

Automated Feature Extraction for Early Detection of Diabetic Retinopathy in Fundus Images

Anonymous CVPR submission

Paper ID ****

Abstract

Automated detection of lesions in retinal images can assist in early diagnosis and screening of a common disease: Diabetic Retinopathy. A robust and computationally efficient approach for the localization of the different features and lesions in a fundus retinal image is presented in this paper. Since many features have common intensity properties, geometric features and correlations are used to distinguish between them. We propose a new constraint for optic disk detection where we first detect the major blood vessels first and use the intersection of these to find the approximate location of the optic disk. This is further localized using color properties. We also show that many of the features such as the blood vessels, exudates and microaneurysms and hemorrhages can be detected quite accurately using different morphological operations applied appropriately. Extensive evaluation of the algorithm on a database of 516 images with varied contrast, illumination and disease stages yields 97.1% success rate for optic disk localization, a sensitivity and specificity of 95.7% and 94.2% respectively for exudate detection and 95.1% and 90.5% for microaneurysm/hemorrhage detection. These compare very favorably with existing systems and promise real deployment of these systems.

1. Introduction

Diabetic retinopathy (DR) is a common retinal complication associated with diabetes. It is a major cause of blindness in both middle and advanced age groups. According to the National Diabetes Information data (US)¹, a total of 20.8 million people i.e. 7 percent of the US population have diabetes out of which only 14.6 million cases are diagnosed. Early detection of the disease via regular screening is particularly important to prevent vision loss. In this process, an automated DR diagnostic system can assist in a big

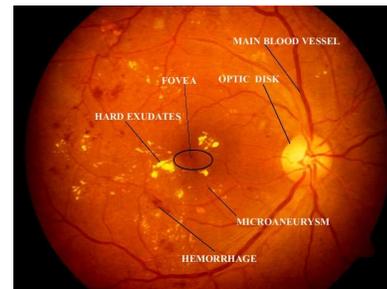


Figure 1. Illustration of various features on a typical retinopathic image.

way since a large population has to be screened and that too repeatedly.

Color fundus images are used by ophthalmologists to study eye diseases like diabetic retinopathy. Figure 1 shows a typical retinal image labeled with various feature components of Diabetic Retinopathy. Microaneurysms are small saccular pouches caused by local distension of capillary walls and appear as small red dots[1]. This may also lead to big blood clots called hemorrhages. Hard exudates are yellow lipid deposits which appear as bright yellow lesions. The bright circular region from where the blood vessels emanate is called the optic disk. The fovea defines the center of the retina, and is the region of highest visual acuity. The spatial distribution of exudates and microaneurysms and hemorrhages, especially in relation to the fovea can be used to determine the severity of diabetic retinopathy.

1.1. Related Work

Sinthaniyothin [12] uses maximum variance to obtain the optic disk center and a region growing segmentation method to obtain the exudates. [4] tracks the optic disk through a pyramidal decomposition and obtains disk localization from a template-based matching that uses the Hausdorff distance measure on the binary edge image. However, the above methods will fail if exudates similar in brightness and size to the optic disk are present.

¹<http://www.diabetes.niddk.nih.gov/dm/pubs/statistics>

054
055
056
057
058
059
060
061
062
063
064
065
066
067
068
069
070
071
072
073
074
075
076
077
078
079
080
081
082
083
084
085
086
087
088
089
090
091
092
093
094
095
096
097
098
099
100
101
102
103
104
105
106
107

[1, 7] used blood vessel intersection property to obtain the optic disk. However, they use the whole blood vessel network which can lead to wrong or inconclusive results because of noise from the fringe blood vessels. In contrast, we use only the main blood vessels, which is more robust.

Statistical classification techniques have been very popular lately for the problem of lesion classification. Exudates have color properties similar to the optic disk while Microaneurysms are difficult to segment due to their similarity in color and proximity with blood vessels. In order to classify detected features, typically, candidate regions are detected using color/morphological techniques and then classification is done on these regions using some classifier. Many classifiers have been tried including Neural Networks [12], PCA [9], Fuzzy C-means clustering [10], SVMs ([16],[2], [15]) and simple Bayesian classification ([15], [13]).

STARE is a complete system for various retinal diseases [6]. The optic disk is detected using blood vessel convergence and high intensity property. In order to determine the features and classification method to be used for a given lesion, a Bayesian probabilistic system is used.

In this paper, we develop methods to automatically detect all of these features in a fundus image using image processing techniques. We show that many of the features such as the blood vessels, exudates and microaneurysms and hemorrhages can be detected quite accurately using different morphological operations applied appropriately. Blood vessels of different thicknesses can be extracted using open and close operations. Exudates appear as bright patches with sharp edges in retinal images and can be extracted using open and close operations using filters of different sizes. Microaneurysms and Hemorrhages (MAHMs) are segmented using morphological filters that exploit their local 'dark patch' property. These are further classified as lesion/non-lesion using a color model extracted from the blood vessels. We propose a new constraint for optic disk detection where we first detect the major blood vessels and then use the intersection of these to find the approximate location of the optic disk. This is further localized using color properties. Detection of the Optic disk, fovea and the blood vessels is used not only for distinguishing them from lesions but also for extracting color information for better lesion detection.

The rest of the paper is organized as follows: Section 2 details the blood vessel extraction algorithm while section 3 elucidates the exudate detection procedure; Section 4 describes the optic disk detection method and section 5 presents the method for detection of fovea, microaneurysms and hemorrhages. Section 6 describes determining the severity of the disease using lesion detection while in section 7, the results of the algorithm over an extensive dataset are presented.

2. Multi-Scale Blood Vessel Extraction

In our approach, color images input from the fundus camera are initially resized to a standard size of 768×576 pixels while maintaining the original aspect ratio. We select the green channel for all our operations because retinal images are almost always saturated in the red channel and have very low contrast in the blue channel.

A closing operation is performed on the green channel image using two different sizes of a structuring element (filter). Closing operation is defined as dilation (Max filter) followed by erosion (Min filter). The formulations of dilation and erosion for gray scale images are as follows.

Dilation:

$$A \oplus B = A_1(x, y) = \max_{i, j \in B} (A(x - i, y - j) + B(i, j))$$

Erosion:

$$A \ominus B = A_2(x, y) = \min_{i, j \in B_1} (A(x - i, y - j) + B_1(i, j))$$

where A is the input image, B and B_1 are the structuring elements or masks used for dilation and erosion respectively.

Dilation in gray scale enlarges brighter regions and closes small dark regions. The erosion is necessary to shrink the dilated objects back to their original size and shape. The dark regions closed by dilation do not respond to erosion. Thus, the vessels being thin dark segments laid out on a brighter background are closed by such a closing operation. A subtraction of the closed images across two different scales (let S_1 and S_2 be the sizes of the structuring elements B_1 and B_2) will thus give the blood vessel segments of the green channel image. The operation is as follows:

$$C' = (A \oplus B_2) \ominus B_2 - (A \oplus B_1) \ominus B_1$$

We use a disk shaped structuring element for morphological operations. The radius of the larger disk (S_2) is fixed at a high value (we use 6 pixels for an image of size 768×576 pixels) so that all the vessels including the main blood vessel get closed. S_1 is chosen adaptively as follows:

1. 1 or 2 pixels below S_2 if we want to obtain only the thicker vessels emanating from the optic disk.
2. Atleast 4 pixels below S_2 to obtain the entire blood vessel network.

Criterion 1 is used for optic disk localization whereas criterion 2 is used in microaneurysms and hemorrhages detection. The image C' is thresholded (90% of the maximum intensity) and median filtered to obtain the binary image of the blood vessels (U). Morphological thinning is then performed on U to obtain the skeleton of the blood

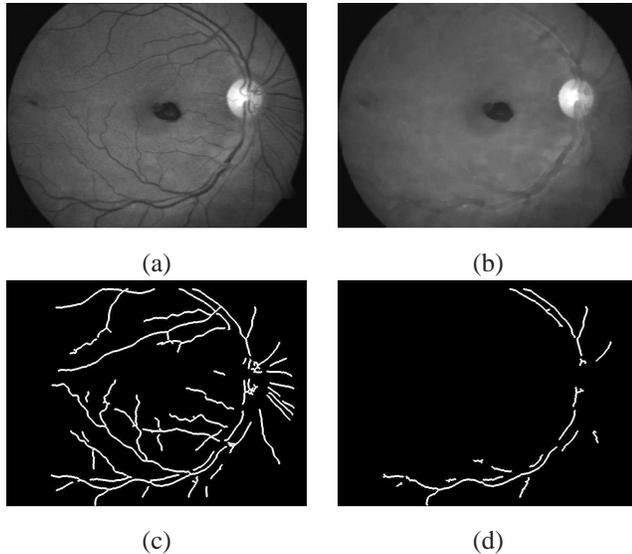


Figure 2. Illustration of the multi-scale vessel extraction algorithm on an image of our dataset: (a) The original green channel image, (b) The green image dilated at scale S_2 , (c) The vessel network obtained at lower scale S_1 (criterion 2), (d) The thicker vessels obtained using higher S_1 (criterion 1).

vessel network. Thinning operation is implemented as $U - (U \ominus B_1 - \overline{U} \ominus B_2)$, where B_1 and B_2 are disjoint structuring elements and \overline{U} is the complement of the image U . Noise can occur in the thinned image usually in the form of dots. A 2×2 median filtering operation is performed to remove the isolated specks of noise. The vessel segments being connected structures are unaffected by this operation. An additional source of noise in retinopathic images could be exudates, the removal of which is detailed in section 6.

Figure 2 shows the results of the vessel extraction algorithm on an image having no exudates. Figure 2 (c) showing the entire vessel network was obtained using criterion 2 while Figure 2 (d) having only the thicker vessels arising from the optic disk was obtained using criterion 1.

3. Exudate Localization and Detection

Exudates appear as bright lesions in retinopathic images and have sharp edges and high contrast with the background. Most of the standard edge detectors like Sobel and Canny add a lot of noise and miss out key edges when used for extracting exudate edges and hence are not suitable for this application. We perform boundary detection for exudates using morphological operations.

Dilation is performed on the green channel (obtained after pre-processing without equalization and normalization) at 2 different scales: S_3 and S_4 , both of which are greater than S_2 which was used for vessel extraction. Hence, at both S_3 and S_4 , the blood vessels do not appear in the di-

lated result. The exudates being bright with sharp edges respond to dilation. Subtraction of the results across the 2 scales gives the boundaries of the exudates: $P = (A \oplus B_4) - (A \oplus B_3)$. The image P is thresholded in intensity to obtain the binary boundaries. The threshold is chosen as α times the maximum intensity in P where α is obtained by training. Hard exudates give closed boundaries in the thresholded result. Short breaks in the contours are connected by smoothing splines. This bridging of short breaks in boundaries is useful for extracting softer exudates. A morphological filling operation is then used to search for regions bounded by closed contours in the result. It is defined as follows:

$$E^k = (E^{k-1} \oplus B) \cap H^c \quad (1)$$

This is iterated starting from $k = 1$ until $E^k = E^{k-1}$. Here, H^c is the complement of the thresholded binary image, B is a 4-connected structuring element and E^0 is an image containing a seed point inside the boundary. In the case of multiple holes, this operation is done in each hole until all the holes are filled.

Morphological filling operation on the binary image thus gives us the candidate exudate patches. However, the candidate regions may contain artifacts. Therefore, a linear classifier is built which uses the brightness and edge properties of exudates. Exudates are bright yellow or white in color and have high intensity in the green channel. We localize the exudate patches more accurately by taking all the candidate regions whose mean intensities in the green channel are greater than a fraction β (obtained by training) of the maximum intensity in the channel. For classifying the patches based on their edge strength, the gradient magnitude image of the green channel is chosen. This gradient magnitude image is thresholded (the absolute threshold γ obtained by training) and the number of white pixels in the thresholded image for each exudate patch is counted. We denote this as the gradient count of each patch. Patches which do not have sufficient gradient count (δ) are discarded.

Patches that satisfy both the brightness criterion and gradient count are retained. In each of the patches classified as exudates, the exact lesion boundary is tracked starting from the pixel with the highest gradient magnitude and completing the contour based on continuity of gradient magnitude and direction. Pixels in the interior of these contours are then accurately classified as exudate pixels. The optic disk which may invariably appear in the result is masked out using the procedure outlined in the next section.

The various thresholds like α , β etc. were obtained by training the exudate algorithm on a set of 50 color fundus images obtained from ophthalmologists. Ground truth images were obtained for this set with exudates pre-marked and the parameter values that gave optimal results (in terms of detecting true exudates and not detecting artifacts) were chosen. The values we obtained are $\alpha = 0.06$, $\beta = 0.52$,

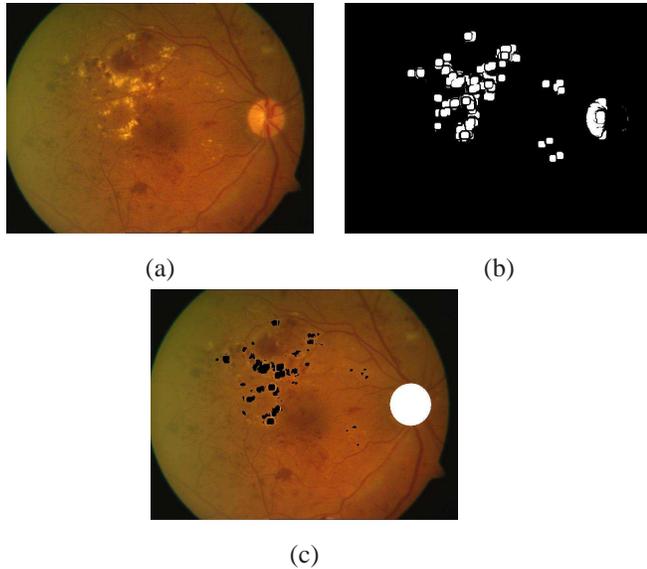


Figure 3. Illustration of the exudate algorithm on an image of our dataset: (a) The original image, (b) Localized candidate regions, (c) Final result with exudate pixels in black color and optic disk masked out.

$\gamma = 3$ and $\delta = 4$.

Figure 3 shows the steps of the exudate detection algorithm on an image of our dataset. The candidate regions obtained after morphological segmentation are shown in Figure 3 (b) and the result after classification, pixel identification and optic disk removal is shown in Figure 3 (c).

4. Optic Disk Detection

The detection of the optic disk in fundus images is a very important task because of its similarity in brightness, color and contrast to the exudates. It invariably appears in exudate detection results and hence there is a need to mask it out. Moreover, the optic disk is an important retinal feature and can be used for registration of retinal images. It can also be used to diagnose other diseases like Glaucoma.

In this paper, we use a constraint that has been overlooked so far which is to detect the optic disk using convergence of only the thicker blood vessels arising from it. This significantly improves the performance compared to existing techniques that use the entire vessel network. Also, this thicker vessel convergence is almost always present in the image as opposed to other features of the optic disk such as color or circular shape. This approach is combined with the high intensity property of disk regions in a cost function to improve the robustness of optic disk detection compared to existing methods.

4.1. Exudate Subtraction

The vessel extraction algorithm detailed in section 3 (using criterion 1) yields the skeleton of the thicker blood vessels. These 'thicker' vessels include the main blood vessel and other smaller but thick vessels emanating from the optic disk. Exudates have the potential of occurring as noise in the vessel segmentation because closing and subtracting the irregularly shaped exudates using regular structuring elements is not an exact operation (i.e. the shape of exudates does not remain the same after the closing operation). Hence, we perform an additional step to overcome this limitation. The result of the exudate detection algorithm is dilated and subtracted from the blood vessel result. This removes any noise due to exudates in the extraction of the main blood vessel. The resulting image after exudate subtraction is then processed to obtain the optic disk center.

4.2. Vessel Convergence

The segments of the thicker blood vessel skeleton are modeled as lines. We transform the vessel image into the Hough space using the Hough Transform (HT) $(x, y) \xrightarrow{HT} (r, \theta)$ to obtain the lines. The dataset of lines thus generated is reduced by eliminating lines with slopes $\theta < 45^\circ$. This can be done as the vessels converging at the optic disk typically have a high slope in the vicinity of the disk. This reduced dataset of lines is intersected pairwise to generate an intersection map. Lines close to each other and with nearly the same slope are not intersected due to a higher triangulation error.

Weighted Convergence Optimization: The intersection map generated from candidate line segments of the thicker vessels is used to find the location of the optic disk. The map is dilated to make the region of convergence more apparent.

Dilation and erosion in binary images are implemented as OR and AND filters (extension of Max and Min filters in gray scale). Let the dilated intersection image be M and the pre-processed green channel image be A . Then, we define a weighted image (J) as follows:

$$J = M + wA$$

where w is obtained as follows:

$$w = \begin{cases} 1, & \text{if } N \leq N_0 \\ (N_0/N)^\beta, & \text{if } N > N_0 \end{cases} \quad (2)$$

where N is the number of high intensity pixels ($Intensity > 200$) in the initial green channel image (before pre-processing), N_0 is the number of pixels corresponding to the size of a normal disk (taken from [12]), β is a power law factor (taken as 2) which rapidly decreases w

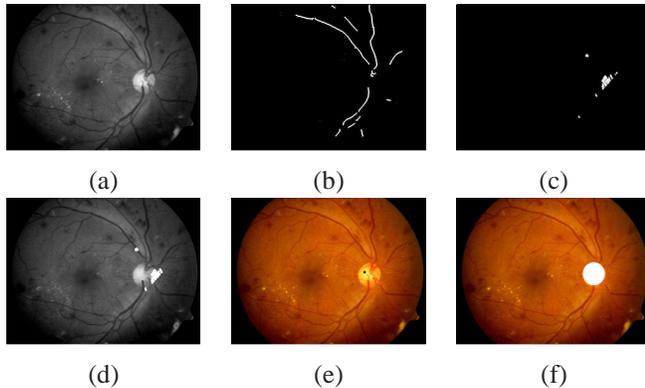


Figure 4. Illustration of the steps of the optic disk algorithm on an image of our dataset: (a) The green channel image, (b) The major vessels obtained using criterion 1, (c) The dilated intersection map, (d) The image J obtained by weighted summation, (e) Final disk center result indicated as a black spot, (f) The disk mask obtained.

as N increases above N_0 . Using this formula, the image A is given less weightage if the number of bright pixels (N) in the green channel is more than a threshold N_0 (sign of bright exudates); else A is given a weightage of 1 (normal images). Image J is then used to obtain the optimal location of the optic disk using a cost function F .

$$F = \sum_{i,j \in W} J(x-i, y-j) \quad (3)$$

where (x, y) is a point in the intersection map and summation is over a circular sliding window (W). F can be computed efficiently using dynamic programming/integral images. The point (x, y) which maximizes F , is taken as the location of the optic disk.

Figure 4 shows the steps of the optic disk localization algorithm on an image of our dataset. The vessels of Figure 4 (b) are obtained after exudate subtraction.

5. Detection of Microaneurysms and Hemorrhages

Microaneurysms are small blood clots which occur due to capillary burst. They are the hardest to detect in retinopathic images. Hemorrhages are bigger clots. Microaneurysms and Hemorrhages (MAHMs) are treated as holes (i.e. small dark blobs surrounded by brighter regions) and morphological filling is performed on the green channel (without pre-processing) to identify them. The filling operation in gray scale is an extension of binary filling used in section 5. The unfilled green channel image is then subtracted from the filled one and thresholded in intensity to yield an image (R) with microaneurysm patches. The threshold (ν) is chosen based on the mean intensity of the

retinal image in the red channel. For a mean intensity of 127 (taken as reference on the intensity scale of 0-255), ν was chosen as 7 (obtained by training). The threshold is incremented/decremented by 1 for every 20 units increase/decrease in the mean.

Blood vessels can also appear as noise in the microaneurysm and hemorrhage detection as they have similar color and contrast to the clots. To remove this additional noise, the full blood vessel network skeleton (section 4, Criterion 2) is first obtained. The resulting blood vessel network is dilated and subtracted from the image R to remove the noise due to vessels. The remaining patches are further classified using intensity properties and a color model based on the detected blood vessels.

5.1. Blood Vessel based Color Model

Microaneurysms and hemorrhages have similar color properties as the blood vessels. Hence, a color model is built for classification of these lesions. We use the blood vessel result obtained before the thinning operation is performed (i.e. image U in section 4). For every candidate microaneurysm/hemorrhage patch, the blood vessel segments in a local neighborhood are obtained. The mean (μ) and standard deviation (σ) of these vessels in the red and green channels are calculated. The pixels in each localized patch whose intensities in the red and green channels are in the range $[\mu - 1.2\sigma, \mu + 1.2\sigma]$ are retained. This helps in the removal of background artifacts.

5.2. Detection and Removal of Fovea

The fovea is a dark region located in the center of the macula region of the retina. It commonly appears in microaneurysm and hemorrhage detection results much as the optic disk does in exudate detection results. The fovea is detected using the location of the optic disk and curvature of the main blood vessel. The main blood vessel is obtained as the thickest and largest blood vessel emanating from the optic disk. The entire course of the main blood vessel is obtained (from the image of the thicker vessels) by looking for its continuity from the optic disk. This vessel is modeled as a parabola ([9]). The vertex of the parabola is taken as the pixel on the main blood vessel that is closest to the center of the optic disk circular mask. The fovea is located approximately between 2 to 3 optic disk diameter (ODD) distance from the vertex, along the main axis of the modeled parabola and is taken as the darkest pixel in this region. The region of the fovea is taken to be 1 optic disk diameter of the detected fovea location and detection of microaneurysms in this region is suppressed.

Figure 5 shows the steps of the MAHMs algorithm on an image of our dataset. We use a size threshold to distinguish between microaneurysm and hemorrhage patches.

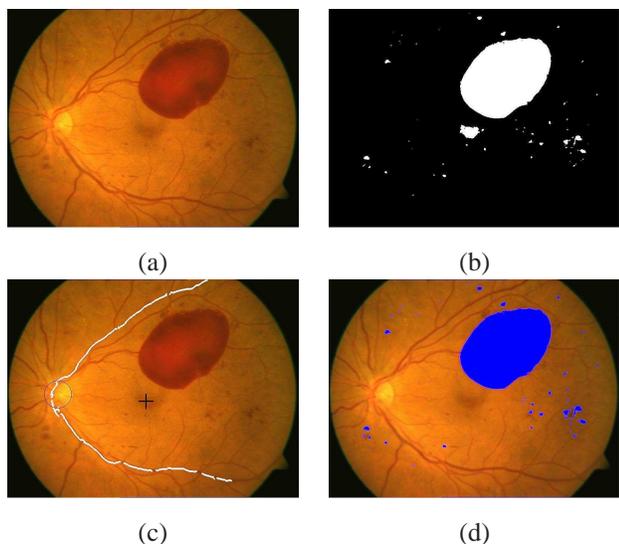


Figure 5. Illustration of the steps of our Microaneurysms and Hemorrhages detection algorithm: (a) An image from our dataset, (b) The image R obtained after morphological filling and thresholding, and removing blood vessels and pixels not satisfying the color model, (c) Fovea detection marked with '+', the main blood vessel and optic disk mask boundary are also shown, (d) The final result after removing fovea and classification, MAHMs marked in blue.

6. Predicting the Severity of Disease

The distribution of the exudate lesions about the fovea can be used to predict the severity of Diabetic macular edema. [9] has divided the retinal image into 10 sub-regions about the fovea. The exudates occurring in the macular region are more dangerous and require immediate medical attention than the ones farther away. Similarly, the size, count and distribution of microaneurysms and hemorrhages is also used to predict the severity of DR. The region around the optic disk is divided into four quadrants for this purpose. The International Council of Ophthalmology² lists 5 levels for Diabetic Retinopathy based on these criteria: none, mild, moderate, severe, and proliferative. Our system uses these criteria in order to classify each image in these categories. Figure 5 contains a very large hemorrhage and many other smaller hemorrhages and microaneurysms and is a case of proliferative Diabetic Retinopathy.

7. Results of Experiments and Discussion

We used a dataset of 516 images for evaluating the algorithm. The images were obtained from diverse sources and hence have sufficient variations in color, illumination and quality. The various sources of the images are as follows: 211 images were obtained from ophthalmologists at 2 local eye hospitals; 81 images were taken from the STARE

²<http://www.icoph.org/standards/pdrclass.html>

Table 1. Results of optic disk localization for specific databases and for the overall normal, abnormal cases.

Source or Type	No. Images	No. Correct	% Success
Hospitals	211	209	99.1
STARE	81	76	93.8
DRIVE	40	40	100
Diaretdb0	130	127	97.7
Red Atlas	54	49	90.7
Overall Normal	112	111	99.1
Overall Abnormal	404	390	96.5

database [7]; 130 images were from the Diaretdb0 database [8]; 40 images were from the DRIVE database [14] and 54 were from the Red Atlas database³.

The images in the dataset were classified by ophthalmologists based on the lesion type (exudates/MAHMs) into those with the lesion and those without it. An image having no lesions is considered normal whereas one that has lesions like exudates, microaneurysms and hemorrhages is considered abnormal. Among the 211 images obtained from eye hospitals, 29 are normal and 182 are abnormal. The DRIVE database contains 33 normal and 7 abnormal images. The STARE database has 30 images which are normal and the other 51 being abnormal. The diaretdb0 database consisting of 130 images has 20 normal and 110 abnormal ones. All the images taken from the Red Atlas database were abnormal.

Of the 516 images, 345 were identified by ophthalmologists as having exudates and 171 did not have any exudates. A total of 348 images were identified as containing microaneurysms/hemorrhages while 168 were free from this lesion type. The entire algorithm was run on the database and results for optic disk localization, exudate detection and MAHMs detection were obtained. The matlab code takes 20 seconds per image on an average to run on a 2 GHz machine with 448 MB RAM.

Table 1 summarizes the results of our optic disk detection algorithm on the dataset. The pixel location obtained prior to placing the circular mask is considered for evaluation purposes. If the pixel location indicated by the optic disk algorithm falls within the boundaries of the optic disk, then it is considered a correct detection and vice-versa. The results in Table 1 are classified based on the source and also based on the nature of the images (normal or abnormal).

In Table 2, the overall optic disk detection rate of our algorithm is presented along with other results from literature. Gagnon et al. [4] obtain 100% optic disk detection but their dataset consists of only normal images without retinopathy. Adam Hoover et. al. [7] evaluate their optic disk detection algorithm exclusively on the STARE database. We obtain better results on the STARE database as indicated by table

³<http://redatlas.org/main.htm>

Table 2. Comparison of results for optic disk localization

Author	No. of Images	% Success
Chutatape et al. [9]	35	99
Gagnon et al. [4]	40	100
Adam Hoover et al. [7]	81	89
Our Method	516	97.1

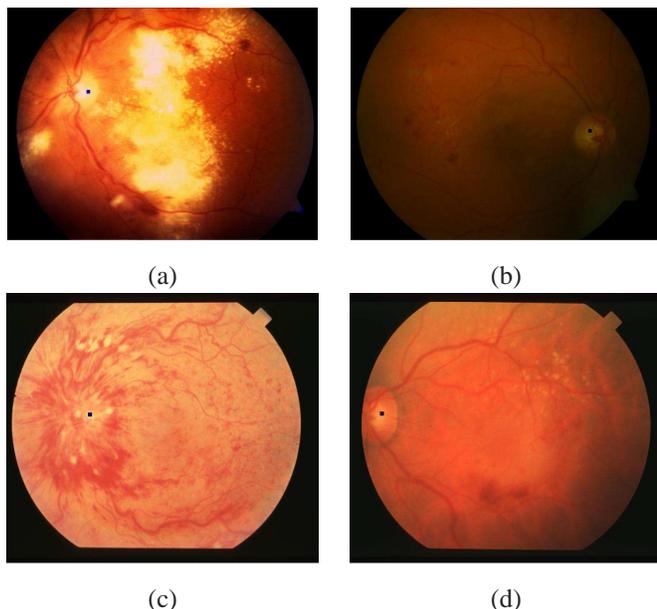


Figure 6. Optic disk localization (without mask) results indicated as black spot on some images of the dataset: (a) Image with big, bright exudates and bright optic disk, (b) An image with non uniform illumination from the diaretdb0 database, (c) An image with lots of blood clots from the STARE database, (d) Another image from the STARE database.

1. Figure 6 shows the optic disk localization results (prior to the masking step) on some images of the dataset. Figures 6 (c) and 6 (d) show 2 cases where our algorithm performs better than [7] on the STARE database.

We obtain sensitivity and specificity results for our algorithm at both the image and pixel levels. For the pixel level results, we obtained ground truth images from the ophthalmologists with lesion (both exudates and MAHMs) pixels marked for all images except those from the diaretdb0 database which already has the ground truth. The results of the algorithm were evaluated with respect to the groundtruth results to determine pixel level sensitivity and sepecificity for each lesion type. The pixel level results were averaged across all images.

In Table 3, the sensitivity and specificity results of our exudate detection algorithm are presented for the 516 images along with other results from the literature. In Table 4, the results for our microaneurysm and hemorrhage detection method are presented along with other results from the

Table 3. Comparison of Results for exudate detection: NI-Number of Images used, NIEx-Number of Images with Exudates, SN-Sensitivity, SP-Specificity (ND*-data not given in the paper)

Author	NI	NIEx	SN(%)	SP(%)
Chutatape [9]	35	28	100	71.0
Sinthanayothin [12]	30	21	88.5	99.7
Sanchez et al.[11]	20	10	100	90.0
Wang et al.[15]	154	54	100	70.0
Garcia et al.[5](image)	50	25	100	84.0
Garcia et al.[5](lesion)	50	25	84.4	62.7
Fleming et al.[3]	13219	300	95.0	84.6
Sopharak et al.[13]	10	ND*	93.38	98.14
Our approach(image)	516	345	95.7	94.2
Our approach(pixel)	516	345	94.6	91.1

Table 4. Comparison of Results for MAHMs Detection: NI-Number of Images used, NIMAHMs-Number of Images with MAHMs, SN-Sensitivity, SP-Specificity

Author	NI	NIMAHMs	SN(%)	SP(%)
Sinthanayothin [12]	30	14	77.5	88.7
Our result (Image)	516	348	95.1	90.5
Our result (Pixel)	516	348	92.0	90.1

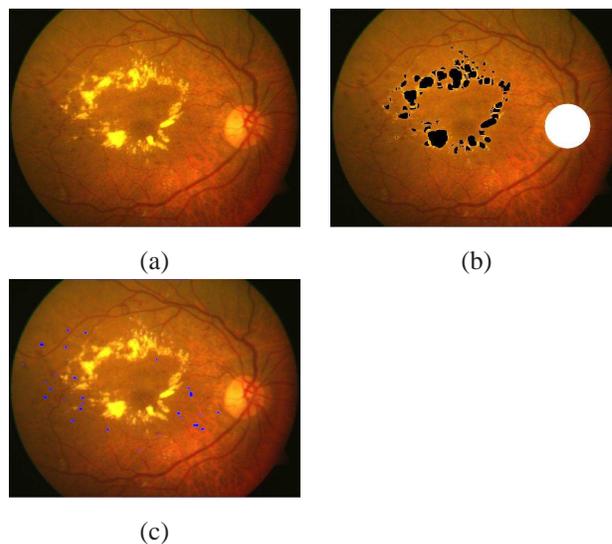


Figure 7. Our exudate and MAHMs detection results (indicated in black and blue color respectively) on an image of the dataset with the optic disk removed using a white mask in the exudate result: (a) The original Image with hard exudates and many microaneurysms, (b) Our exudate detection result, (c) Our MAHMs detection result (the result has been dilated to make the MAHM patches visible).

literature. Figure 7 shows the exudate and MAHM detection results on an image of our dataset. Image 7 (a) contains hard exudates and lots of microaneurysms/hemorrhages which are detected by our algorithm as shown in Figures

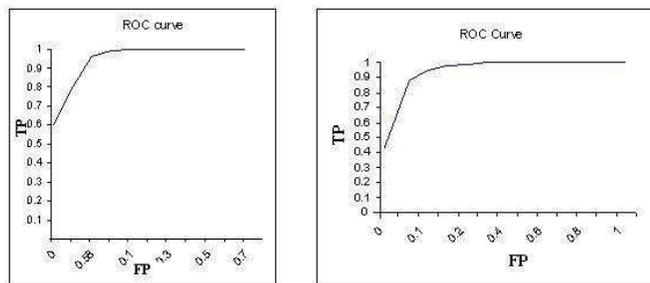


Figure 8. Plot of ROC curves for various lesion types: (a) The ROC curve for exudates, (b) The ROC curve for MAHMs.

7 (b) and 7 (c) respectively. Some of the exudates in Figure 7 (b) are clustered in the macula and hence need immediate medical attention. Based on the size, count and distribution of MAHMs, our algorithm correctly predicted the level of severity of diabetic retinopathy for this case as moderate.

Figures 8 (a) and 8 (b) show the ROC curves of our algorithm for exudates and MAHMs respectively. The curves are obtained by varying various parameter thresholds (α , β , γ , δ for exudates and ν , ζ for MAHMs respectively) while running the algorithm on the database and plotting the True Positive rate (TP i.e. sensitivity) versus the False Positive rate (FP i.e. 1-True Negative rate). Here, the TP and FP are obtained at the image level for each lesion type. The true positive rates reach close to 100% for both exudates and MAHMs at very low false positive rates, thus demonstrating the robustness of our algorithm for lesion detection.

8. Conclusion

In this paper, an efficient framework for early detection of Diabetic Retinopathy has been developed. The optic disk is tracked by combining the blood vessel convergence and high disk intensity properties in a cost function. We show that, as opposed to most methods that use learning techniques, geometrical relationships of different features and lesions can be used along with simple morphological operations in order to obtain a very robust system for analysis of retinal images. Our techniques may further be combined with some learning methods for possibly even better results.

References

- [1] K. Akita and H. Kuga. A computer method of understanding ocular fundus images. *Pattern Recognition*, 15(6):431–443, 1982. 1, 2
- [2] O. Chutatape and X. Zhang. Detection and classification of bright lesions in color fundus images. In *International Conference on Image Processing*, pages 139–142, 2004. 2
- [3] A. D. Fleming, S. Philip, K. A. Goatman, G. J. Williams, J. A. Olson, and P. F. Sharp. Automated detection of exudates for di-

- abetic retinopathy screening. *Physics in Medicine and Biology*, 52:7385–7396, Dec. 2007. 7
- [4] L. Gagnon, M. Lalonde, M. Beaulieu, and M.-C. Boucher. Procedure to detect anatomical structures in optical fundus images. In *Proc. SPIE Medical Imaging: Image Processing*, pages 1218–1225, 2001. 1, 6, 7
- [5] M. Garcia, R. Hornero, C. Sanchez, M. Lopez, and A. Diez. Feature extraction and selection for the automatic detection of hard exudates in retinal images. *Engineering in Medicine and Biology Society, 2007. EMBS 2007. 29th Annual International Conference of the IEEE*, pages 4969–4972, Aug. 2007. 7
- [6] M. Goldbaum, S. Moezzi, A. Taylor, S. Chatterjee, J. Boyd, E. Hunter, and R. Jain. Automated diagnosis and image understanding with object extraction, object classification, and inferring in retinal images. *International Conference on Image Processing*, 3:695–698, Sept. 1996. 2
- [7] A. Hoover and M. Goldbaum. Locating the optic nerve in a retinal image using the fuzzy convergence of the blood vessels. *IEEE Trans. on Medical Imaging*, 22:951–958, Aug. 2003. 2, 6, 7
- [8] T. Kauppi, V. Kalesnykiene, J. K. Kamarainen, L. Lensu, I. Sorri, H. Uusitalo, H. Kalviainen, and J. Pietila. Diaretddb: Evaluation database and methodology for diabetic retinopathy algorithms. Technical report, Lappeenranta University of Technology, Lappeenranta, Finland, 2006. 6
- [9] H. Li and O. Chutatape. A model-based approach for automated feature extraction in fundus images. In *Proc. IEEE International Conference on Computer Vision*, page 394, 2003. 2, 5, 6, 7
- [10] A. Osareh. *Automated Identification of Diabetic Retinal Exudates and the Optic Disc*. PhD thesis, Univ. of Bristol, Jan. 2004. 2
- [11] C. I. Sanchez, A. Mayo, M. Garcia, M. I. Lopez, and R. Hornero. Automatic image processing algorithm to detect hard exudates based on mixture models. *Engineering in Medicine and Biology Society, 2006. EMBS '06. 28th Annual International Conference of the IEEE*, pages 4453–4456, Sept. 2006. 7
- [12] C. Sinthanayothin, J. F. Boyce, T. H. Williamson, H. L. Cook, E. Mensah, S. Lal, and D. Usher. Automated detection of diabetic retinopathy on digital fundus images. *Diabetic Medicine*, 19:105–112, 2002. 1, 2, 4, 7
- [13] A. Sopharak, K. Thet Nwe, Y. A. Moe, M. N. Dailey, and B. Uyyanonvara. Automatic exudate detection with a naive bayes classifier. In *International Conference on Embedded Systems and Intelligent Technology (ICESIT)*, pages 139–142, Feb. 2008. 2, 7
- [14] J. Staal, M. D. Abramoff, M. Niemeijer, M. A. Viergever, and B. V. Ginneken. Ridge-based vessel segmentation in color images of the retina. *IEEE Trans. on Medical Imaging*, 23(4):501–509, 2004. 6
- [15] H. Wang, W. Hsu, G. K. G., and L. M. L. An effective approach to detect lesions in color retinal images. In *In Proc. Conf. Comp. Vision Pattern Rec.*, pages II: 181–186, 2000. 2, 7
- [16] X. Zhang and O. Chutatape. Top-down and bottom-up strategies in lesion detection of background diabetic retinopathy. In *In Proc. Conf. Comp. Vision Pattern Rec.*, pages 422–428, 2005. 2