

The Role of Domain Knowledge in the Detection of Retinal Hard Exudates

Wynne Hsu^{1,2} P M D S Pallawala¹ Mong Li Lee¹ Kah-Guan Au Eong³

¹School of Computing, National University of Singapore
{whsu, pererapa, leeml}@comp.nus.edu.sg

²Singapore-MIT Alliance, National University of Singapore
4 Engineering Drive 3, Singapore 117576
whsu@sma.nus.edu.sg

³Department of Ophthalmology, Tan Tock Seng Hospital, Singapore
kah_guan_au_eong@ttsh.com.sg

Abstract

Diabetic retinopathy is a major cause of blindness in the world. Regular screening and timely intervention can halt or reverse the progression of this disease. Digital retinal imaging technologies have become an integral part of eye screening programs worldwide due to their greater accuracy and repeatability in staging diabetic retinopathy. These screening programs produce an enormous number of retinal images since diabetic patients typically have both their eyes examined at least once a year. Automated detection of retinal lesions can reduce the workload and increase the efficiency of doctors and other eye-care personnel reading the retinal images and facilitate the follow-up management of diabetic patients. Existing techniques to detect retinal lesions are neither adaptable nor sufficiently sensitive and specific for real-life screening application. In this paper, we demonstrate the role of domain knowledge in improving the accuracy and robustness of detection of hard exudates in retinal images. Experiments on 543 consecutive retinal images of diabetic patients indicate that we are able to achieve 100% sensitivity and 74% specificity in the detection of hard exudates.

1. Introduction

Diabetic retinopathy is a serious complication of diabetes mellitus and a major cause of blindness worldwide [12]. Early detection and timely treatment of diabetic retinopathy can halt or reverse the progression of the disease and prevent blindness. For this reason, almost all diabetic patients should have their retina in both eyes examined at least once every year even if they have no visual symptom. A fundus camera is used to capture images of the retina that are then read and graded by doctors. Given the number of diabetic patients screened yearly, the number of retinal images generated is large and the majority of them are normal. Automated detection

of abnormal retinal images can reduce the workload of doctors reading the images and improve the follow-up management of diabetic patients.

One important manifestation of diabetic retinopathy is the development of hard exudates in the retina. Hard exudates occurring in the macula can cause significant visual impairment. Existing algorithms to detect hard exudates [3,6,14] are not sufficiently sensitive and specific for clinical application. The main obstacle is the extreme variability in the color of the retinal images that depends on factors such as degree of pigmentation in the retinal pigment epithelium and choroid in the eye, size of the pupil, and illumination. These factors, in turn, affect the appearance of hard exudates in the retinal images.

In this paper, we use domain knowledge to improve the reliability of automated hard exudate detection. Domain knowledge of light- or bright-colored lesions such as drusen and cotton wool spots is used to differentiate true hard exudates from these other lesions. High intensity artifacts near large retinal blood vessels as a result of reflection of light are often mistaken as hard exudates. Domain knowledge of the locations of retinal blood vessels is used to remove such artifacts. In addition, image quality has a great impact on the features of retinal lesions, especially hard exudates. Knowledge of how the size and intensity difference of hard exudates vary according to the degree of overexposure and underexposure is used to identify the often missed hard exudates in overexposed and underexposed retinal images. Experiments on 543 consecutive digital retinal images of diabetic patients indicate that our approach is able to achieve 100 % sensitivity and 74% specificity.

The rest of the paper is organized as follows. Section 2 gives the background for our work. Section 3 presents the proposed knowledge based hard exudate detection algorithm. Section 4 shows the results of our experiments. Section 5 compares our approach with a recent work on retinal image analysis. Finally, we conclude in Section 6.

2. Preliminaries

In this section, we will first give the clinical domain knowledge for the development of certain lesions in retinal images. We will then review existing hard exudate detection techniques and highlight their shortcomings.

2.1 Lesions in Retinal Images

Diabetic retinopathy is characterized by microvascular leakage and microvascular occlusion in the retina. Weakening of retinal capillary walls and loss of pericyte support, give rise to microaneurysm formation. Serum lipoproteins leak from these microaneurysms and are deposited as hard exudates in the retina. Hard exudates vary in size and have a yellow waxy appearance with distinct margins. Microvascular occlusion in diabetic retinopathy leads to infarction of the retinal nerve fibers. This forms white fluffy opaque lesions known as cotton wool spots along the nerve fiber layer.

Drusen are small, discrete, yellow-white, slightly elevated lesions which are usually an early manifestation of age-related macular degeneration but are unrelated to diabetic retinopathy. Drusen consists of focal collections of hyaline material located between the basal lamina of the retinal pigment epithelium and the inner collagenous layer of Bruch's membrane. Drusen are less yellow in color compared to hard exudates. We will show in the subsequent section how this domain knowledge is used to differentiate hard exudates from other lesions.

2.2 Related Work

A number of studies have been carried out to automatically detect hard exudates based on their size, shape, texture, etc. [11] shows that using these features in isolation is insufficient to detect hard exudates accurately. Several attempts have been made to detect hard exudates using histogram segmentation. If the background color of a retinal image is sufficiently uniform, a simple and effective method to separate exudates from the background is to select proper thresholds [10,15,18]. [11,14] investigate dynamic thresholding by dividing a retinal image into sub-images and computing a threshold based on the local histogram. This technique is useful only if the sub-images contains unimodal hard exudates.

[6, 9] examine how the 3-D color feature space can be mapped into 3-D spherical coordinates. Unfortunately, hard exudates, drusen, and cotton wool spots overlap in this feature space and are not linearly separable. [6] attempts to use Mahalanobis classifier and jackknife linear discriminant methods to detect exudates but the results were inconclusive. [3] applies restrictions on the exudates and test images and uses the Mahalanobis classifier with

features extracted from shape, color and edge of the hard exudates to achieve high sensitivity and specificity.

[17] employs the minimum distance discriminant classification to detect lesions in retinal images. But it is not robust for different types of retinal images and user intervention is required on a case-by-case basis. The approach is not scalable for analyzing large numbers of retinal images in a real-life screening application.

3. Domain Knowledge Based Hard Exudate Detection

We will now present a domain knowledge based approach to detect hard exudates in digital retinal images. The retinal images are obtained from a Topcon digital fundus camera and stored in 24-bit TIFF format with 768x576 pixels resolution. This consists of 8-bits red, green and blue layers with 256 levels each. Figure 1 shows examples of digital retinal images with different background colors. There are three steps in the hard exudate detection algorithm:

- (1) Use median filtering to compute intensity difference map.
- (2) Use dynamic clustering to determine lesion clusters.
- (3) Apply domain knowledge to identify true hard exudates.

We will explain each step in the following subsections.

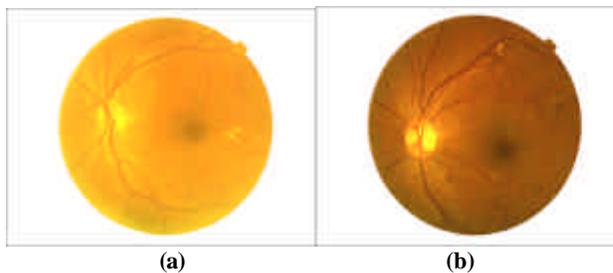


Figure 1. Digital retinal images with (a) bright yellowish background and (b) with dark reddish background.

3.1 Compute Intensity Difference Map

The retinal lesions of interest in our current study are manifested as bright intensity regions in the retinal images. [13] shows that the green layer contains the most information on the brightness and structure of exudates compared to the other two layers. An optic disc recognition algorithm [5] is first used to remove the optic disc from the retinal images. Smoothing operation is applied on the green layer to remove random noises. Next, median filtering [1,7] is applied to the smoothed image. Figure 2 shows the results of median filtering on the retinal images in Figure 1. To highlight the lesions areas, the median filtered image is subtracted from the green layer of the original smoothed digital retinal image to

obtain a difference map. Figure 3 shows the difference maps computed for Figure 1.

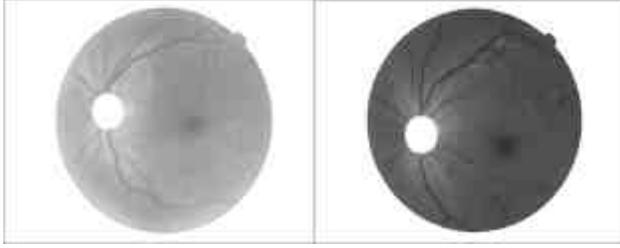


Figure 2. Median filtered images of Figure 1.

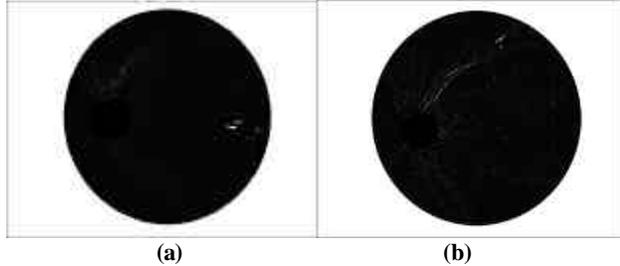


Figure 3. Difference map of Figure 1.

From the output, we see that median filtering and subsequent subtraction is effective in isolating the lesion areas. We conducted experiments to determine the optimal size for the median filter. Table 1 shows the results. We observe that a median filter of size 17x17 is sufficient to give 100% sensitivity. Hence this filter is selected for this algorithm.

| Filter Size | Sensitivity |
|-------------|-------------|
| 13x13 | 87% |
| 15x15 | 93% |
| 17x17 | 100% |
| 19x19 | 100% |

Table 1. Comparison of various median filters.

We can distinguish two types of objects in the difference map: the brighter lesion patches, and the darker background features. These form separate clusters in the brightness space of the difference map. The exact extent or location of each group in brightness space cannot be pre-determined due to the non-uniform qualities of input retinal images. But there is always a clear demarcation between the two groups.

3.2 Determine Lesion Clusters

Dynamic clustering techniques [2,4] have been used to automatically and accurately determine clusters without prior thresholds or input parameters. This technique separates a set of data points into reasonable groups or clusters according to the attributes of data. Here, “a cluster” means the set of data within the cluster are more

similar to each other than the data from other clusters in specific aspects/features. Each cluster has its own “cluster center” in the feature space. An important measurement of similarity for data is “distance” between cluster centers and between points inside one cluster. In our case, the feature space is the difference map. The distance measure is the Minkowski metrics defined by

$$d_{ij}(X_i, X_j, r) = \left\{ \sum_{k=1}^n |X_{ik} - X_{jk}|^r \right\}^{1/r}$$

where n is dimension of feature space; r is optional parameter; d_{ij} is distance between pixel X_i and X_j in n -dimensional feature space. When $r=1$, this distance becomes “city block” metric; and when $r=2$, it denotes the Euclidean distance, etc. In our application, the distance measure is simply the difference in the gray values between pixel i and pixel j in the difference map.

We need to pre-determine the number of clusters and provide the initial cluster centers [16]. Since we are looking for lesions versus non-lesions, the number of clusters is two. The lesion cluster is located at higher intensity range while the background cluster is located at lower intensity range. The maximum and minimum gray levels in the difference map are calculated and are termed *max* and *min* respectively. These values are used as the initial cluster centers $CLSTR_{lesion}$ and $CLSTR_{others}$.

The entire clustering process is iterative. Let q be the clustering iteration number, $CLSTR(q)_{lesion}$ and $CLSTR(q)_{others}$ be the cluster centers in iteration q . Then the clustering algorithm proceeds as follows:

- (1) Calculate the *max* and *min* values of a $M_h \times M_v$ pixel difference map.
- (2) Set $q = 1$, $CLSTR(q)_{lesion} = max$, $CLSTR(q)_{others} = min$.
- (3) For $i = 1, 2, \dots, M_h \times M_v$,
 Let the intensity level of pixel i be X_i
 $D1 = distance(X_i, CLSTR(q)_{lesion})$
 $D2 = distance(X_i, CLSTR(q)_{others})$
 If $D1 < D2$
 Then assign pixel i to the *lesion* cluster
 Else assign pixel i to *others* cluster
- (4) Update cluster centers according to the equations:

$$CLSTR(q+1)_{lesion} = \frac{1}{|CLSTR(q)_{lesion}|} \sum_{i \in lesion} X_i$$

$$CLSTR(q+1)_{others} = \frac{1}{|CLSTR(q)_{others}|} \sum_{i \in others} X_i$$
- (5) Increment the iteration and repeat Steps (3) and (4) until stopping conditions are met.

Two common stopping conditions are:

- (a) The cluster centers stay within a limited range between iteration q and iteration $q + \Delta q$ ($\Delta q > 1$), i.e., the cluster centers do not vary much from one iteration to another;
- (b) Number of iteration exceeds a user-defined loop threshold say T_{loop} .

Initial investigations show that dynamic clustering was able to successfully detect all exudates, but it is unable to differentiate true hard exudates from artifacts and other lesions. In fact, a significant number of artifacts that appear along the large retinal vessels are detected as hard exudates. These artifacts were primarily from the glistening background very close to the large vessels. Figure 4 show the results of the clustering for Figure 1.

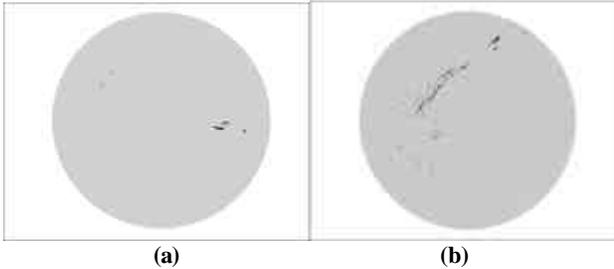


Figure 4. Clustered image of Figure 1 containing hard exudates and artifacts.

3.3 Apply Domain Knowledge

3.3.1 Knowledge of retinal vessels

In order to improve the accuracy of our detection algorithm, we incorporate knowledge about blood vessels into the detection algorithm to minimize false identifications of artifacts near the vessels as exudates. A vessel recognition algorithm is used to determine major blood vessels in the retinal images. Experiments show that the vessels are mostly distributed in the lower intensity levels of the green layer. A reverse video operation is performed to bring the vessels into the high intensity levels. Next, a median filtering technique using a 20x20 median filter is applied to suppress the vessels information while preserving the background information. To highlight the blood vessels, we subtract the median filtered image from the original inverted image. This produces an image that contains vessels in a uniform background. Figure 5 shows the results.

Dynamic clustering is used to differentiate the vessels from the background. Clustering was done with two initial cluster centers: one for vessels and the other for background. The cluster centers are set after repeated experiments to determine to the best initial centers values. The experiment was done on 210 digital retinal images of varying qualities. All the vessels in a given retinal image are classified into large vessels, primary vessels and secondary vessels. Large vessels are the vessels that emerge from the optic disc. Primary vessels are formed by complete division of large vessels or branches of large vessels that are more than 3 pixels in diameter. Secondary vessels are the branches of the primary vessels or branches of large vessels that are less than 3 pixels in diameter. Table 2 summarizes the results.

We selected the initial cluster center for the vessel to be Mean - SD and the initial cluster center for the background to be Mean + SD. Figure 6 shows the results of applying dynamic clustering with the selected initial cluster centers. Note that a large number of vessel fragments is produced. The following heuristic is used to link vessel fragments: Vessel fragments are connected to the closest edge of the another vessel fragment if the distance between two fragments are less than 3 pixels and the gradient of the fragment vessels are compatible to be a part of a completed vessel. These constructed vessels are then connected to the large vessels if they are 3 pixels or less apart. Any unconnected regions were removed from the image. Figure 7 shows the completed retinal vessel structure. Armed with this knowledge of the location of retinal vessels, we are able to eliminate a large number of false identification of hard exudates due to artifacts.

| | Accuracy (%) $INITIALCENTRE_{vessel} = Mean - 2*SD$ $INITIALCENTRE_{others} = Mean + SD$ | Accuracy (%) $INITIALCENTRE_{vessel} = Mean - SD$ $INITIALCENTRE_{others} = Mean + SD$ | Accuracy (%) $INITIALCENTRE_{vessel} = Mean - SD/2$ $INITIALCENTRE_{others} = Mean + SD$ |
|-------------------|--|--|--|
| Large vessels | 93.65% | 99.24% | 78.41% |
| Primary vessels | 64.51% | 84.57% | 68.8% |
| Secondary vessels | 36.41% | 60.72% | 56.45% |

Table 2. Results of the vessel recognition algorithm



Figure 5. Vessel information highlighted image.



Figure 6. Clustered image.



Figure 7. Completed retinal vessel structure.

3.3.2 Knowledge of Drusen and Cotton Wool Spots

To further improve the detection accuracy, we study the characteristics of drusen, cotton wool spots and hard exudates. The 3-D spherical coordinates [8] in the equations below give the best results in differentiating hard exudates from drusen and cotton wool spots:

$$r = (R^2 + G^2 + B^2)^{1/2}$$

$$\theta = \text{Arctan}(G/R)$$

$$\varphi = \text{Arccos}(B/r)$$

Let r_c be the mean r value of all the pixels belonging to lesion cluster c and r_{bc} be the mean r value of all the pixels belonging to the immediate background of c . We define the immediate background of c to be an area that is 7 times the size of c , with c at the center of this area. Since hard exudates have a much higher intensity contrast than cotton wool spots or drusen, this contrast information is captured by $r_{\text{difference}}$ where $r_{\text{difference}} = r_c - r_{bc}$. Figure 8 shows the distribution of $r_{\text{difference}}$ for the hard exudates and other lesions e.g. drusen and cotton wool spots. There are some overlaps among the three types of lesions primarily due to the overexposed and underexposed digital retinal images.

A study of the image quality reveals that if the average intensity of the original retinal image is above 110 or below 70, the quality of the image deteriorates and the lesions features alters. We use this information to differentiate between true hard exudates from drusen and cotton wool spots. In good quality images where average intensity between 70 and 110, if the value of $r_{\text{difference}}$ is 25 or more, we classify it as hard exudates. In poor quality images where average intensity above 110 or below 70, if the value of $r_{\text{difference}}$ is above 13, we will classify it as hard exudates. Figure 9 shows the results.

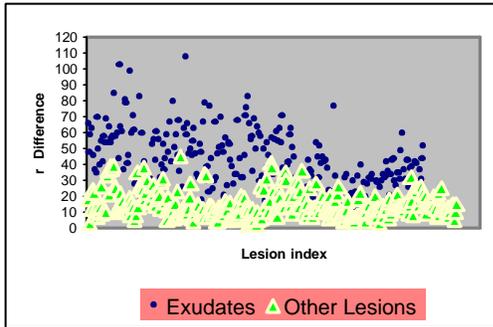


Figure 8. $r_{\text{difference}}$ for hard exudates and other lesions.

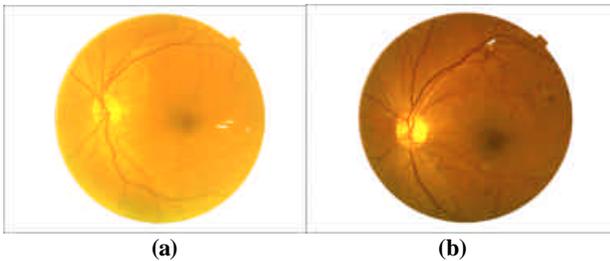


Figure 9. Hard exudates detected for Figure 1.

4. Experiment Results

We evaluate the robustness and reliability of our detection algorithm on 543 consecutive retinal images of diabetic patients who attended a routine diabetic retinopathy screening program. These images contain artifacts and varying combinations of one or more retinal lesions including hard exudates, drusen, cotton wool spots, hemorrhages and microaneurysms. Clinically, the retinal images range from normal to advanced or proliferative diabetic retinopathy. Some of the images are blurred because of the coexistence of cataract in the eye. The quality of these retinal images varies from very poor to good. A retinal specialist first reviews these images and identifies the locations and types of lesions present in the images. We compare the results of our hard exudate detection algorithm with that given by the retinal specialist. Two accuracy measures are used:

- Sensitivity* - Ratio of number of images detected to be abnormal to total number of abnormal images.
- Specificity* - Ratio of number of images detected to be normal to total number of normal images.

Table 3 shows the accuracy of the algorithm in identifying hard exudates without the use of domain knowledge. Table 4 shows the accuracy of the algorithm when domain knowledge is applied. Our hard exudate detection algorithm is able to achieve 100% sensitivity with or without domain knowledge. With the help of domain knowledge, the specificity improves from 41% (215/512) to 74% (380/512). This translates to a 33% savings without compromising the sensitivity.

| | | Automated Detection of Hard Exudate | |
|--------------------|----------------------|-------------------------------------|---------|
| | | Absent | Present |
| Retinal Specialist | Hard Exudate Absent | 215 | 297 |
| | Hard Exudate Present | 0 | 31 |

Table 3. Hard exudate detection without domain knowledge.

| | | Automated Detection of Hard Exudate | |
|--------------------|----------------------|-------------------------------------|---------|
| | | Absent | Present |
| Retinal Specialist | Hard Exudate Absent | 380 | 132 |
| | Hard Exudate Present | 0 | 31 |

Table 4. Hard exudate detection with domain knowledge.

5. Discussion

The most similar work to ours that has been reported is by B. M. Ege, Ole K. Hejlesen, and Ole V. Larsen [3]. They show promising results in detecting hard exudates using Mahalanobis classifier and statistical classification. However, we are unable to obtain either their test images

or their executable version for comparison. Hence, we can only analyze their work qualitatively to demonstrate the differences between our work and theirs.

In their algorithm, they restrict the hard exudate area to a maximum of 300 pixels and pre-define the hard exudate diameter to be between 1-6 pixels. From our experiments, it was noted that such criteria are unacceptable as there are a number of hard exudates in our retinal images that are much larger than the specified restrictions. Furthermore, it is difficult to define the diameters of hard exudates since their shapes are highly variable. The algorithm uses thresholding in preprocessing the images. Area growing technique is then applied to estimate the shape of the lesion following thresholding. We implemented a simple version of the algorithm and our results indicate that the area growing technique (even with very small threshold) tends to grow rapidly and gives a wrong estimate of the exudate size and shape. Finally, in their work, only 143 test images are used and these images are selected using several criteria (e.g. lesions are not connected to other structures, lesions must be compatible with early diabetic retinopathy, images without fibrotic lesions etc.). In our work, we tested our algorithm on 543 consecutive retinal images of diabetic patients who attended a routine diabetic retinopathy screening program. No restriction has been applied in the selection of test images. The result truly demonstrates the robustness and reliability of our algorithm.

6. Conclusion

In this paper, we have proposed a novel hard exudate detection algorithm that incorporates relevant domain knowledge to achieve the robustness and reliability needed for clinical application. We have shown that the algorithm is able to achieve 100% sensitivity and 74% specificity in the detection of hard exudates in digital retinal images derived from a routine diabetic retinopathy screening program. In our future work, we will investigate how we can integrate further domain knowledge to help identify other abnormal lesions such as hemorrhage and neovascularisation. The ultimate aim is to build a comprehensive retinal image analysis system that will facilitate the screening of diabetic patients at risk of blindness from diabetic retinopathy and improve of the management of diabetic retinopathy.

Acknowledgements

We thank Wang Huan for her implementation of an earlier version of the algorithm. This project is supported by National University of Singapore academic research project (RP991613).

References

- [1] J. T. Astola and T.G. Campbell. On Computation of the Running Median, in *IEEE Trans. Acoustics, Speech, and Signal Processing*, ASSP-37(4), 572-574, 1989.
- [2] New Approaches in Classification and Data Analysis, edited by H.H. Bock, et al Springer-Verlag, 1994.
- [3] B. M. Ege, O. K. Hejlesen, O. V. Larsen, K. Moller, et al. Screening for Diabetic Retinopathy Using Computer Based Image Analysis and Statistical Classification, in *Computer Methods and Programs in Biomedicine*, Vol. 62, 165-175, 2000.
- [4] B. Everitt. *Cluster Analysis*, in Heinemann Educational Books Ltd., UK, 1974.
- [5] K. G. Goh, W. Hsu, M. L. Lee, ADRIS: An Automatic Diabetic Retinal Image Screening System, *Medical Data Mining and Knowledge Discovery*, Springer-Verlag. 2000.
- [6] M. H. Goldbaum, N. P. Katz, M. Nelson. The Discrimination of Similarly Colored Objects in Computer Images of the Ocular Fundus, in *Investigative Ophthalmology and Visual Sciences*, Vol. 31, 617-623, 1990.
- [7] T. S. Huang, G. T. Yang, Q. Y. Tang. A Fast Two-Dimensional Median Filtering Algorithm, in *IEEE Trans. Acoustics, Speech, and Signal Processing*, ASSP-27(1), 13-18, 1979.
- [8] J. K. Kasson, W. Plouffe. An Analysis of Selected Computer Interchange Color Spaces, in *ACM Trans. on Graphics*, Vol. 11, 373-405, 1992.
- [9] N. Katz, M. Goldbaum, M. Nelson, S. Chaudhuri. An Image Processing System for Automatic Retina Diagnosis, in *Three-Dimensional Imaging and Remote Sensing Imaging, SPIE*, Vol. 902, 131-137, 1988.
- [10] S. C. Lee, Y. Wang. A General Algorithm for Recognizing Small, Vague, and Imager-Alike Objects in a Nonuniformly Illuminated Medical Diagnostic Image, in *32nd Asilomar Conference on Signals, Systems & Computers*, Vol. 2, 941-943, 1998.
- [11] L. Leistriz, D. Schweitzer. Automated Detection and Quantification of Exudates in Retinal Images, in *SPIE*, Vol. 2298, 690-696, 1994.
- [12] K.H. Lim. Registration of New Blindness in Singapore for 1985-1995, in *Singapore Medical Journal*, Vol. 40, No.2, Feb. 1999.
- [13] Z. Liu, C. Opas, S.M. Krishnan. Automatic Image Analysis of Fundus Photograph, in *Proceedings of IEEE International Conference on Engineering in Medicine and Biology*, Vol. 2, 524-525, 1997.
- [14] R. Philips, J. Forrester, P. Sharp. Automated Detection and Quantification of Retinal Exudates, in *Graefe's Archive for Clinical & Experimental Ophthalmology*, Vol. 231, 90-94, 1993.
- [15] R. P. Philips, T. Spencer, P. G. B. Ross, P. F. Sharp, J. V. Forrester. Quantification of Diabetic Maculopathy by Digital Imaging of the Fundus, in *Eye*, Vol. 5, 130-137, 1991.
- [16] H. Spath. *Cluster Dissection and Analysis*, Ellis Horwood Limited, 1985.
- [17] H. Wang, W. Hsu, K. G. Goh, M. L. Lee. An Effective Approach to Detect Lesions in Color Retinal Images, in *IEEE Conference on Computer Vision and Pattern Recognition*, South Carolina, USA, June 2000.
- [18] N. P. Ward, S. Tomlinson, et al. Image Analysis of Fundus Photographs, in *Ophthalmology*, Vol. 96, No.1, Jan. 1989.