed in Section 4.

Chapter 4

Results and Contributions

The three major contributions made by this work as mentioned in Section1.2 are discussed in this chapter. First, the performances of the techniques developed for identification of each of the anatomical components are analyzed. Then, a diagnosis tool to assist ophthalmologists in DR diagnosis developed in this work is presented. Finally, the overall accuracy of detection of DR from fundus image based on the identified components are demonstrated.

4.1 Results of Component Extraction

4.1.1 Detection of Optic Disc

The trained YOLO model for optic disc detection was tested over a completely unobserved fundus image dataset by IDRiD [65]. The dataset contains 81 images with pixel level labeling of optic disc. For performance evaluation of our technique, the pixels marked after segmentation were compared with the ground truth in the label image. The segmentation was considered correct if the intersection over union between the detected region and the ground truth was 95% or more. Following this evaluation methodology, an accuracy of 100% ws achieved for the test dataset.

Compared to existing works in the literature, The proposed method in this work for detection of optic disc has been able to identify the optic disc in some special cases where the

existing methods in the literature were unable to do so. These special cases along with the challenges of each case are discussed elaborately as followed.

Special Case 1: Images where optic disc is not the brightest region

Most segmentation and thresholding techniques used for optic disc detection make the assumption that, optic disc is the brightest region in a fundus image. These methods threshold the fundus image at a particular value, possibly selected by experimentation, and then considers only the regions that are brighter than the selected value. Though true for most cases, some special images were found in the Ubicomp dataset [54] where optic disc is not the brightest region in the entire fundus image. The detection technique developed in this work could locate and segment out the optic disc successfully from such special images where many of the state of the art techniques could not. A sample fundus image for this case and the corresponding detection output of the trained model is demonstrated in Figure 4.1.

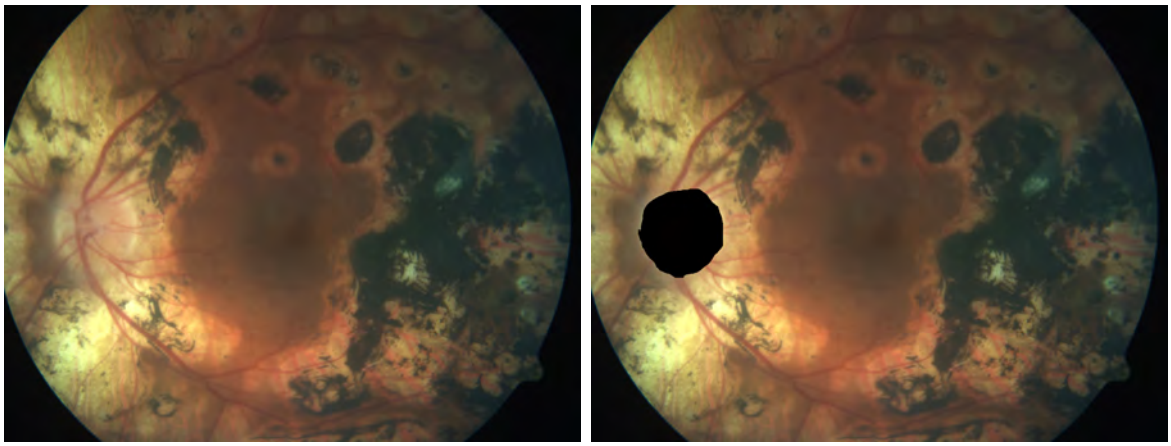

(a) Sample fundus image (b) Detected Optic Disc

Figure 4.1: A special case for fundus image where optic disc is not the brightest region.

Special Case 2: Images with lesions similar to the optic disc

Optic disks are nearly circular or elliptical in shape and have a definite range of size. Many existing techniques for detecting optic disc leverage this phenomenon and identify any bright region having such shape and size as optic disc. However, the dataset used for training the

model in this work contained some special images with bright lesions that are similar in shape and size to the optic disc. Many of the existing techniques are prone to make false positive detections of optic disc for such cases. The detection method developed in this work is free from making such false positives. A sample fundus image having such a bright lesion and the corresponding output of the detection model is illustrated in Figure 4.2

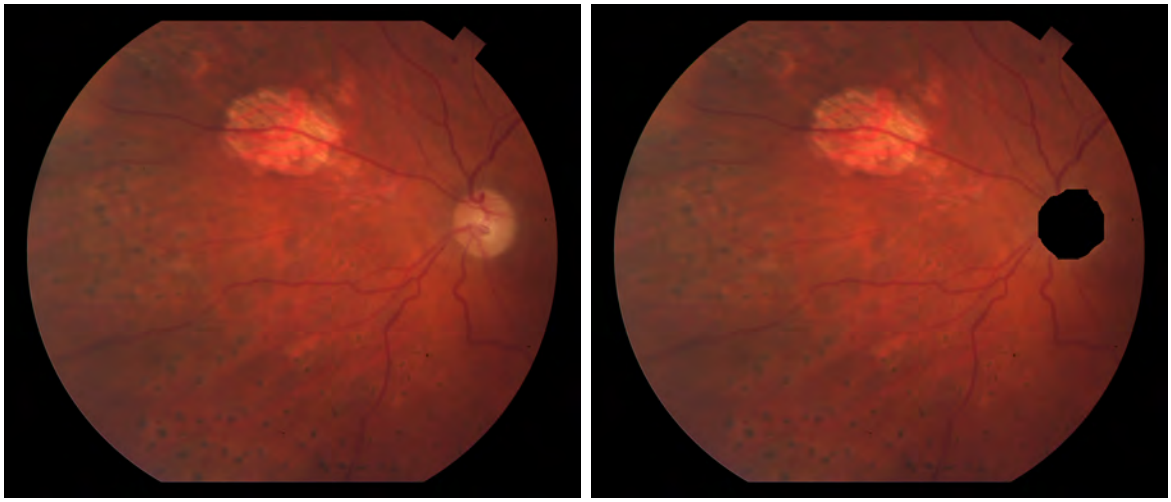

(a) Sample fundus image (b) Detected Optic Disc

Figure 4.2: A special case for fundus image where lesions similar in size, shape, and color to the optic disc is present.

Special Case 3: Images where optic disc is absent

Fundus images can be taken from different angles with different fields of view. Sometimes a particular portion of the retina is intended to be captured in stead of the entire retina. As a result, there can be images without the optic disc present in it. Existing detection techniques based on segmentation presume that the optic disc must be present in the image. This may result in false detections if bright lesions are present in an image without optic disc. Such an image is shown in Figure 4.3 where optic disc is absent but bright lesions of similar shapes, which were falsely identified as optic disc by some of the state of the art methods, are present. The detection method developed in this work did not make any such false detection of optic disc.

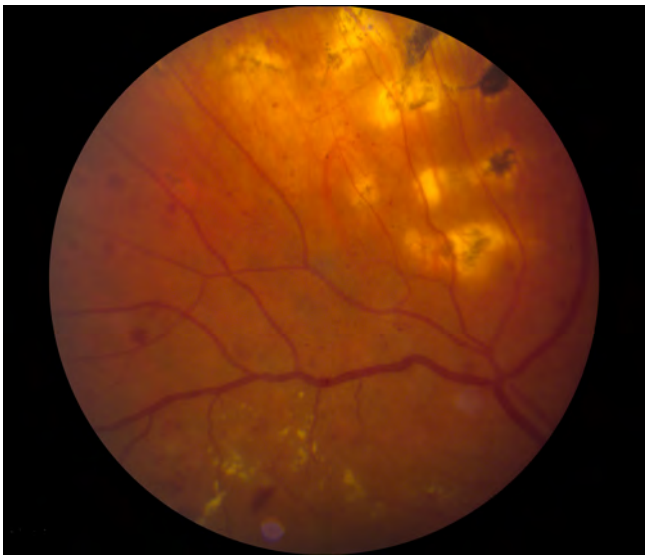

Figure 4.3: A special case for fundus image where optic disc is absent.

4.1.2 Detection of Blood Vessels

The output probability image generated by the RNN network for detection of blood vessel was first thresholded at a value t. Thresholding at a small value of t would identify pixels with small probability as vessel pixel. As a result, true positive detections would be increased. However, this will also identify microaneurysms and other non-vessel dark pixels as vessel-pixel and thus increase false positive detections. In contrast, thresholding at a large value of t resulted in increase of true negatives since pixels having low probabilities were discarded. However, some of the vessel pixels were found to be discarded along with resulting in an increase in false negative detection. Two parameters: sensitivity and specificity to evaluate the performance of the blood vessel detection method where:

$$Sensitivty = \frac{TruePositive}{TruePositive + FalseNegative}$$

$$Specificity = \frac{TrueNegative}{TrueNegative + FalsePositive}$$

The detection method was tested with 20 previously unobserved test fundus images. The value of t was varied from 0.50 up to 1.00 with an increase of value by 0.05 at every step. Specificity and sensitivity at every such value of t was calculated. High sensitivity of detection

Value of t	**Average Sensitivity for 20 Test Images**	**Average Specificity for 20 Test Images**
0.50	0.987	0.725
0.55	0.973	0.727
0.60	0.952	0.761
0.65	0.941	0.772
0.70	0.913	0.810
0.75	0.875	0.863
0.80	0.783	0.937
0.85	0.747	0.958
0.90	0.624	0.982
0.95	0.556	1.000

Table 4.1: Sensitivity and Specificity obtained by the Blood vessel detection technique for different values of thresholding parameter t.

was obtained at lower values of t with low specificity. On the other hand, higher value of t produced low sensitivity along with high specificity of detection. The average sensitivity and specificity obtained for 20 test images at different values of t as mentioned are listed in Table 4.1. The trade-off between sensitivity and specificity can be observed from the Receiver Operating Characteristics (ROC) Curve shown in Figure 4.4. There is a positive correlation between specificity and the complemented value of sensitivity which actually implies a negative correlation between specificity and sensitivity. From 20 test images, a sensitivity of 98.7% and specificity of 72.5% was obtained at the point of highest sensitivity. Conversly, A sensitivity of 55.6% and specificity of 100% was obtained at the point of highest specificity.

4.1.3 Detection of Exudates

Four different classifiers were used for the detection of exudates from extracted regions of interest. They are: Naïve Bayes, kNN, Support Vector Machine, and Artificial Neural Network.

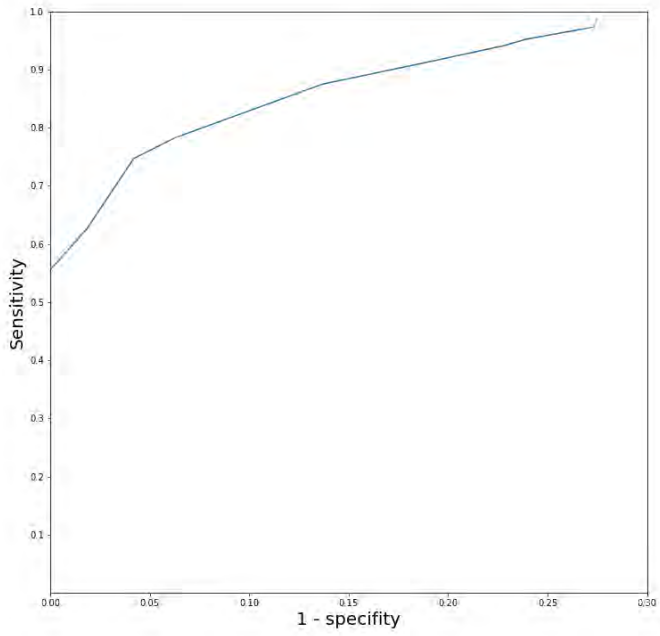

Figure 4.4: ROC curve for the blood vessel detection technique.

Table 4.2 shows the accuracy obtained by each classifier.

Classifier	**Accuracy**
Naïve Bayes	63.70%
Support Vector Machine (SVM)	74.91%
k Nearest Neighbor (kNN)	73.20%
Artificial Neural Network (ANN)	76.34%

Table 4.2: Accuracy obtained by different classifiers for the detection of exudates.

The best accuracy was obtained with an ANN classifier consisting of two hidden layers with 12 and 15 neurons in them respectively. The specificity and sensitivity achieved were 85.2% and 71% respectively by this classifier. To detect exudate from an input fundus image, first regions of interest were extracted from the image following procedures mentioned in Section 3 and then the extracted regions were classified to be exudate or non-exudate by the trained ANN

classifier. A sample fundus image and the exudates detected from it following the discussed method is shown in Figure 4.5

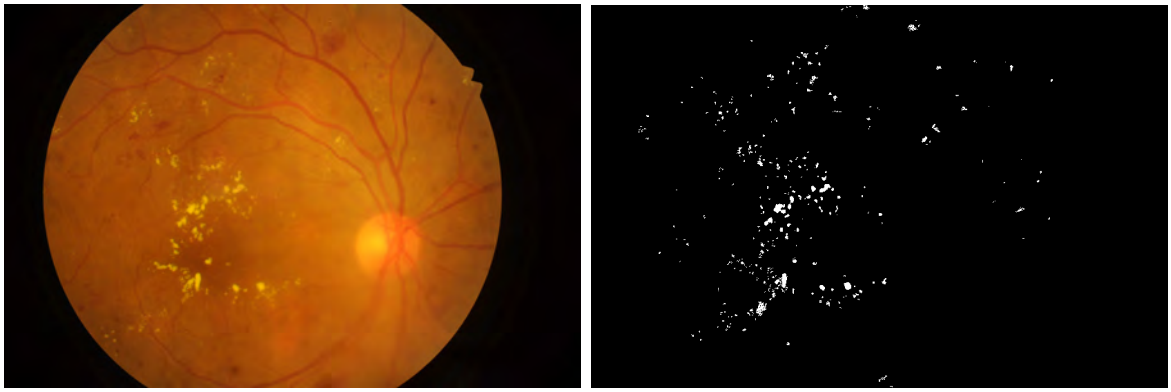

(a) Sample fundus image (b) Exudates marked in white

Figure 4.5: Output generated by the exudate detection technique applied on a fundus image.

4.1.4 Detection of Dark Lesions

Detection of dark lesions was performed first by extracting candidate regions from a fundus image and then classifying the regions. Four different classifiers were trained for the classification task and the classifier with the best performance was chosen. The accuracy obtained by the four classifiers: Naïve Bayes, kNN, Support Vector Machine, and Artificial Neural Network are demonstrated in Table 4.3.

Classifier	**Accuracy**
Naïve Bayes	56.80%
Support Vector Machine (SVM)	69.33%
k Nearest Neighbor (kNN)	69.20%
Artificial Neural Network (ANN)	71.60%

Table 4.3: Accuracy obtained by different classifiers for the detection of microaneurysms and hemorrhages.

The best accuracy was obtained with an ANN classifier consisting of two hidden layers with 20 and 15 neurons in them respectively. The specificity and sensitivity achieved were 76.2%

and 69% respectively by this classifier. To detect dark lesions from an unobserved fundus image, first regions of interest were extracted from the image following procedures mentioned in Section 3 and then the extracted regions were classified to be lesions or benign regions by the trained ANN classifier. A sample fundus image and the dark lesions detected from it following the discussed method is shown in Figure 4.6

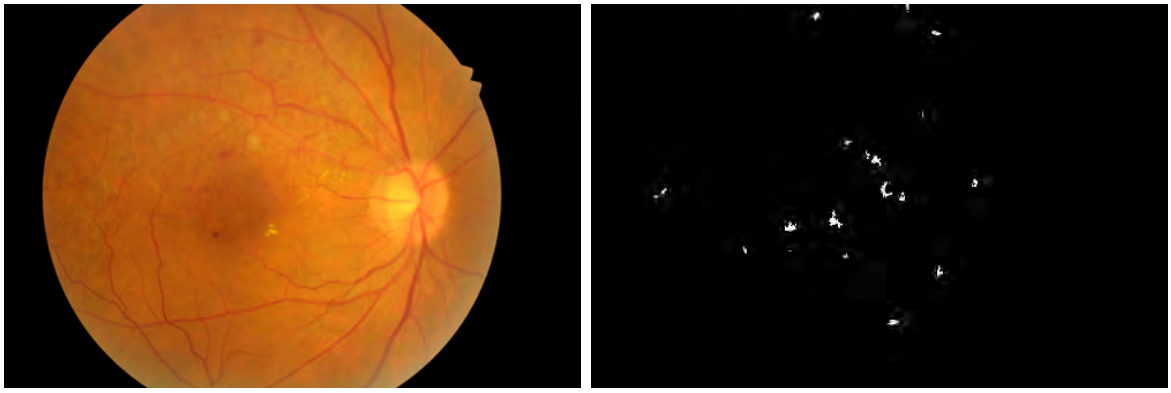

(a) Sample fundus image (b) Dark lesions marked in white

Figure 4.6: Output generated by the dark lesion detection technique applied on a fundus image.

4.2 A Diagnosis Tool

A software tool was developed to provide a visual representation of different lesions present in an input fundus image as well as a statistical summary of them. The tool, developed in python programming language, takes a fundus image as input and generates two binary label images marking the exudates and dark lesions present in the fundus image respectively. The four quadrants and two circles with radius D and $2D$ are also incorporated in the label images. In addition, it also provides a summary of the percentage of different lesions present in each quadrant and inside the drawn circles. Ophthalmologists can get a better idea about the lesions in an image using this tool. Providing a quantitative information about the lesions, this tool can come into a great assistance to reduce manual labor of ophthalmologists while diagnosing DR severity. Figure 4.7 shows the statistics calculated ad output generated by the tool for a sample fundus image.

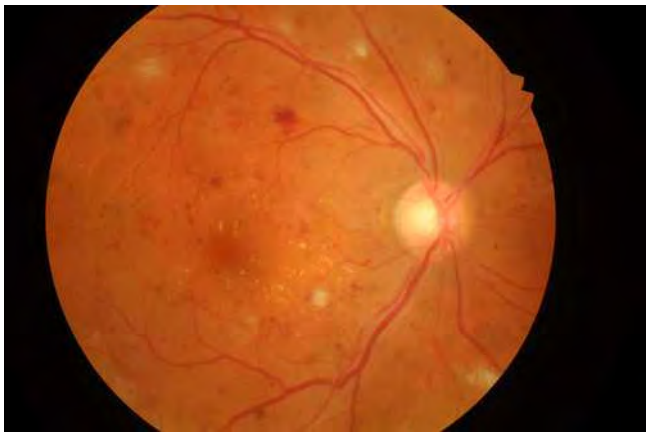

(a) Sample fundus image

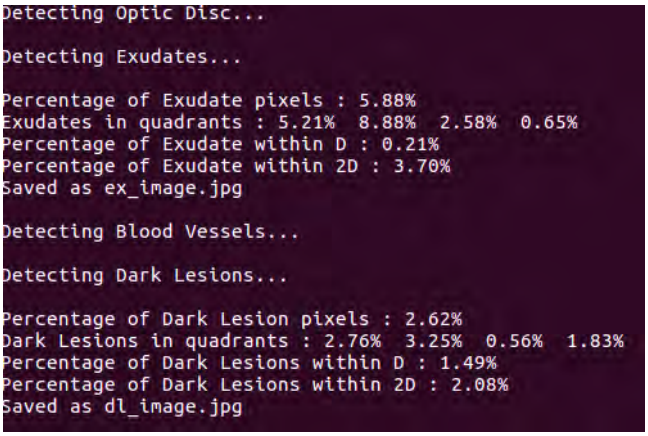

(b) Statistics generated by developed tool

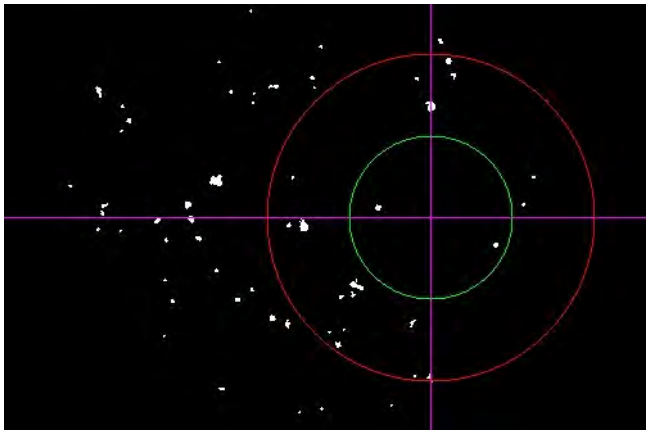

(c) Dark Lesions detected by the tool

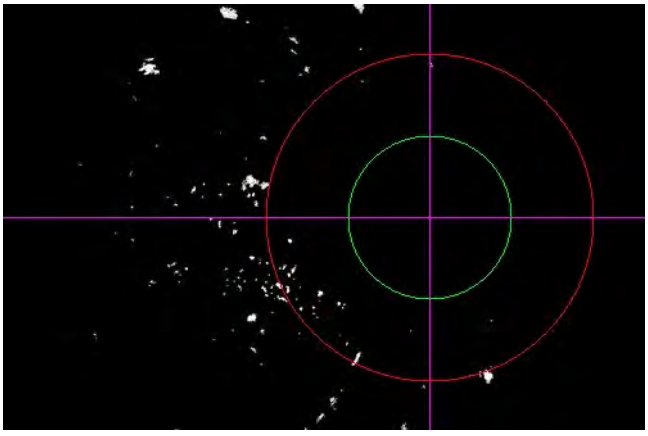

(d) Exudates detected by the tool

Figure 4.7: Output generated by the diagnosis tool developed in this work

4.3 Results of DR Detection

Two different classifiers : SVM and ANN were trained for the detection of DR severity into five stages as proposed by Wilkinson [3]. The performance of each classifier was measured by the number of correct prediction made by it in detecting the DR severity of 103 images in an unobserved test dataset. Number of images belonging to each DR severity level for both the training and test dataset are depicted in Table 4.4. The number of correct prediction made by the SVM classifer was 56 which implies a detection accuracy of 54.36%. ANN classifier produced a better performance by correctly detecting DR severity of 61 test images with a detection accuracy of 59.22%. Table 4.5 and Table 4.6 show the confusion matrix for each of the classifiers.

DR Severity	Train Dataset	Test Dataset
0	135	34
1	20	5
2	136	32
3	74	19
4	49	13

Table 4.4: Number of images for each DR severity level in the train and test dataset used for detection.

From Table 4.4 it can be observed that, number of images available for training fluctuate significantly for different DR severity levels. For example, there were 136 images belonging to severity level 2 while severity level 1 had only 20 train images. This might have affected the capability of the classifiers to identify each severity level correctly irrespective to its value. Also, all misclassification errors are not the same in terms of evaluation of the models. As an instance, an image with severity level 4 misclassified as level 3 should be considered to have smaller error than the same image if misclassified as severity level 0. We calculated the root mean square error (rmse) for each classifier to account for this phenomenon. If N is the number of test images, then root mean square error can be defined as:

$$rmse = \sqrt{\frac{\sum_{i=1}^{N}(PredictedSeverity - ActualSeverity)^2}{N}}$$

The calculated root mean square error for the SVM classifier and ANN classifier were 1.14 and 1.34 respectively.

Actual DR Stage \ Predicted DR Stage	0	1	2	3	4
0	21	0	11	1	1
1	3	0	2	0	0
2	5	3	19	4	1
3	0	0	6	10	3
4	0	0	4	3	6

Table 4.5: Confusion matrix for classification of 103 test fundus images into five stages of Diabetic Retinopathy by a Support Vector Machine (SVM) classifier.

Actual DR Stage \ Predicted DR Stage	0	1	2	3	4
0	23	1	6	4	0
1	2	1	2	0	0
2	5	0	21	5	1
3	4	1	4	9	1
4	2	0	3	1	7

Table 4.6: Confusion matrix for classification of 103 test fundus images into five stages of Diabetic Retinopathy by an Artificial Neural Network (ANN) classifier.

Chapter 5

Conclusion

This work has attempted to detect the severity level of Diabetic Retinopathy after identifying five different anatomical components from fundus images. There are very few works existing in the literature demonstrating the detection task. In fact, this work is the first of its kind on the particular dataset used for detection based on extracted components. Retrospective methods have used comparatively smaller datasets with little variation in illumination and texture. A major challenge in the process was lack of labeled data for all the components as well as DR severity. As a result, different datasets were used to train models for identifying different components. After that, the models were combined to extract features from the image dataset used for DR severity detection. This also supports the generalization capability of the component extraction models irrespective of image dataset. However, there is still room for improvement and can be incorporated in future studies.

Since DR severity level was detected based on features related to extracted components, errors occurring in each component extraction step cumulated to contribute a larger error rate in the final severity detection phase. Therefore, improvement in the performance of each of the component extraction technique is expected to contribute in improvement of the severity detection performance. The lesions identified from fundus images in this work are also helpful in the diagnosis of some other retinal disease such as glaucoma, macular edema, etc. Therefore, the developed techniques in this work can also be utilized in studies related to the diagnosis of the mentioned diseases.

The severity detection model was trained with a dataset containing 400 images. Availability of a larger labeled dataset would help reducing model under-fitting and thus improve the severity detection accuracy. Moreover, incorporation of domain knowledge regarding the relevance of each component for different DR stage would also help this cause. A unified system for identifying each component and then detecting the DR severity from fundus image of a patient can be a promising doorway for future research works in this field.

Bibliography

[1] Janet L Leasher, Rupert RA Bourne, Seth R Flaxman, Jost B Jonas, Jill Keeffe, Kovin Naidoo, Konrad Pesudovs, Holly Price, Richard A White, Tien Y Wong, et al. Global estimates on the number of people blind or visually impaired by diabetic retinopathy: a meta-analysis from 1990 to 2010. *Diabetes care*, 39(9):1643–1649, 2016.

[2] Yingfeng Zheng, Mingguang He, and Nathan Congdon. The worldwide epidemic of diabetic retinopathy. *Indian journal of ophthalmology*, 60(5):428, 2012.

[3] CP Wilkinson, Frederick L Ferris III, Ronald E Klein, Paul P Lee, Carl David Agardh, Matthew Davis, Diana Dills, Anselm Kampik, R Pararajasegaram, Juan T Verdaguer, et al. Proposed international clinical diabetic retinopathy and diabetic macular edema disease severity scales. *Ophthalmology*, 110(9):1677–1682, 2003.

[4] Nathan Silberman, Kristy Ahrlich, Rob Fergus, and Lakshminarayanan Subramanian. Case for automated detection of diabetic retinopathy. In *Proceedings of the AAAI Spring Symposium: Artificial Intelligence for Development, California, USA, March 22-24*, 2010.

[5] Sobha Sivaprasad, Bhaskar Gupta, Roxanne Crosby-Nwaobi, and Jennifer Evans. Prevalence of diabetic retinopathy in various ethnic groups: a worldwide perspective. *Survey of ophthalmology*, 57(4):347–370, 2012.

[6] Chanjira Sinthanayothin, James F Boyce, Tom H Williamson, Helen L Cook, Evelyn Mensah, Shantanu Lal, and David Usher. Automated detection of diabetic retinopathy on digital fundus images. *Diabetic medicine*, 19:105–112, 2002.

[7] Wong Li Yun, U Rajendra Acharya, Yedatore V Venkatesh, Caroline Chee, Lim Choo Min, and E Yin Kwee Ng. Identification of different stages of diabetic retinopathy using retinal optical images. *Information sciences*, 178(1):106–121, 2008.

[8] Bernhard M Ege, Ole K Hejlesen, Ole V Larsen, Karina Møller, Barry Jennings, David Kerr, and David A Cavan. Screening for diabetic retinopathy using computer based image analysis and statistical classification. *Computer methods and programs in biomedicine*, 62(3):165–175, 2000.

[9] Karthikeyan Ganesan, Roshan Joy Martis, U Rajendra Acharya, Chua Kuang Chua, Lim Choo Min, EYK Ng, and Augustinus Laude. Computer-aided diabetic retinopathy detection using trace transforms on digital fundus images. *Medical & biological engineering & computing*, 52(8):663–672, 2014.

[10] R John Winder, Philip J Morrow, Ian N McRitchie, JR Bailie, and Patricia M Hart. Algorithms for digital image processing in diabetic retinopathy. *Computerized medical imaging and graphics*, 33(8):608–622, 2009.

[11] Chanjira Sinthanayothin, James F Boyce, Helen L Cook, and Thomas H Williamson. Automated localisation of the optic disc, fovea, and retinal blood vessels from digital colour fundus images. *British Journal of Ophthalmology*, 83(8):902–910, 1999.

[12] James Lowell, Andrew Hunter, David Steel, Ansu Basu, Robert Ryder, Eric Fletcher, Lee Kennedy, et al. Optic nerve head segmentation. *Medical Imaging, IEEE Transactions on*, 23(2):256–264, 2004.

[13] Thomas Walter and Jean-Claude Klein. Segmentation of color fundus images of the human retina: Detection of the optic disc and the vascular tree using morphological techniques. In *International Symposium on Medical Data Analysis*, pages 282–287. Springer, 2001.

[14] Saiprasad Ravishankar, Arpit Jain, and Anurag Mittal. Automated feature extraction for early detection of diabetic retinopathy in fundus images. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 210–217. IEEE, 2009.

[15] Marc Lalonde, Mario Beaulieu, and Langis Gagnon. Fast and robust optic disc detection using pyramidal decomposition and hausdorff-based template matching. *IEEE transactions on medical imaging*, 20(11):1193–1200, 2001.

[16] Dongbo Zhang and Yuanyuan Zhao. Novel accurate and fast optic disc detection in retinal images with vessel distribution and directional characteristics. *IEEE J. Biomedical and Health Informatics*, 20(1):333–342, 2016.

[17] Javad Rahebi and Fırat Hardalaç. A new approach to optic disc detection in human retinal images using the firefly algorithm. *Medical & biological engineering & computing*, 54(2-3):453–461, 2016.

[18] RA Welikala, Muhammad Moazam Fraz, Jamshid Dehmeshki, Andreas Hoppe, V Tah, S Mann, Thomas H Williamson, and Sarah A Barman. Genetic algorithm based feature selection combined with dual classification for the automated detection of proliferative diabetic retinopathy. *Computerized Medical Imaging and Graphics*, 43:64–77, 2015.

[19] Subhasis Chaudhuri, Shankar Chatterjee, Norman Katz, Mark Nelson, and Michael Goldbaum. Detection of blood vessels in retinal images using two-dimensional matched filters. *IEEE Transactions on medical imaging*, 8(3):263–269, 1989.

[20] Muhammad Moazam Fraz, Paolo Remagnino, Andreas Hoppe, Bunyarit Uyyanonvara, Alicja R Rudnicka, Christopher G Owen, and Sarah A Barman. Blood vessel segmentation methodologies in retinal images–a survey. *Computer methods and programs in biomedicine*, 108(1):407–433, 2012.

[21] Bob Zhang, Lin Zhang, Lei Zhang, and Fakhri Karray. Retinal vessel extraction by matched filter with first-order derivative of gaussian. *Computers in biology and medicine*, 40(4):438–445, 2010.

[22] AD Hoover, Valentina Kouznetsova, and Michael Goldbaum. Locating blood vessels in retinal images by piecewise threshold probing of a matched filter response. *IEEE Transactions on Medical imaging*, 19(3):203–210, 2000.

[23] José Ignacio Orlando, Marcos Fracchia, Valeria del Río, and Mariana del Fresno. Retinal blood vessel segmentation in high resolution fundus photographs using automated feature parameter estimation. In *13th International Conference on Medical Information Processing and Analysis*, volume 10572, page 1057210. International Society for Optics and Photonics, 2017.

[24] Joes Staal, Michael D Abràmoff, Meindert Niemeijer, Max A Viergever, and Bram Van Ginneken. Ridge-based vessel segmentation in color images of the retina. *IEEE transactions on medical imaging*, 23(4):501–509, 2004.

[25] Lili Xu and Shuqian Luo. A novel method for blood vessel detection from retinal images. *Biomedical engineering online*, 9(1):14, 2010.

[26] Carmen Alina Lupascu, Domenico Tegolo, and Emanuele Trucco. FABC: retinal vessel segmentation using adaboost. *IEEE Transactions on Information Technology in Biomedicine*, 14(5):1267–1274, 2010.

[27] Sameh A Salem, Nancy M Salem, and Asoke K Nandi. Segmentation of retinal blood vessels using a novel clustering algorithm (racal) with a partial supervision strategy. *Medical & biological engineering & computing*, 45(3):261–273, 2007.

[28] Lucia Espona, María J Carreira, MG Penedo, and Marcos Ortega. Retinal vessel tree segmentation using a deformable contour model. In *Pattern Recognition, 2008. ICPR 2008. 19th International Conference on*, pages 1–4. IEEE, 2008.

[29] Bashir Al-Diri, Andrew Hunter, and David Steel. An active contour model for segmenting and measuring retinal vessels. *IEEE Transactions on Medical imaging*, 28(9):1488–1497, 2009.

[30] Yong Yang, Shuying Huang, and Nini Rao. An automatic hybrid method for retinal blood vessel extraction. *International Journal of Applied Mathematics and Computer Science*, 18(3):399–407, 2008.

[31] Huazhu Fu, Yanwu Xu, Damon Wing Kee Wong, and Jiang Liu. Retinal vessel segmentation via deep learning network and fully-connected conditional random fields. In *Biomedical Imaging (ISBI), 2016 IEEE 13th International Symposium on*, pages 698–701. IEEE, 2016.

[32] Kevis-Kokitsi Maninis, Jordi Pont-Tuset, Pablo Arbeláez, and Luc Van Gool. Deep retinal image understanding. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 140–148. Springer, 2016.

[33] Shuangling Wang, Yilong Yin, Guibao Cao, Benzheng Wei, Yuanjie Zheng, and Gongping Yang. Hierarchical retinal blood vessel segmentation based on feature and ensemble learning. *Neurocomputing*, 149:708–717, 2015.

[34] Thomas Walter, J-C Klein, Pascale Massin, and Ali Erginay. A contribution of image processing to the diagnosis of diabetic retinopathy-detection of exudates in color fundus images of the human retina. *IEEE transactions on medical imaging*, 21(10):1236–1243, 2002.

[35] Wynne Hsu, PMDS Pallawala, Mong Li Lee, and Kah-Guan Au Eong. The role of domain knowledge in the detection of retinal hard exudates. In *Computer Vision and Pattern Recognition, 2001. CVPR 2001. Proceedings of the 2001 IEEE Computer Society Conference on*, volume 2, pages II–II. IEEE, 2001.

[36] M Esmaeili, H Rabbani, AM Dehnavi, and A Dehghani. Automatic detection of exudates and optic disk in retinal images using curvelet transform. *IET image processing*, 6(7):1005–1013, 2012.

[37] C Eswaran, Ahmed Wasif Reza, and Subhas Hati. Extraction of the contours of optic disc and exudates based on marker-controlled watershed segmentation. In *Computer Science and Information Technology, 2008. ICCSIT'08. International Conference on*, pages 719–723. IEEE, 2008.

[38] Huiqi Li and Opas Chutatape. Automated feature extraction in color retinal images by a model based approach. *IEEE Transactions on biomedical engineering*, 51(2):246–254, 2004.

[39] Alireza Osareh, Bita Shadgar, and Richard Markham. A computational-intelligence-based approach for detection of exudates in diabetic retinopathy images. *IEEE Transactions on Information Technology in Biomedicine*, 13(4):535–545, 2009.

[40] C Jayakumari and T Santhanam. Detection of hard exudates for diabetic retinopathy using contextual clustering and fuzzy art neural network. *Asian journal of information technology*, 6(8):842–846, 2007.

[41] Pavle Prentašić and Sven Lončarić. Detection of exudates in fundus photographs using deep neural networks and anatomical landmark detection fusion. *Computer methods and programs in biomedicine*, 137:281–292, 2016.

[42] Oscar Perdomo, John Arevalo, and Fabio A González. Convolutional network to detect exudates in eye fundus images of diabetic subjects. In *12th International Symposium on Medical Information Processing and Analysis*, volume 10160, page 101600T. International Society for Optics and Photonics, 2017.

[43] Alan D Fleming, Sam Philip, Keith A Goatman, John A Olson, and Peter F Sharp. Automated microaneurysm detection using local contrast normalization and local vessel detection. *IEEE transactions on medical imaging*, 25(9):1223–1232, 2006.

[44] Atsushi Mizutani, Chisako Muramatsu, Yuji Hatanaka, Shinsuke Suemori, Takeshi Hara, and Hiroshi Fujita. Automated microaneurysm detection method based on double ring filter in retinal fundus images. In *Medical Imaging 2009: Computer-Aided Diagnosis*, volume 7260, page 72601N. International Society for Optics and Photonics, 2009.

[45] Sherif Abdelazeem. Micro-aneurysm detection using vessels removal and circular hough transform. In *Radio Science Conference, 2002.(NRSC 2002). Proceedings of the Nineteenth National*, pages 421–426. IEEE, 2002.

[46] M Usman Akram, Shehzad Khalid, and Shoab A Khan. Identification and classification of microaneurysms for early detection of diabetic retinopathy. *Pattern Recognition*, 46(1):107–116, 2013.

[47] Akara Sopharak, Bunyarit Uyyanonvara, Sarah Barman, et al. Automatic microaneurysm detection from non-dilated diabetic retinopathy retinal images using mathematical morphology methods. *IAENG International Journal of Computer Science*, 38(3):295–301, 2011.

[48] Li Tang, Meindert Niemeijer, Joseph M Reinhardt, Mona K Garvin, and Michael D Abràmoff. Splat feature classification with application to retinal hemorrhage detection in fundus images. *IEEE Transactions on Medical Imaging*, 32(2):364–375, 2013.

[49] Meindert Niemeijer, Bram Van Ginneken, Joes Staal, Maria SA Suttorp-Schulten, and Michael D Abràmoff. Automatic detection of red lesions in digital color fundus photographs. *IEEE Transactions on medical imaging*, 24(5):584–592, 2005.

[50] Mark JJP van Grinsven, Bram van Ginneken, Carel B Hoyng, Thomas Theelen, and Clara I Sánchez. Fast convolutional neural network training using selective data sampling: application to hemorrhage detection in color fundus images. *IEEE transactions on medical imaging*, 35(5):1273–1284, 2016.

[51] Varun Gulshan, Lily Peng, Marc Coram, Martin C Stumpe, Derek Wu, Arunachalam Narayanaswamy, Subhashini Venugopalan, Kasumi Widner, Tom Madams, Jorge Cuadros, et al. Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *Jama*, 316(22):2402–2410, 2016.

[52] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016.

[53] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014. URL: `http://arxiv.org/abs/1409.1556`.

[54] *The Ubiquitous Lab, Marquette University, USA*. Website: `http://ubicomp.mscs.mu.edu`. Last accessed 11:51 am, January 17, 2019.

[55] Website: `https://sites.google.com/view/ishtiyaqueahmad/datasets`. Last accessed 11:54 am, January 17, 2019.

[56] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016.

[57] *ImageNet*. Website: `http://www.image-net.org/`. Last accessed 11:57 am, January 17, 2019.

[58] Joseph Redmon. Darknet: Open source neural networks in c. `http://pjreddie.com/darknet/`, 2013–2016.

[59] Nobuyuki Otsu. A threshold selection method from gray-level histograms. *IEEE transactions on systems, man, and cybernetics*, 9(1):62–66, 1979.

[60] *STARE (STructured Analysis of the Retina)*. Website: `http://cecas.clemson.edu/~ahoover/stare/`. Last accessed 11:51 am, September 17, 2018.

[61] J.J. Staal, M.D. Abramoff, M. Niemeijer, M.A. Viergever, and B. van Ginneken. Ridge based vessel segmentation in color images of the retina. *IEEE Transactions on Medical Imaging*, 23(4):501–509, 2004.

[62] Randitha Bandara and Manjula Giragama. A retinal image enhancement technique for blood vessel segmentation algorithm. 2017.

[63] François Chollet et al. Keras. `https://keras.io`, 2015.

[64] Theano Development Team. Theano: A Python framework for fast computation of mathematical expressions. *arXiv e-prints*, abs/1605.02688, 2016.

[65] Prasanna Porwal, Samiksha Pachade, Ravi Kamble, Manesh Kokare, Girish Deshmukh, Vivek Sahasrabuddhe, and Fabrice Meriaudeau. Indian diabetic retinopathy image dataset (IDRiD): A database for diabetic retinopathy screening research., 2018.