

Detection of Age-related Macular Degeneration (AMD) in Human Eye's Retina Using Image Processing and Analysis Methods

TECHNICAL UNIVERSITY OF CRETE

Department of Electronics Engineering and Computer Engineering

Konstantinos E. Rapantzikos

April 2000

Accepted on the recommendation of

Prof. Dr. M. Zervakhs, thesis advisor

Prof. Dr. E. Christodoulou, co-examiner

As. Prof. Dr. E. Petrakhs, co-examiner

Acknowledgments

I wish to extend my sincerest thanks to my advisor Dr. Michalis Zervakis for giving me the opportunity to work under his guidance

To my parents for their love and affection

To Dr. Costas Balas for providing me with helpful information concerning my case study

To everyone else for helping me complete this work

KONSTANTINOS E. RAPANTZIKOS

Contents

1. INTRODUCTION	5
1.1 Age-Related Macular Degeneration	5
1.2 Medical Examination	5
1.3 Image Acquisition	7
1.4 Problem Definition:	7
1.5 Related Research	9
2. CRITICAL OVERVIEW OF ALGORITHMS	11
2.1 The Illumination Model and its Compensation	11
2.2 Contrast Enhancement Techniques	18
2.2.1 Global vs. Local Enhancement	18
2.2.2 Local Enhancement Techniques	19
2.2.3 Effectiveness of Local Enhancement Operators	25
2.3 Thresholding Techniques	29
2.3.1 Otsu Algorithm	30
2.3.2 Thresholding According to Distribution's Curvature	31
2.3.3 Relaxation Method	32
2.3.4 Histogram Local Adaptive Thresholding (HALT):	34
2.3.5 Results of Thresholding Schemes	35
3. MAIN ALGORITHM FOR DRUSEN DETECTION	38
3.1 Illumination Model Compensation	38
3.2 Enhancement	38
3.3 First Level Thresholding	39
3.4 Local Thresholding and Elimination of Sparse Pixels	39
3.5 Expansion	44
3.6 Overall Algorithm	44
4. EXAMPLES	45
4.1 Illumination Model Compensation	45
4.2 Enhancement and Global Thresholding	46
4.3 Morphological Dilation, HALT and Median Filtering	47
4.4 Other Examples	48
4.5 General conclusion on the application	51
5. CONCLUSION AND FUTURE WORK	52
APPENDIX A HISTOGRAM PROPERTIES	55
Features for Histogram Analysis	55
APPENDIX B OTHER DETECTION APPROACHES	58
B.1 Detection Based on Gradient Information	58
B.2 Otsu thresholding scheme	62
B.3 Template matching and Curvature estimation	63

<i>APPENDIX C NUMERICAL FEATURES FOR CLASSIFICATION OF DRUSEN</i>	<i>65</i>
REFERENCES	66

1. INTRODUCTION

1.1 Age-Related Macular Degeneration

Age-related macular degeneration (AMD) is a disease that causes progressive damage to the macula, a small, specialized part of the eye that allows us to see fine details clearly. When the macula malfunctions, people experience blurring or darkness in the center of their vision and tasks such as reading and driving are affected. Some common ways vision-loss is detected are when words on a page look blurred, a dark or empty area appears in the center of vision, or straight lines look distorted.

There are two forms of AMD including dry (also called atrophic, non-neovascular, or nonexudative) and wet (also called exudative). Dry AMD is the more common form of the disease and accounts for 90% of all AMD. The key identifier for dry AMD is small, round, white-yellow deposits called drusen that build up in the macula. The dry form currently cannot be treated with medication or surgery, but magnifying and telescopic lenses may be used to take advantage of most of the remaining eyesight. Vitamins and supplements may be helpful in slowing the progress of AMD.

The dry form may or may not progress to the wet form. The wet form is less common but more severe than the dry form. It accounts for approximately 10% of all AMD but 90% of all blindness from the disease. This form is characterized by choroidal neovascularization (CNV), the development of abnormal blood vessels beneath the retinal pigment epithelium (RPE) layer of the retina. These vessels can bleed and cause macular scarring which can result in profound loss of central vision. Currently, laser treatment can be performed to stop the blood-vessel growth. This treatment does not restore vision, but it may minimize or delay the loss of eyesight. Since laser photocoagulation can damage healthy surrounding tissue, treatment is typically not performed until the disease has progressed to a stage where the laser treatment causes less damage than the disease.

AMD is the leading cause of irreversible vision loss in people over 65 in the U.S. Although the cause of AMD is not completely understood, it has been identified that age is the greatest risk factor and there is a hereditary associated with the disease. AMD occurs more often in Caucasians than in Hispanics or African-Americans. In addition, there is a higher incidence in females than males.

1.2 Medical Examination

A thorough examination by an eye doctor is the best way to determine if one has macular degeneration, or if he/she is at risk for developing the condition. The typical eye exam evaluates the health of eyes through a number of tests.

The exam begins by testing the visual acuity or the sharpness of vision. There are several different tests for visual acuity. After these visual tests, the front part of the eyes is examined to determine if everything is healthy. The doctor may put anesthetic

drops in the eyes before measuring the pressure in each eye. Then, drops are administered and cause pupils to dilate. This allows the doctor to examine the retina through the enlarged pupil. The drops typically take between 20 and 45 minutes to work, and wear off in about 4 hours. While the pupils are dilated, it is usually difficult to read, and bright lights may be uncomfortable. Some patients use sunglasses after dilation to reduce light sensitivity.

The slit lamp is a microscope that gives the examiner a magnified view of the retina. Using the slit lamp, the doctor looks for drusen and other areas of the retina that appear suspicious or abnormal. After the dilating drops are administered and generated their effect, the eye doctor seats the patient at the slit lamp, which is a special microscope that enables the doctor to examine the different parts of the eye under magnification. When used with handheld lenses or special contact lenses, the slit lamp gives the examiner a highly magnified view of the retina.

The examiner looks for drusen and other areas of the retina that might appear suspicious or abnormal. Since choroidal neovascularization (the new blood vessel growth found in the "wet" form of macular degeneration) occurs beneath the retina, the blood vessels themselves are not usually visible. But the examination can reveal clues, such as elevation of the retina, or fluids behind the retina, that suggest the presence of choroidal neovascularization (CNV). In these cases, further testing may be necessary.

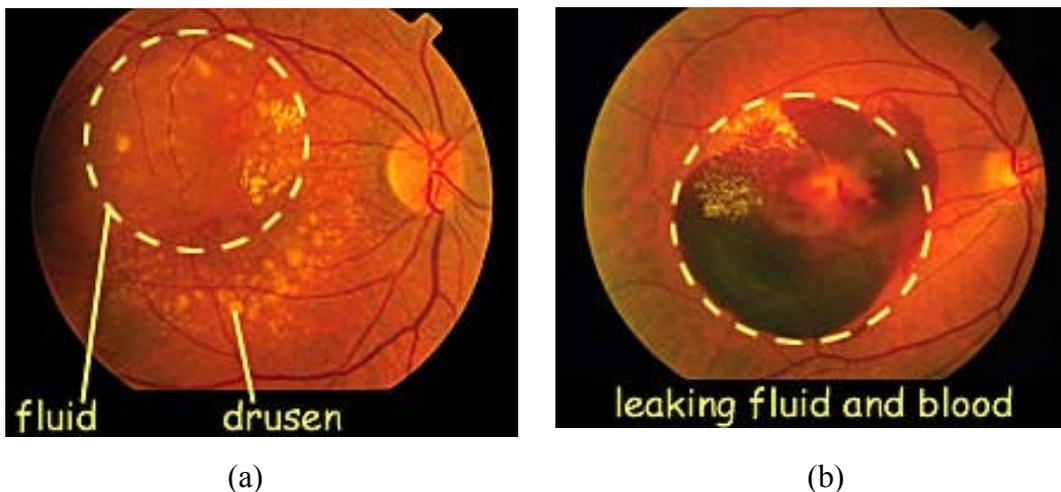


Fig. 1.1 Retinal photographs showing (a) dry and (b) wet form of AMD

Fig. 1.1-(a) shows many drusen and fluid under the retina, in a case that is suggestive of choroidal neovascularization. Fig. 1.1-(b) shows fluid and blood beneath the retina, which suggests the presence of wet form AMD or choroidal neovascularization (CNV). Additional testing will be required for complete diagnosis and treatment.

This examination process is tedious for both the doctor and the patient. Depth vision cameras have been developed to capture the retinal images for accurate off-line analysis. One particularly useful device is the Fundus camera developed in Spectral Imaging Technologies & Biomedical Diagnostics Lab-Institute of Electronic Structure & Laser (I.E.S.L.).

1.3 Image Acquisition

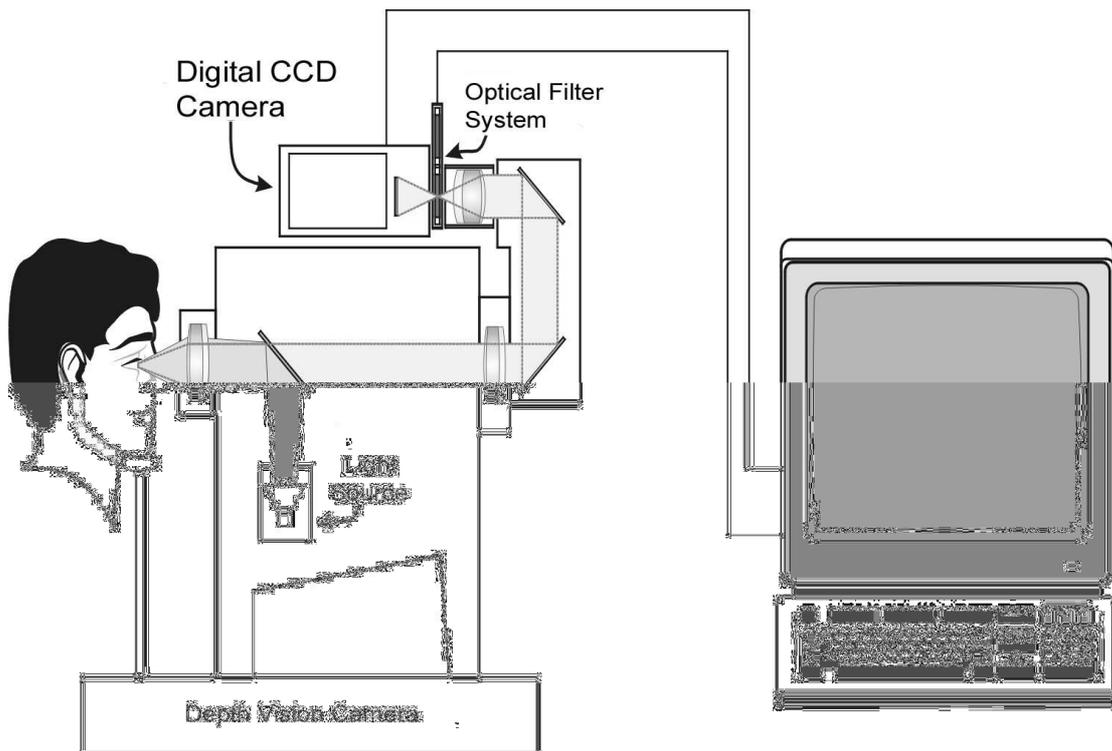


Fig. 1.2 Schematic of Fundus camera

The use of Fundus camera for optical imaging of retina is illustrated in Fig. 1.2. A digital camera is adapted to Fundus camera in combination with optical and digital filters, to capture images from eye's retina.

1.4 Problem Definition:

In this thesis we attempt to develop an automated system for analyzing the images from a depth-vision camera. The goal is to isolate the drusen areas, detect their presence and measure their spatial extent.

In examining for AMD, our region of interest is the nearly central part of the retina and actually the area, which does not contain the optic nerve. Drusen appear in the form of random shaped spots inside the retina. It is very difficult to localize drusen by just looking at a retina's image, because many drusen are vague and "intermixed" with vessels.

The images to be analyzed are complex mainly due to the vessels' and fibres' presence. Moreover, a problem that makes the analysis difficult is the non-uniform illumination caused by the different surface curvature (different light absorption). In brief, the following problems must be considered:

- Non-uniform illumination.
- Vessels interact with drusen.
- Contrast is low \Rightarrow drusen are difficult to be discriminated from background in terms of their intensity.
- Drusen are of different sizes (large and small).
- Small drusen near to each other tend to create larger areas that can be mistaken as large drusen.
- Drusen tend to spread (blur) around their location.

By taking the previous facts under consideration, the following tasks must be performed:

- The non-uniform illumination effect has to be moderated (Fig. 1.3).
- Each image must be further enhanced, in order to create more distance between the intensities of spots and background.
- The resultant image has to be thresholded so that the diseased areas can be accurately isolated.
- Finally, a refinement operation might be necessary for obtaining the whole area of a detected spot.

These issues are examined in the following chapters of the thesis and an algorithm is proposed for the effective detection of drusen.

More specifically the chapters of the thesis are organized as follows. Chapter 2 presents an overview of enhancement and thresholding techniques. Chapter 3 develops and analyzes the proposed approach for drusen detection. Chapter 4 presents examples of the application of the algorithm proposed on actual retinal images. Chapter 5 concludes this thesis with a summary and further directions opened up by this research.

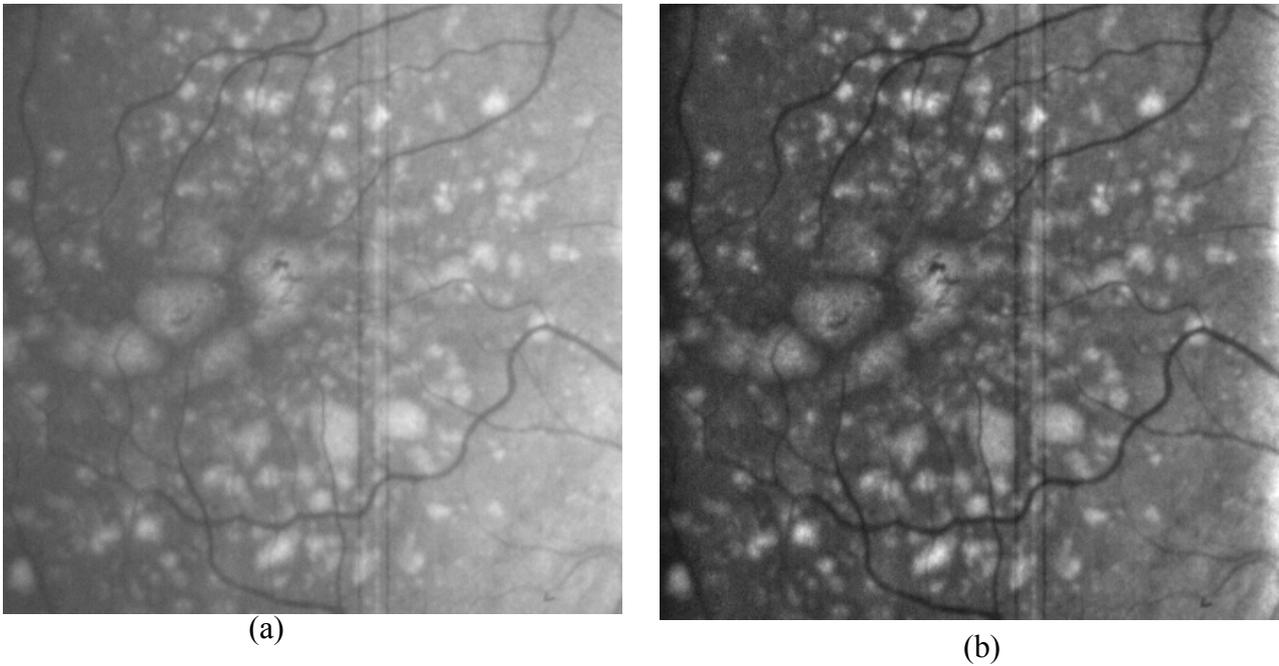


Fig. 1.3 (a), image with non-uniform illumination, (b), image after non-uniform illumination correction

1.5 Related Research

During our search for related problems we came across similar enhancement, thresholding and detection issues. Some of them are reviewed here in relation to our work.

Cork is a natural material produced in the Mediterranean countries. Cork stoppers are used to seal wine bottles. The cork stoppers are grouped into eight classes according to the degree of defects on the cork surface. These defects appear in the form of random-shaped holes, cracks, and others and they are not specifically defined to a particular shape or size. Thus, a “good” enhancement and classification process has to be built in order to reduce the rejection rate [9]. An iterative local enhancement and adaptive thresholding scheme was used in [9] to define and separate cracks from background. Classification is achieved with the use of a fuzzy MLP with a backpropagation-training algorithm.

Breast cancer is a major cause of fatality among all cancers for women. However, the etiologies of this kind of cancer are unclear and no single dominant cause has yet emerged. Early detection is maybe the only way, till now, to start treatment before cancer is spread to other parts of the body. The proposed detector, in [11], consists of image preprocessing, feature extraction, cancer detection and classification (decision). Image preprocessing performs noise reduction by median, band-pass and Gaussian filters, and enhancement of features extracted from the mammogram by band-pass and wavelet filters. A sharpening filter is then applied to maximize the contrast value between the masses and the local background. Mean, variance, skewness, and kurtosis are statistical features used for the detection phase through an RBFNN (Radial Basis Function Neural Network) for classification purposes.

Early detection and removal of skin cancer can also lead to patient’s survival. A segmentation of skin cancer images is proposed in [10], which consists of the

following steps. During the preprocessing phase the color image is transformed in such a way that the intensity of a pixel shows the distance of it from the background. An initial segmentation is obtained by the application of a threshold value determined from the average intensity of high gradient pixels. Shrinkage and expansion of a closed elastic curve, initialized at the approximate boundary of skin cancer, is applied so as to obtain best fit.

Localization of the prostate in ultrasound images is crucial in urology. The objective of the proposed method in [13] is edge detection and localization of prostate in such images. Initially, the image is smoothed and first (edge strength) and second derivatives (edge location) are calculated. An approach that uses large filters in homogeneous areas and smaller ones in areas with some gray level transitions is presented and evaluated. An improvement of the previous method is achieved by adaptive filtering, in [12], using local standard-deviation calculations. Refinement of the edge localization is achieved by adjusting the filter size for areas with a local standard deviation above the average.

A novel automated image analysis system is built, in [14], in order to differentiate immunohistochemically stained cells from background. Initially, global thresholding algorithms are applied to find an approximate threshold at which cells could be separated. A refinement algorithm follows, so as to erode edge pixels of the regions. However, resulting cells' regions are overlapping and a new decomposition method is designed, in order to segment cells correctly.

The characteristics and functionality of blood vessels play an important role in clinical and research studies on a number of diseases. An algorithm for vessel detection is presented in [15] that uses a set of linear filters. These filters are sensitive to vessels of different orientation and thickness. Thresholding with hysteresis, in [16], is used to differentiate actual vessels from noise or misleading regions.

During our research on thresholding techniques, we came across several publications. A class of methods uses entropy's principle for threshold decision. References [17], [18], [19] use two-dimensional entropies, while Wong and Sahoo, in [20], propose a threshold selection method based on the maximum entropy principle. Another class of techniques, in [5,21,22], based directly on the histogram's shape proved to be more suitable to our problem. Otsu proposes a non-parametric and unsupervised method for automatic threshold selection in [5], while two other methods, namely Kittler's - Illingworth's and extension of Lloyd's, are presented in [21] and [22] respectively.

Our first approach on detecting drusen, after enhancing and thresholding the images, was the use of circle/ellipse detectors. Hough Transform, which is described in [23,24], is used for line and curve detection. Variants of HT exist and make use of the edge orientation, the circle radii, a complex accumulator with the phase proportional to the log of radius, etc. A modified CHT (Circular Hough Transform) is presented in [25]. A two-dimensional Hough transform in conjunction with radius histogramming is used for circle recognition in [27]. A new circle/ellipse detector that adopts a hybrid scheme, which consists of a genetic algorithm phase and a local search phase, is proposed in [26].

2. CRITICAL OVERVIEW OF ALGORITHMS

2.1 The Illumination Model and its Compensation

· The illumination-reflectance model:

The term *image* refers to a two-dimensional (2-D) light-intensity function, denoted by $f(x, y)$, where the value or amplitude of f at spatial coordinates (x, y) gives the intensity (brightness) of the image at that point. As light is a form of energy, $f(x, y)$ must be nonzero and finite, that is

$$0 < f(x, y) < \infty. \quad (2.1-1)$$

The images people perceive in everyday visual activities normally consist of light reflected from objects. The basic nature of $f(x, y)$ may be characterized by two components: (1) the amount of source light incident on the scene being viewed and (2) the amount of light reflected by the objects in the scene. Appropriately, they are called the *illumination* and *reflectance components*, and are denoted by $i(x, y)$ and $r(x, y)$, respectively. The functions $i(x, y)$ and $r(x, y)$ combine as a product to form $f(x, y)$:

$$f(x, y) = i(x, y)r(x, y) \quad (2.1-2)$$

where

$$0 < i(x, y) < \infty \quad (2.1-3)$$

and

$$0 < r(x, y) < 1. \quad (2.1-4)$$

· Homomorphic filtering:

The illumination-reflectance model can be used as the basis for a frequency domain procedure that is useful for improving the appearance of an image by simultaneous brightness range compression and contrast enhancement. It is particularly effective in cases of large intensity variations of the background where object differences are diffused within the background changes and objects are hard to be identified. In terms of its illumination (background) and reflectance components (objects), an image is expressed by means of the relation ? q. (2.1-2).

This equation cannot be used directly to operate separately on the spectra of illumination and reflectance because the Fourier transform of the product of two functions is not separable. In other words,

$$F\{f(x, y)\} \neq F\{i(x, y)\}F\{r(x, y)\}$$

Let, however, define

$$\begin{aligned} z(x, y) &= \ln f(x, y) \\ &= \ln i(x, y) + \ln r(x, y). \end{aligned} \quad (2.1-5)$$

Then,

$$\begin{aligned} F\{z(x, y)\} &= F\{\ln f(x, y)\} \\ &= F\{\ln I(x, y)\} + F\{\ln r(x, y)\} \end{aligned} \quad (2.1-6)$$

or

$$Z(u, v) = I(u, v) + R(u, v) \quad (2.1-7)$$

where $I(u, v)$ and $R(u, v)$ are the Fourier transforms of $\ln i(x, y)$ and $\ln r(x, y)$, respectively.

If we process $Z(u, v)$ by means of a filter function $H(u, v)$ then

$$\begin{aligned} S(u, v) &= H(u, v)Z(u, v) \\ &= H(u, v)I(u, v) + H(u, v)R(u, v) \end{aligned} \quad (2.1-8)$$

where $S(u, v)$ is the Fourier transform of the result. In the spatial domain,

$$\begin{aligned} s(x, y) &= F^{-1}\{S(u, v)\} \\ &= F^{-1}\{H(u, v)I(u, v)\} + F^{-1}\{H(u, v)R(u, v)\}. \end{aligned} \quad (2.1-9)$$

By letting

$$i'(x, y) = F^{-1}\{H(u, v)I(u, v)\} \quad (2.1-10)$$

and

$$r'(x, y) = F^{-1}\{H(u, v)R(u, v)\} \quad (2.1-11)$$

Eq.(2.1-9) can be expressed in the form

$$s(x, y) = i'(x, y) + r'(x, y) \quad (2.1-12)$$

Finally, as $z(x, y)$ was formed by taking the logarithm of the original image $f(x, y)$, the inverse operation yields the desired enhanced image $g(x, y)$; that is,

$$\begin{aligned} g(x, y) &= \exp[s(x, y)] \\ &= \exp[i'(x, y)] \exp[r'(x, y)] \\ &= i_0(x, y) r_0(x, y) \end{aligned} \quad (2.1-13)$$

where

$$i_0(x, y) = \exp[i'(x, y)] \quad (2.1-14)$$

and

$$r_0(x,y) = \exp[r'(x,y)] \quad (2.1-15)$$

are the illumination and reflectance components of the output image. The whole process is summarized in Fig. 2.1.

This enhancement approach is based on a special case of a class of systems known as *homomorphic systems*. The key to the approach is the separation of the illumination and reflectance components, achieved by Eq. (2.1-7), so that *homomorphic filter* $H(u, v)$ can operate on these components separately, as indicated in Eq. (2.1-8).

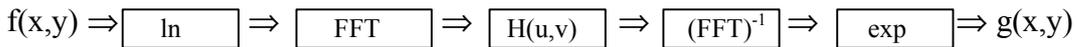
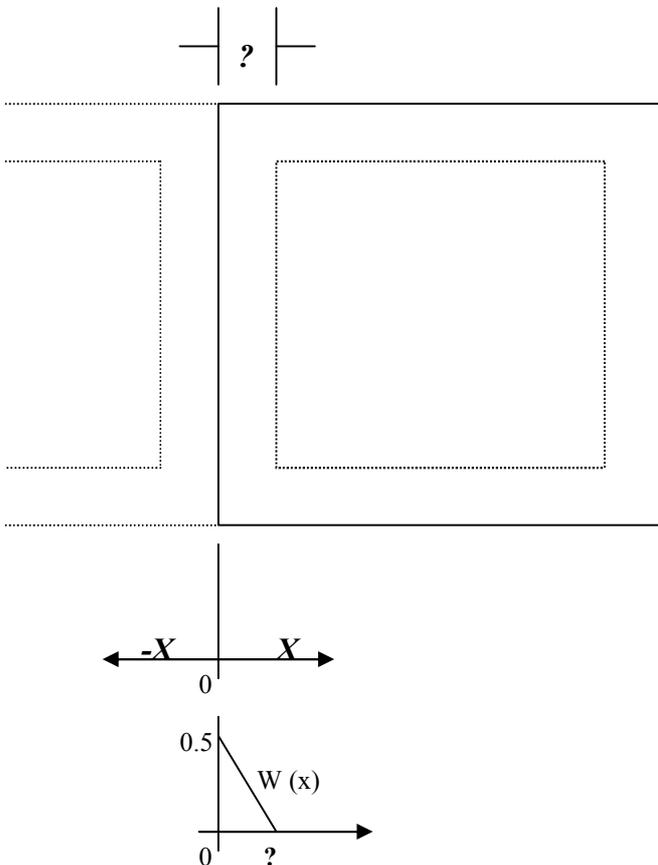


Fig. 2.1 Steps of homomorphic filtering

The illumination component of an image is generally characterized by slow spatial variations, while the reflectance component tends to vary abruptly, particularly at the junctions of dissimilar objects. These characteristics lead to associating the low frequencies of the Fourier transform of the logarithm of an image with illumination and the high frequencies with reflectance. Although these associations are rough approximations, they can be used to advantage in image enhancement.

The previous method produces an undesirable effect when applied to images with illumination varying strongly from side to side (dark to bright or reverse). As a consequence, these images have different illumination at each side. Homomorphic filtering enhances such differences. For smoother transitions at the edges of the image we need to impose some similarity at the horizontal and vertical borders, which results in a horizontal and vertical averaging at border areas (Fig 2.2). Border smoothing is achieved through windowing with $W(x)$.



filtering enhances such differences. For smoother transitions at the edges of the image we need to impose some similarity at the horizontal and vertical borders, which results in a horizontal and vertical averaging at border areas (Fig 2.2). Border smoothing is achieved through windowing with $W(x)$.

Two procedures of line scanning and filtering:

$$(1 - W(x)) * f(x) + W(x) * f(-x) \quad \left. \vphantom{(1 - W(x))} \right\} \text{ For } y=0 \dots N-1$$

$$(1 - W(y)) * f(y) + w(y) * f(-y) \quad \left. \vphantom{(1 - W(y))} \right\} \text{ For } x=0 \dots N-1$$

Fig. 2.2 Border smoothing

It becomes obvious that a good deal of control can be gained over the illumination and reflectance components with the homomorphic filter. This control requires the specification of a filter function $H(u, v)$ that affects differently the low- and high-frequency components of the image's spectrum. Fig. 2.3 shows a cross section of such a function. A complete specification of $H(u, v)$ is obtained by rotating the cross section 360° about the vertical axis. If the parameters γ_L and γ_H are chosen so that $\gamma_L < 1$ and $\gamma_H > 1$, the filter function shown in Fig. 2.3 tends to decrease the low frequencies and amplify the high frequencies, which is the desirable operation of this filter as described previously.

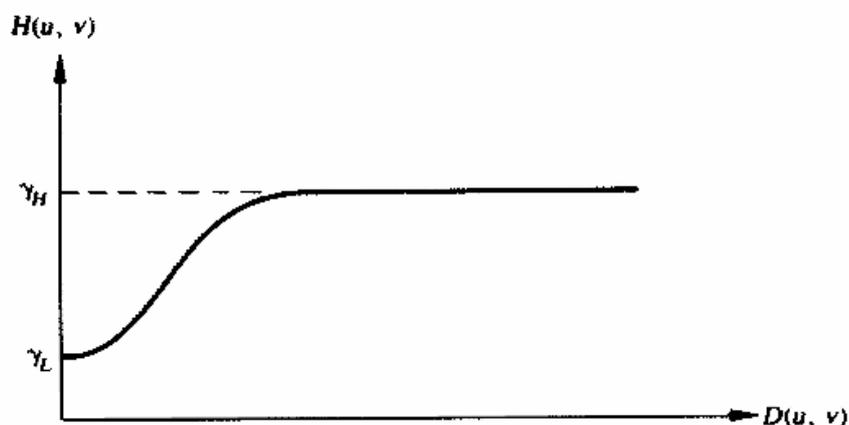


Fig. 2.3 Cross section of a circularly symmetric filter function for use in homomorphic filtering. $D(u, v)$ is the distance from the origin (from [1].)

The selection of the filter's cutoff frequency is crucial. The overall illumination is actually a slowly varying 2D function caused e.g. by surface curvature. We want to attenuate this component so that the high frequency detail will be easily detected. Nevertheless, illumination is spread out to a range of low frequencies, so that the appropriate cutoff point for separating the operation in the frequency domain is not obvious. As an example, consider the image in Fig.2.5-(a) that illustrates the problem. Selection of a very low cutoff frequency will not moderate the non-uniform illumination effect as seen in (Fig. 2.5 – (c)). On the contrary, if the cutoff frequency is high, then noise (high frequencies) is retained and the result is of questionable value, as seen in (Fig. 2.5-(b)). A closer look on spectra of images under consideration facilitates the appropriate selection. Fig. 2.4 shows the plot of one horizontal line of image's FFT. We observe a common abrupt fall in a certain region, so that we can select the appropriate cutoff, (Fig. 2.4-(b)), as the point of first significant FFT curvature change, or the first point where the magnitude of its derivative falls below a threshold. This cutoff point derives the image in Fig. 2.5–(d), which has compensated for illumination changes by the background.

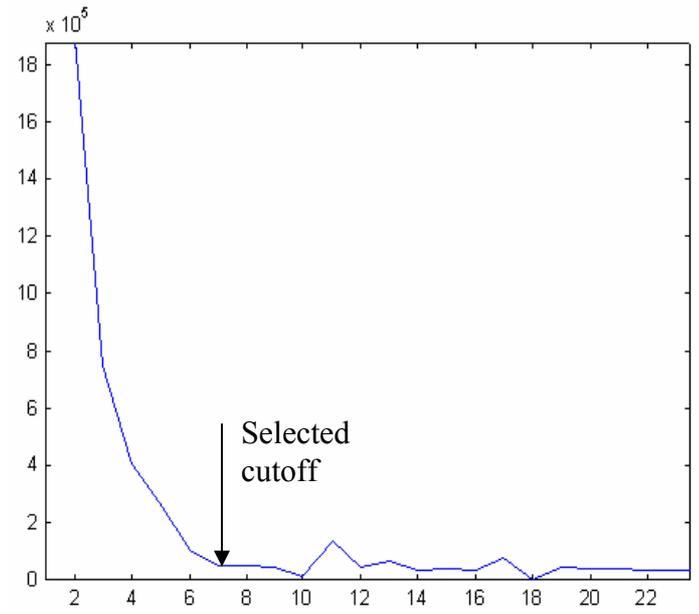
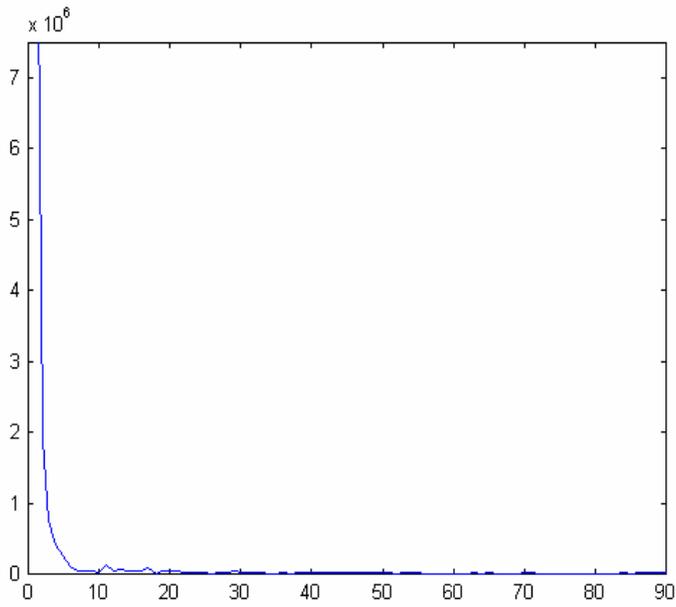
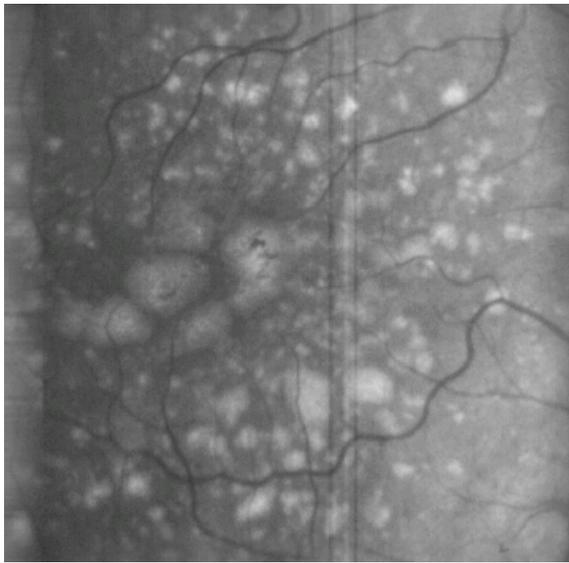
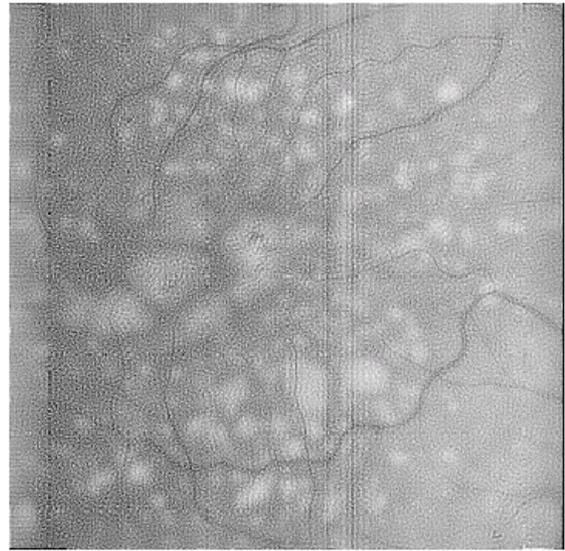


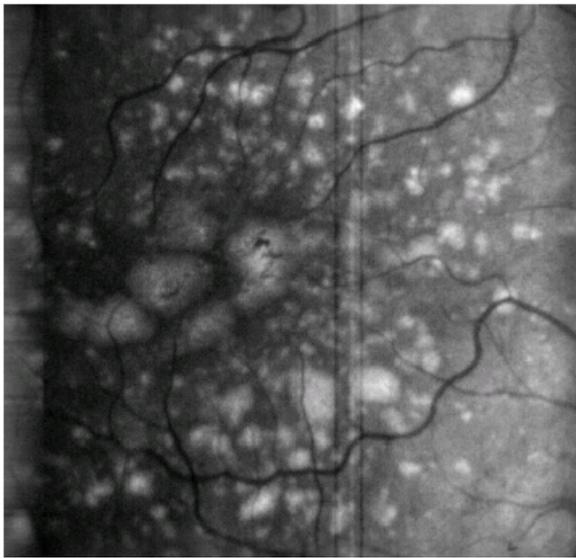
Fig. 2.4 (a) Fourier spectrum of image showing low and high cutoff frequencies; (b) zoomed plot showing the approximate position of selected cutoff



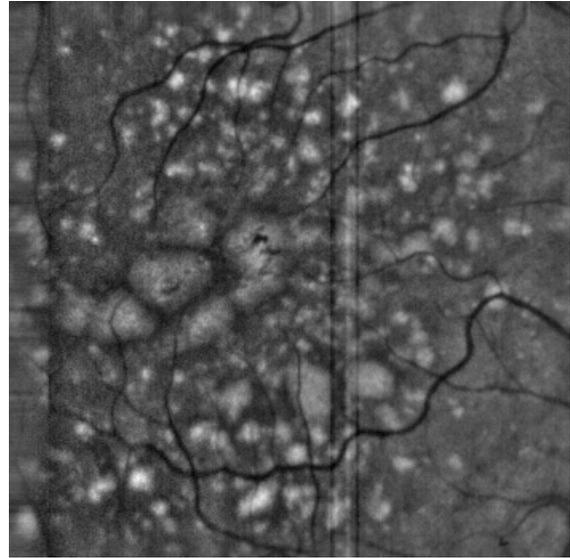
(a)



(b)



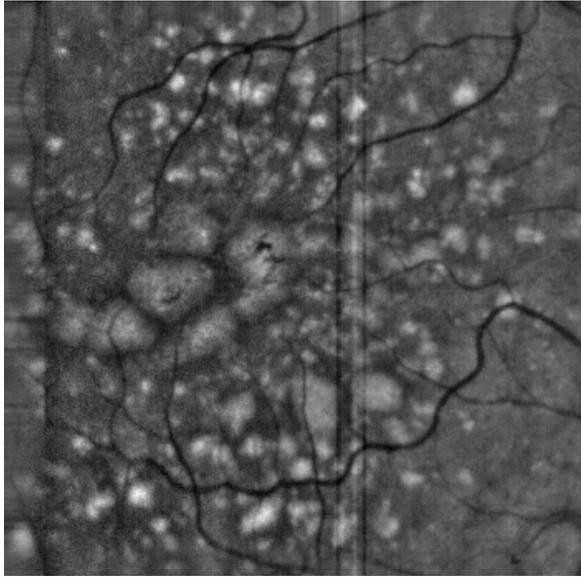
(c)



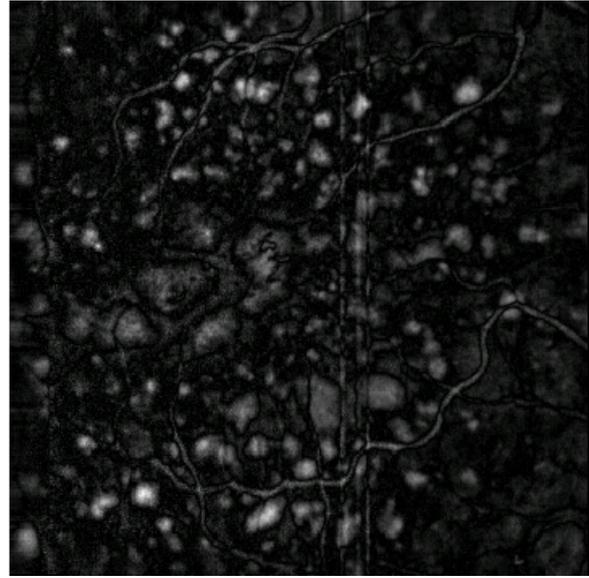
(d)

Fig. 2.5 (a) original image after modification of edges through averaging; (b) result of applying homomorphic with a high cutoff; (c) use of a low cutoff; (d) use of an intermediate cutoff (the selected one)

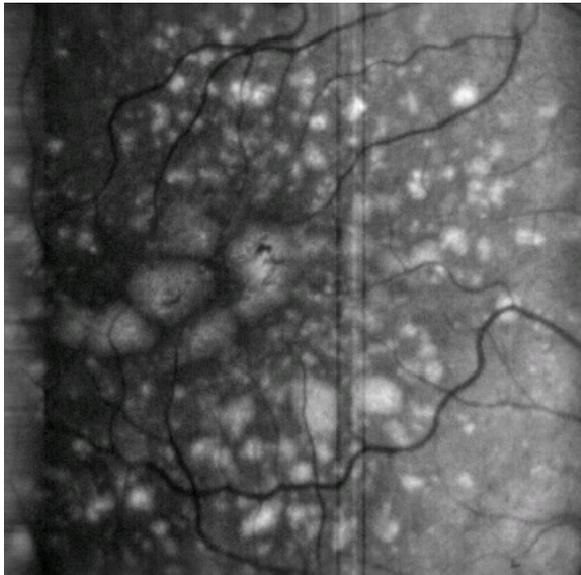
The selection of appropriate γ_L and γ_H is challenging, since a trade-off between preservation of low and amplification of high frequencies exist. In order to get this trade-off over, we provide results, shown in Fig. 2.6, obtained with various differences between values of γ_L and γ_H .



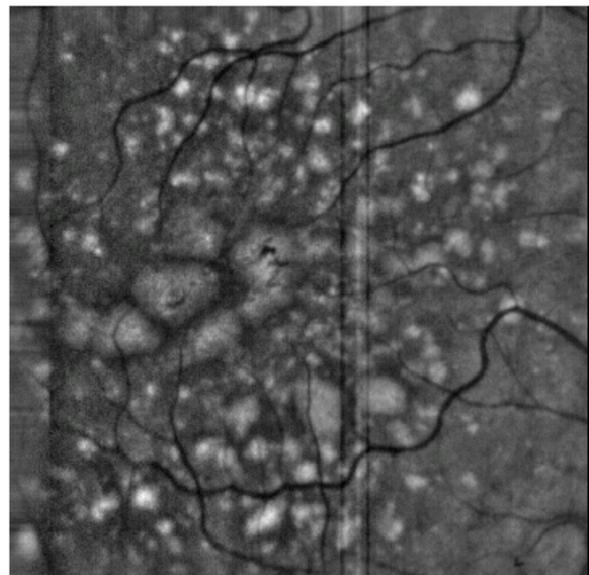
(a)



(b)



(c)



(d)

Fig. 2.6 (a) original image; (b) $\gamma_L=0.01$ $\gamma_H=1.5$; (c) $\gamma_L=0.9$ $\gamma_H=1.1$; (d) $\gamma_L=0.1$ $\gamma_H=1.2$;

The image in Fig. 2.6-(a) is badly illuminated resulting in darker right side and lighter right side. When the difference ($\gamma_H - \gamma_L$) is large, as in Fig. 2.6-(b), low frequencies are attenuated and uniform areas (large drusen) are not clearly distinguished from the background. This means that further enhancement would either shrink or break these areas apart. When the difference is small, as shown in Fig. 2.6-(c), the non-uniform illumination effect is not moderated, because low frequencies (responsible for overall illumination) are not adequately attenuated. Best result is obtained in Fig. 2.6-(d), as the difference in illumination is corrected and large drusen are clearly distinguishable.

2.2 Contrast Enhancement Techniques

2.2.1 Global vs. Local Enhancement

The contrast of an image is a property based on human perception abilities. An approximate definition of contrast is [37]

$$c = \frac{F - B}{F + B}$$

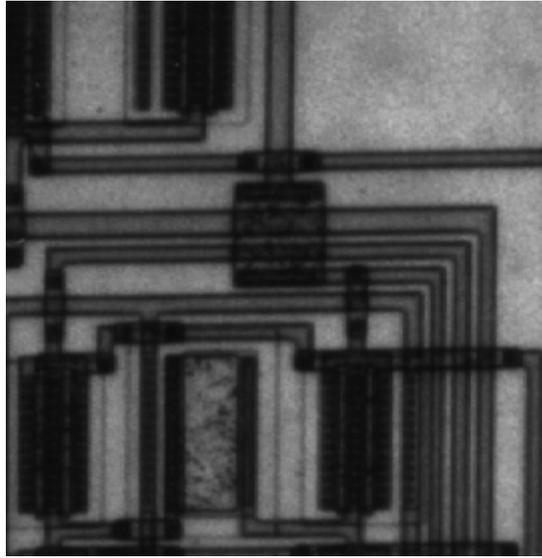
where F and B are the mean gray levels of two regions whose contrast is evaluated. Image processing defines several methods suitable for overall enhancement, like global histogram equalization - specification. Beyond global enhancement, it is often necessary to enhance details over small areas of an image.

Global enhancement techniques are affected by the overall distribution in the image and stretch illumination differences that are widely spread within the image. Actually, they separate strong concentrations in the distribution of the image. Small intensity differences in a local region are not accentuated.

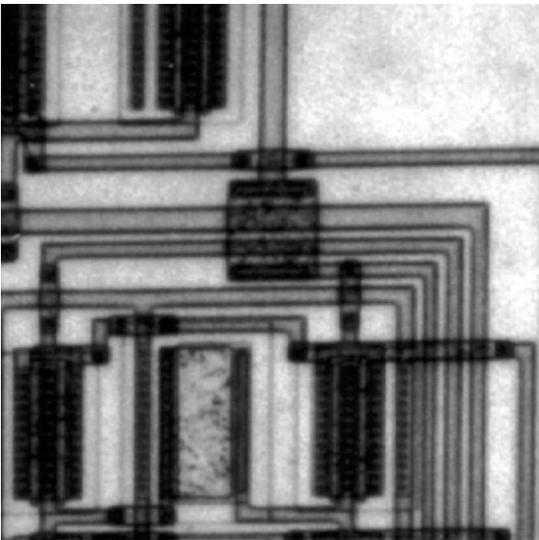
The general idea can be expressed in the framework of three cases:

- a) no enhancement for very small gray level differences between neighborhoods (caused probably by quantization noise or very small gray level variance);
- b) moderate to strong enhancement applied if the contrast between regions is small but outside the range of quantization contrast;
- c) no contrast enhancement is applied if the contrast is already sufficient.

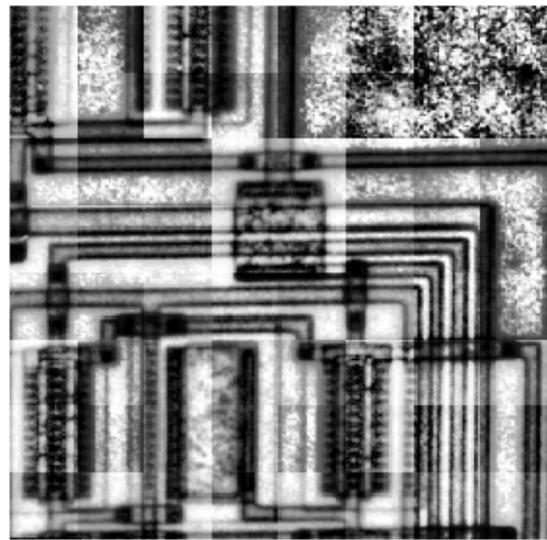
Global histogram processing techniques, like histogram equalization, are easily adaptable to local enhancement. The procedure is to define a sliding or a non-sliding window and apply the technique at smaller parts of the image. Features from neighboring pixels' intensities, like mean and standard deviation (variance), are often used at typical local transformations. Mean is a measure of average brightness and the variance is a measure of contrast (high contrast if pixels deviate significantly from their neighbors). Fig. 2.7 presents an example, where the same image is processed by both global and local technique. Fig. 2.7-(b) illustrates the overall enhancement of intensity variations obtained by global histogram equalization, while in Fig. 2.7-(c) we notice the excessive contrast enhancement in small regions.



(a)



(b)



(c)

Fig. 2.7 (a) original image; (b) result of global histogram equalization; (c) result of local histogram equalization using a small neighborhood about each pixel

2.2.2 Local Enhancement Techniques

The general form of a local transformation is expressed as follows:

$$g(x, y) = A(x, y) * [f(x, y) - m(x, y)] + m(x, y)$$

where $A(x, y) = k \frac{M}{s(x, y)}$ $0 < k < 1$

- $m(x, y)$: gray level mean in neighborhood centered at (x, y)
- $s(x, y)$: gray level standard deviation of local neighborhood
- M : global image mean

Application of the local gain factor $A(x, y)$ to the difference between $f(x, y)$ and the local mean amplifies local variations. Because $A(x, y)$ is inversely proportional to the standard deviation, areas with low contrast receive larger gain. The mean is added back to restore the average intensity level of the image in the local region. In practice, by adding back only a fraction of the local mean and restricting the variations of $A(x, y)$ between two limits (A_{\min} , A_{\max}) we can balance large deviations of intensity in isolated regions.

Several variations of the previous equation have been established in the literature.

· ***Adaptive contrast enhancement filter (ACE)***

The ACE filter is used to adjust the contrast differently in different regions of the image. Thus, regions with low standard deviation are enhanced, while regions with high standard deviation (e.g. those containing edges) retain their actual contrast. The adaptive contrast filter is based on the following equation [1]:

$$g(x, y) = k1 * \frac{M\{I(x, y)\}}{\sigma_l(x, y)} * (I(x, y) - m(x, y)) + k2 * m(x, y)$$

- $M\{I(x, y)\}$: mean of the entire image $I(x, y)$
- σ_l : local standard deviation
- m : local mean
- $k1, k2$: constants between 0 and 1

· ***Adaptive contrast enhancement filter II (ACE2)***

The ACE2 filter is a spatial domain method for contrast and dynamic range modifications with less limitation on linear contrast stretching [28].

$$g(x, y) = k2 * m(x, y) + k1 * [I(x, y) - m(x, y)]$$

- $I(x, y)$: pixel brightness value of the original image
- $m(x, y)$: arithmetic mean brightness value of an $(n \times n)$ window that is centered on the pixel position (r, c)
- $k1$: local gain factor
- $k2$: local mean factor

· ***Wallis statistical differencing***

The form of statistical differencing for enhancement is $g(x, y) = \frac{f(x, y)}{S(x, y)}$ where

$S(x, y)$ is the standard deviation estimated at pixel's neighborhood. Wallis suggested a generalization of this operator in which the enhanced image is forced to desired first order and second order moments [2]. The operator is defined by

$$g(x, y) = [f(x, y) - m(x, y)] \left[\frac{AS_d}{AS(x, y) - S_d} \right] + [rM_d + (1 - r)m(x, y)]$$

M_d : desired mean

S_d : desired standard deviation

A : gain factor that prevents overly large output values when $S(x, y)$ is small

r : mean proportionality factor controlling the ratio of the edge to background deviation

· *Exponential ACE*

The exp_ACE's algorithm is able to enhance the contrast and dynamic range of the image. It uses linear brightness stretching to modify the dynamic range of an image as:

$$g(x, y) = M * \left[\frac{I(x, y)}{M} \right]^{k1} + \left[\frac{m(x, y)}{I(x, y)} \right]^{k2}$$

M : the number of gray levels(255)

m : the local mean

$k1$: local gain factor

$k2$: local mean factor

The application of exponentials ($k1, k2$), controls the contrast gain to each region. Thus, if pixel's value is much higher than local mean almost no gain is applied, since $g(x, y)$ is inversely proportional to pixel's value. The first part of $g(x, y)$ adds back a value and compensates for large deviations of intensity.

· *Logarithmic ACE*

The log_ACE's algorithms are able to enhance the contrast and dynamic range of the image. Log_ACE's algorithm uses linear brightness stretching to modify the dynamic range of an image as:

$$g(r, c) = k1 * [\ln(I_{bar}(x, y) - \ln(m_bar(x, y)))] + k2 * m_bar(x, y) \quad [29]$$

$$m_bar = 1 - \frac{m(x, y)}{M} : \text{normalized complement of local mean}$$

$$l_bar = 1 - \frac{I(x, y)}{M} : \text{normalized complement of image}$$

M : number of gray levels (typically 255)

$k1$: local gain factor

$k2$: local mean factor

· **Histogram specification**

In histogram specification, the general idea is to increase the distance between gray levels of objects of interest and background, so that a single threshold can provide “good” segmentation results. Point operators (Fig. 2.8) can achieve histogram modification and succeed in separating different areas in a histogram.

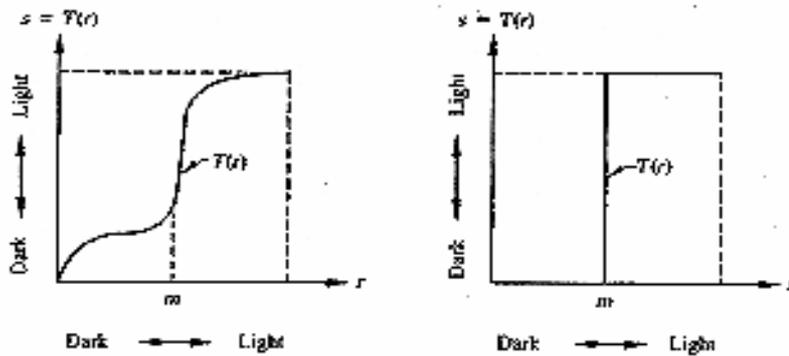


Fig. 2.8 Point operators

Nevertheless, such approaches apply the same transformation in all image areas, irrespective of the particular local distribution. The next method defines a local stretching transformation based on histogram specification.

· **Iterative local enhancement**

Point operators act only on the gray scale irrespective of the spatial location and the neighborhood features of that pixel. As a result, they often amplify noise. The method proposed in [9] could be seen as an adaptive histogram modification technique where the adaptation is performed according to the overall features of the local neighborhood. In order to achieve this, a linear transformation for each pixel is defined on a local area surrounding it. To refine the enhancement result, the size of the local area is iteratively reduced and act as a moving window. The goal is to obtain an enhanced image with an easy-to-threshold histogram (Fig. 2.9).

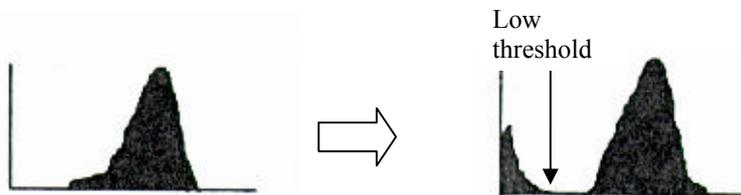


Fig. 2.9 (a) Histogram of original image; (b) histogram of enhanced image

The transformation can be expressed as follows:

$$A_k(x, y) = A_{k-1}(x, y) + m \{ave(x, y, W_k) - A_{k-1}(x, y)\}$$

$$ave(x, y, W_k) = \frac{1}{W_k^2} \sum_{i \in N_x(W_k)} \sum_{j \in N_y(W_k)} A_{k-1}(i, j)$$

$$N_x = \left\{ x - \left\lceil \frac{W_k}{2} \right\rceil, \dots, x + \left\lceil \frac{W_k}{2} \right\rceil \right\}$$

$$N_y = \left\{ y - \left\lceil \frac{W_k}{2} \right\rceil, \dots, y + \left\lceil \frac{W_k}{2} \right\rceil \right\}$$

$$0 < \mathbf{m} < 1$$

$$k = 1, 2, \dots, N$$

(N = total # of iterations)

· **Multilevel Histogram Equalization (MLE):**

Although, the previous methods may work well for a certain group of parameters, it's hard to select these parameters. In addition, because of its iterative scheme the previous technique is time consuming. On the contrary, exponential filtering is fast but is not capable of detecting single drusen and tends to combine them in bigger ones.

Histogram equalization reassigns the brightness values of pixels, so that their distribution spans the entire dynamic range of the image. In many cases, this spreads out the values in regions where different objects are observed, expressing detail in high differences of brightness. Essentially, histogram specification spreads out the peaks of the histogram and combines small intensity concentrations around the peaks into single values. Thus, it does not enhance detail of small spatial extent, even though it stretches overall intensity differences.

Equalization of the entire image could be useless when the image presents a trend of intensity variation along its spatial extent,. The non-uniform illumination effect, which is part of our problem, prohibits the use of global histogram equalization. Moreover, the existence of detail (objects) of varying extent and contrast necessitates the use of locally adaptive algorithms that adapt their performance without changing their design parameters. We propose here an enhancement scheme that takes advantage of both global and local approaches. In fact, it is a hierarchical (multilevel) scheme that progresses from the entire image to smaller regions. Due to the expected intensity similarity in small areas, the windows considered are non-overlapping. Compared with a sliding window approach, our scheme results in smaller computational complexity and larger speed of operation, without compromising on the local enhancement ability owing to its multilevel nature. In this approach, problems could arise using windows that are small enough to fit inside a drusen's region. This case can produce non-desirable misleading contrast variations as shown in Fig. 2.10-(c). It is desirable to use spatial windows larger than any drusen in the image. Considering this effect the algorithm proceeds as follows. The 1st stage of equalization uses a window equal to images's size (global). The 2^d stage splits the image into non-overlapping windows and applies the same operation to each part (block) of the previous result that is larger than any defected area. At any stage i a window w^i is further processed by smaller non-overlapping windows if and only if some application driven criteria are met.

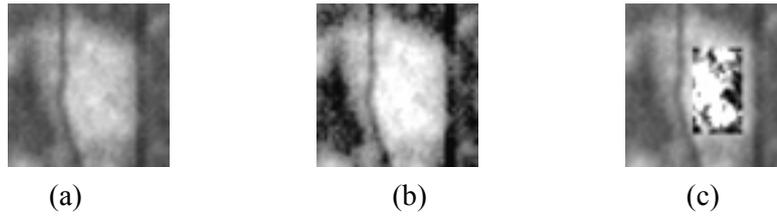


Fig. 2.10 (a) Original window containing one relative large drusen; (b) Histogram equalization using the entire window; (c) Histogram equalization using a smaller window inside the drusen's area

The first “pass” (Fig. 2.11-(a)) is responsible for enhancing the brightest parts of the image, including small, bright drusen and central parts of larger drusen. However, vague anomalies and dark areas that belong to spread drusen must be further enhanced, in order to be detected. The second stage of equalization, as shown in Fig. 2.11-(b), contributes in generating more distance between those “hidden” anomalies and surrounding areas. In our application we always proceed to the second stage. Nevertheless, due to the relatively large drusen experienced in all images tested, further enhancement is meaningless.

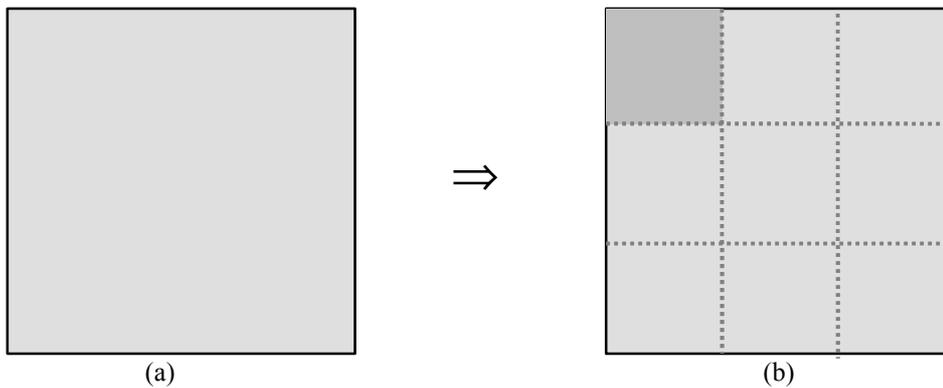


Fig. 2.11 (a), 1st level of histogram equalization (global) applied to entire image; (b) 2^d level of histogram equalization applied to regions of

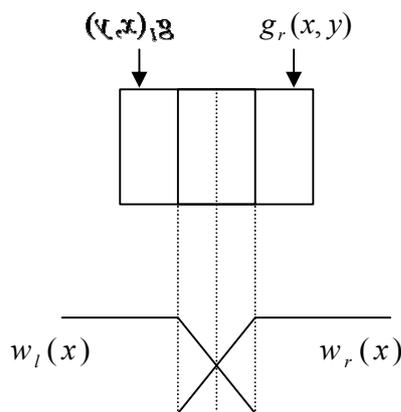


Fig. 2.12 Scheme for enhancing an image using non-overlapping windows with border interaction

A common problem of block processing is the undesirable generation of blocking effects. Usually, that means the generation of vertical or horizontal lines at block boundaries. To overcome such problems we utilize overlapping windows with a certain amount of overlap, as illustrated in Fig. 2.12.

The formula for obtaining the net result for each of the windows in the horizontal direction is written as

$$g(x, y) = w_l(x) * g_l(x, y) + w_r(x) * g_l(x, y)$$

2.2.3 Effectiveness of Local Enhancement Operators

Initial images:

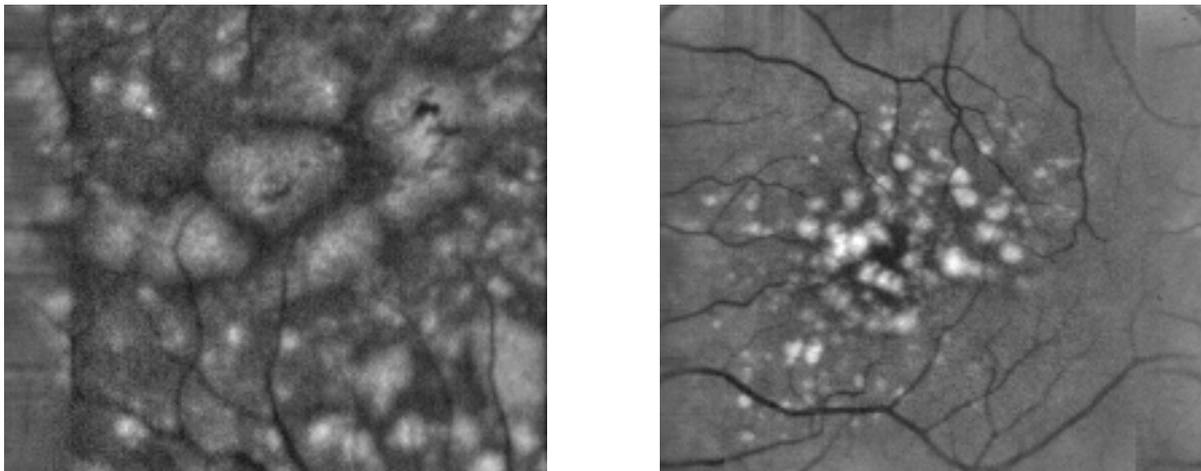
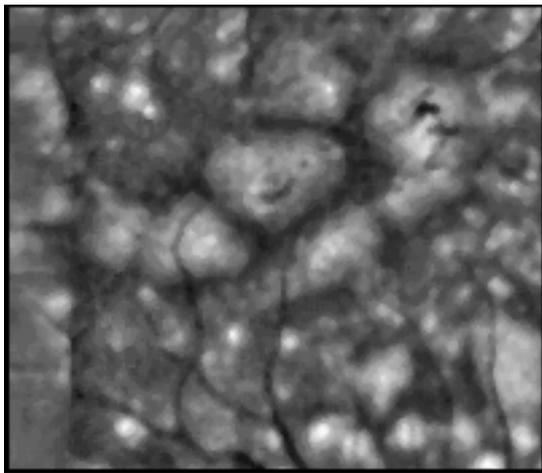


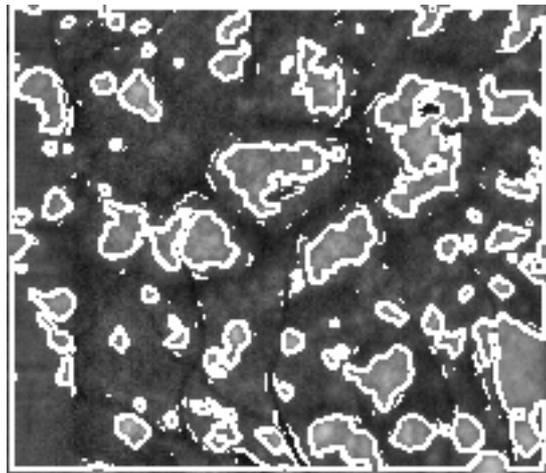
Fig. 2.13 (a) Initial image with large and small drusen; (b) initial image with small dense drusen

For comparison purposes, two representative images are tested. The first image contains relative large and small drusen that are close to each other, while the second one consists of small dense drusen at the central part and sparse ones elsewhere. Generally, extracting the entire area of large drusen is difficult because of the relative wide range of present gray-levels. Smaller drusen are either bright enough to make detection easy, or darker, with intensities similar to that of the background, which make their detection difficult. The results of the iterative method were thresholded using a large value, as specified from the method. Images derived from exponential filtering were manually thresholded to obtain almost optimal results and finally, multilevel equalization is applied and is thresholded by the technique presented in the next section. The results of segmentation are superimposed on the original images for comparison and are presented along with the enhanced images in figures 2.14-15-16.

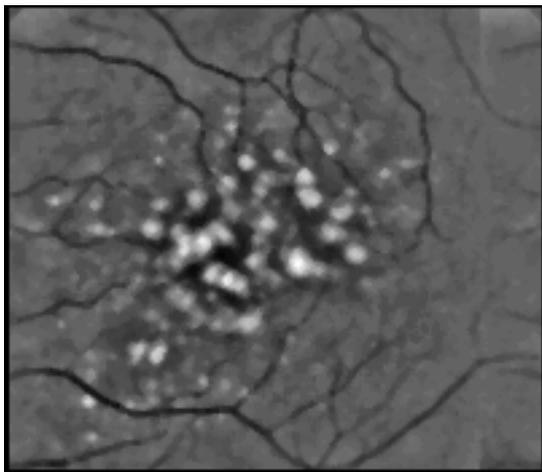
Iterative method:



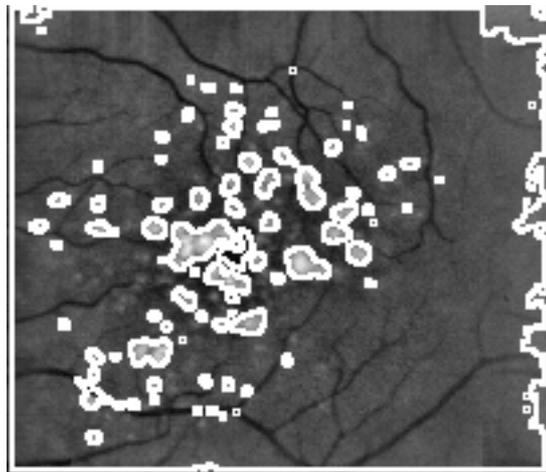
(a)



(b)



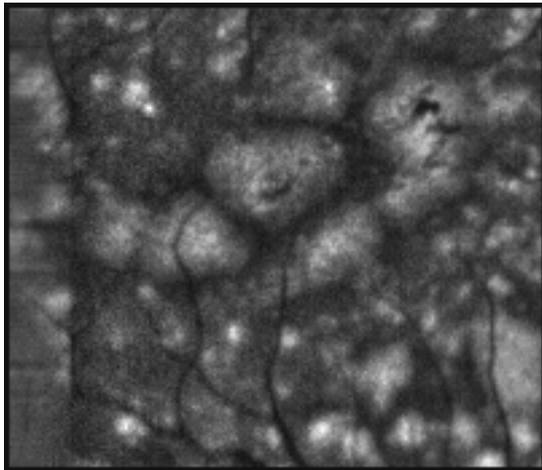
(c)



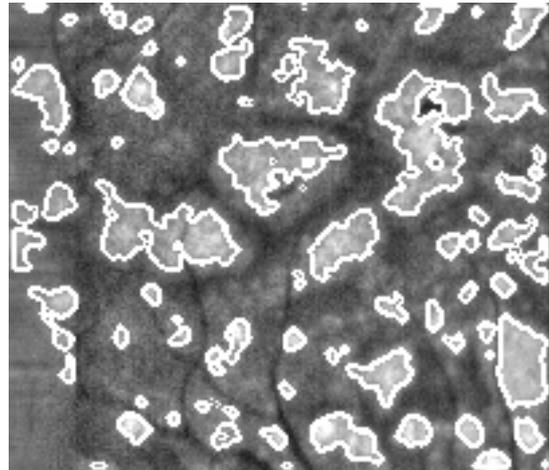
(d)

Fig. 2.14 (a),(c) enhanced images; (b),(d) detected regions over original images thresholded at 90th gray level

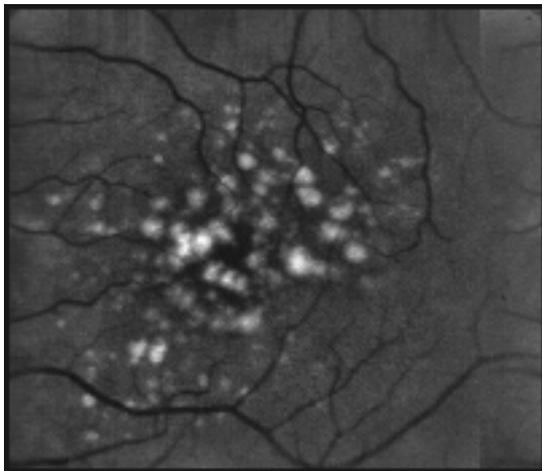
Exponential filter:



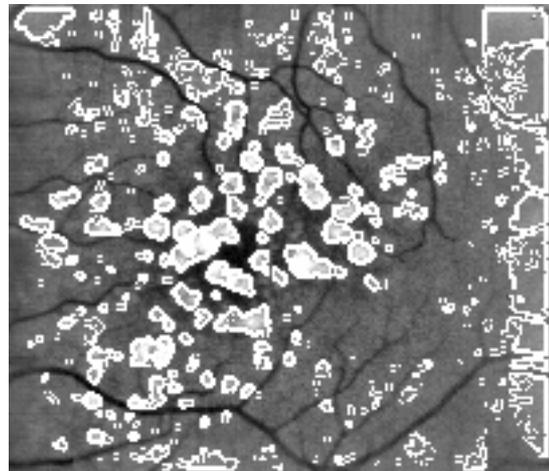
(a)



(b)



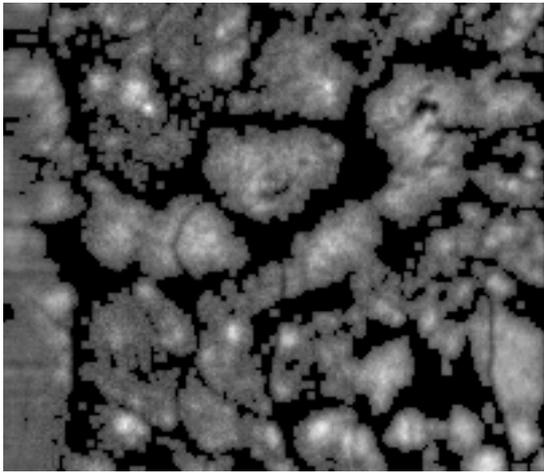
(c)



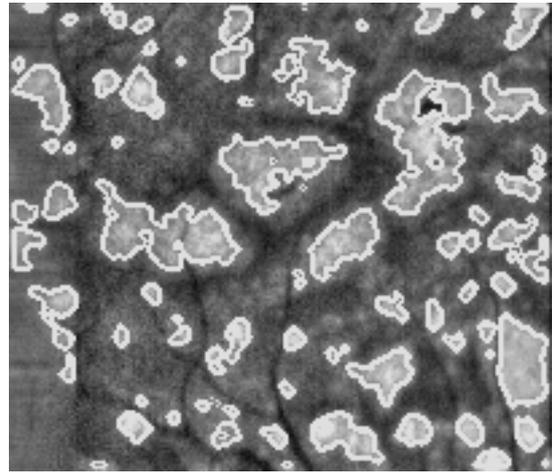
(d)

Fig. 2.15 (a),(c) enhanced images; (b),(d) detected regions over original images thresholded at 80% of corresponding histogram

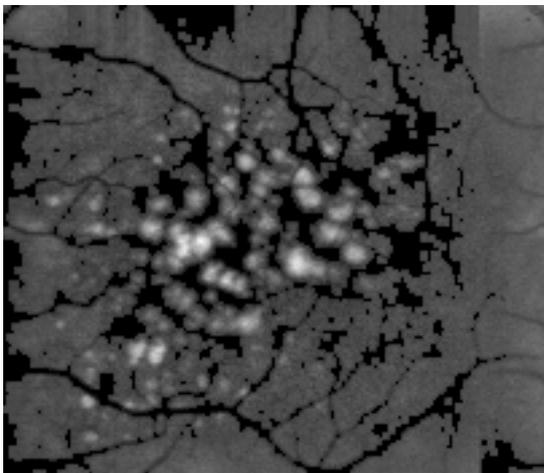
Multilevel equalization:



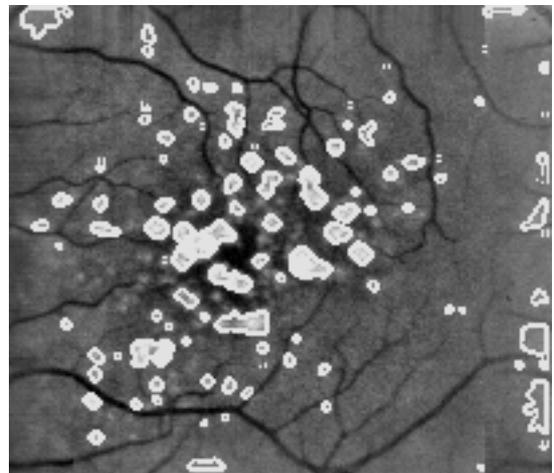
(a)



(b)



(c)



(d)

Fig. 2.16 (a),(c) enhanced images; (b),(d) detected regions over original images thresholded using HALT

Iterative local enhancement and exponential contrast enhancement are neighborhood operators that access pixels in an area around the central pixel, perform a calculation on those values and then derive a new value for the central pixel. Their main drawback is the required parameters, which are, in general, hard to specify and affect the algorithm's performance drastically. Both of them proved to be inefficient in our case study, mainly because there is not a group of constant parameters that will produce an adequate enhancement result for the set of images available. The exponential method performs a one-pass filtering, so it's relative fast. Iterative enhancement is based on a procedure that performs multiple passes according to the specified parameters. Consequently, it's a very slow and memory-consuming method, since the previous step's image has to be saved before continuing with enhancement. Our proposed multilevel equalization approach is much faster and does not require temporary memory for storing intermediate results.

The iterative method succeeds in detecting all the drusen that are contrasted from the background "by first glance", as shown in Fig. 2.14. The algorithm manages to detect large drusen thoroughly, but fails in separating some small, dark drusen. The latter case is obvious at the second test image, where vague drusen are either not detected or joined with false edges belonging to vessels. Another issue is the failure to separate nearby drusen, resulting in contours that contain more than one drusen. Fault detection is also present in regions where the background is lighter than usual or, as mentioned, where drusen are near to vessels.

Exponential filtering detects most evident drusen, as illustrated in Fig. 2.15, but is sensitive to small brightness variations. This results in false detection of areas that surround real anomalies and belong mostly to the background. The latter remark is evident in processing the second image, where many false negatives are detected.

The proposed approach of multilevel histogram equalization is capable of detecting and separating properly most of the anomalies, as shown in Fig. 2.16. It is much less sensitive to small brightness variations (located at background regions) and produces isolated contours containing areas with similar bright gray values (single drusen). Fault detection is not experienced in the images tested and hard to distinguish drusen are also located. An undesirable situation could arise out of the presence of noise, generated during the acquisition phase. Equalization in areas containing noise will result in further enhancement of noise and present objects. Although noise is not common among this kind of images, (we experienced this problem only once in spite of our extended test set of images) the proposed algorithm involves certain steps to compensate for this problem, as explained in Chapter 4.

2.3 Thresholding Techniques

Gray level thresholding is the simplest segmentation process. Its effectiveness derives from the fact that many objects are characterized by nearly constant reflectivity or light absorption of their surfaces. The enhancement process emphasizes and amplifies object similarities compared with their background. However, local enhancement operators also amplify intensity variations within the same objects. For this reason

global thresholding may not be adequate to derive object detail due to object differences at various intensity levels. In image processing applications it is often crucial to make an appropriate threshold selection, so as to obtain a fine segmentation. To achieve this, global thresholding may be performed on smaller blocks of the image rather than on the entire image.

Haralick and Shapiro [4] have established the following quantitative guideline for a “good” image segmentation scheme: “Regions of an image segmentation should be uniform and homogenous with respect to some characteristic such as gray tone or texture. Region interiors should be simple and without many small holes. Adjacent regions of a segmentation should have significantly different values with respect to the characteristic on which they are uniform. Boundaries of each segment should be simple, not ragged, and must be spatially accurate.”

A threshold can be determined interactively or through an unsupervised non-parametric method. The use of histogram to select a threshold is a very common technique. Frequently, simple observations are enough for detecting a good threshold point. The lowest point between two histogram peaks is often a good threshold value. The problem of detecting this point automatically consists of two steps: locating the two peaks and finding the lowest point in between. A parametric approach to identify these peaks is by approximating them with Gaussian curves. Thus, Gaussians are used to fit the histogram and the largest two are used to identify the major peaks. However, this method is slow and not promising because of the unknown number of normal distributions and the distances among them. Several segmentation techniques focus on two distributions assuming that object and background pixels have different mean levels and are random numbers drawn from one of two normal distributions. These distributions have their own standard deviations and variances. One of the most important and older methods that uses this assumption is the Otsu [5] one.

It’s now obvious how difficult it is to find a single “good” threshold for the gray level image in our application, because of the non uniform illumination, the different background color or the non uniform surface curvature. As an alternative option, we can define thresholds for smaller image regions, either based on local properties or local histogram information.

The histogram of a local area is more focused and informative regarding the separation of included objects. Moreover, in our approach local thresholding is imposed by local (adaptive) enhancement. Most of thresholding techniques are easily adjusted to a localized scheme without algorithmic changes. The following sections consider well established thresholding techniques and develop one new approach suitable for our application.

2.3.1 Otsu Algorithm

Otsu [5] proposed an automatic optimal threshold selection method from gray level histograms using a discriminant criterion. The method is unsupervised and nonparametric, properties that are desirable in most image processing applications.

In an ideal case an image's histogram has a deep valley between two peaks representing objects and background. However, that is not the case for most realistic scenes, since noise or unequal heights of the two peaks produce a non-ideal gray level distribution. Some techniques attempt to overcome this problem by approximating the histogram in the least square sense by a sum of Gaussian distributions, and then apply statistical decision procedures [6]. The implementation of these techniques requires many and, sometimes, unstable calculations. However, the conclusion from the application of these threshold selection methods is that Gaussian distributions are good for approximating the modes of a histogram.

Towards this direction, Otsu's approach, is based on discriminant analysis. The threshold operation is regarded as the partitioning of the pixels of an image into two Gaussian classes C_0 and C_1 (e.g., objects and background) discriminated at gray level t . That is $C_0 = \{0, 1 \dots t\}$ and $C_1 = \{t+1, t+2 \dots l-1\}$. Let $\mathbf{s}_w^2, \mathbf{s}_B^2, \mathbf{s}_T^2$ be the within-class variance, between-class variance, and the total variance, respectively. An optimal threshold can be determined by minimizing one of the following (equivalent) criterion functions with respect to t :

$$l = \frac{\mathbf{s}_B^2}{\mathbf{s}_w^2}, \quad h = \frac{\mathbf{s}_B^2}{\mathbf{s}_T^2}, \quad k = \frac{\mathbf{s}_T^2}{\mathbf{s}_w^2}.$$

Of the above three criterion functions, h is the simplest. Thus, the optimal threshold t^* is defined as:

$$t^* = \underset{t \in G}{\text{Arg Min}} h,$$

where

$$\begin{aligned} \mathbf{s}_T^2 &= \sum_{i=0}^{l-1} (i - \mathbf{m}_T)^2 p_i, & \mathbf{m}_T &= \sum_{i=0}^{l-1} i p_i, \\ \mathbf{s}_B^2 &= w_0 w_1 (\mathbf{m}_1 \mathbf{m}_0)^2, & w_0 &= \sum_{i=0}^t p_i, & w_1 &= 1 - w_0, \\ \mathbf{m}_1 &= \frac{\mathbf{m}_T - \mathbf{m}_0}{1 - w_0}, & \mathbf{m}_0 &= \frac{\mathbf{m}_0}{w_0}, & \mathbf{m} &= \sum_{i=0}^t i p_i. \end{aligned}$$

Extension to multilevel thresholding using Otsu method is straightforward.

2.3.2 Thresholding According to Distribution's Curvature

The use of a single global threshold is almost always inadequate for realistic images. A single threshold is unsuccessful when e.g. more than one objects with different gray levels and a non-uniform background exists. A solution to this problem could be the use of a multiple threshold technique. Different objects are recognized from different

lobes in a multimodal histogram. The histogram of an image with gray levels in $[0, L-1]$ is a discrete function $p(r_k) = \frac{n_k}{n}$, where r_k is the k th gray level, n_k is the number of pixels in the image with that gray level, n is the total number of pixels in the image, and $k = 0, 1, 2, \dots, L-1$. Under the ergodicity assumption $p(r_k)$ provides an estimate of the probability of occurrence of gray level r_k . A plot of this function for all values of k can provide a crude description of image's appearance. Boukharouba *et al.* [7,8] proposed a method that uses the intrinsic properties of the cumulative distribution function of an image to derive threshold values. In this method, the curvature of the distribution function is examined prior to the threshold values. The distribution function $F(k)$ at point k is given by

$$F(k) = \frac{\sum_{g=0}^k h(g)}{\sum_{g=0}^{l-1} h(g)},$$

where $h(i)$ is the histogram distribution and $f(i)$ the cumulative distribution.

The curvature of F is then defined by

$$C(x) = F''(x) [1 + (F'(x))^2]^{-3/2},$$

where F' and F'' are the first and second derivatives of F , respectively. The zeros of the curvature determine the thresholds, as well as the gray level to be assigned to each class.

Other properties of the histogram are often used for threshold selection. Some of the most important properties are presented in Appendix A.

These thresholding techniques can be applied to either global or local segmentation performed on the entire image or on smaller windows, respectively. In our problem we need to discriminate small regions of varying intensities within the image. Global thresholding fails to discriminate small local differences. Therefore, the method we use is a local thresholding technique. It is based on local histogram analysis as presented in the following section.

2.3.3 Relaxation Method

Relaxation was introduced by Southwell [30,31] to improve the convergence of recursive solution for systems of linear equations. In image segmentation the pixels of an image are first probabilistically classified into "light" and "dark" classes, based on their gray levels. Then by looking at every pixel's neighborhood these probabilities are adjusted, so as to become very high for light and dark regions.

Rosenfeld and Smith,[32], suggested following formula for initial classification of pixels:

If d and l are the darkest and lightest gray levels and g_i is the gray level of a pixel x_i , then if $g_i > m$, let

$$p_{i,light}^0 = \frac{1}{2} + \frac{1}{2} \frac{g_i - m}{l - m}$$

and if $g_i \leq m$, let

$$p_{i,dark}^0 = \frac{1}{2} + \frac{1}{2} \frac{m - g_i}{m - d}$$

Another approach, which assumes that a histogram can be divided into two Gaussian distributions is suggested by Fekete *et al.*[33].

Finally, previous probabilities must be updated according to a pixel's neighborhood, in order to obtain a "good" segmentation. So, if \mathcal{C} is the set of class labels (e.g., the classes of dark and light pixels), then a compatibility coefficient, $r_{ij}(\mathcal{C}, \mathcal{C}')$, between a pixel x_i with label $\mathcal{C} \in \mathcal{C}$ and another pixel x_j with label $\mathcal{C}' \in \mathcal{C}$ is defined such that

$$r_{ij}(\mathcal{C}, \mathcal{C}') = \begin{cases} -1 & \text{if } \mathcal{C} \text{ and } \mathcal{C}' \text{ are } \textit{incompatible} \\ 0 & \text{if } x_i \text{ and } x_j \text{ are } \textit{independent} \\ 1 & \text{if } \mathcal{C} \text{ and } \mathcal{C}' \text{ are } \textit{compatible} \end{cases}$$

Zucker *et al.*[34] propose the following equation for updating the probabilities:

$$p_i^{k+1}(\mathcal{C}) = \frac{p_i^k(\mathcal{C})[1 + q_i^k(\mathcal{C})]}{\sum_{\mathcal{C}' \in \Lambda} p_i^k(\mathcal{C}') [1 + q_i^k(\mathcal{C}')]},$$

$$q_i^k(\mathcal{C}) = \frac{1}{8} \sum_{x_j \in N_i} \sum_{\mathcal{C}' \in \Lambda} r_{ij}(\mathcal{C}, \mathcal{C}') p_j^k(\mathcal{C}'),$$

where N_i is the 8-neighbor of x_i .

However, Pavlidis [35] proves that the previous formula is not mathematically correct. The above scheme violates the natural expectation that the labeling should not change if neighboring pixels are independent. Thus, Peleg [36] suggests another formula:

$$p_i^{k+1}(\mathcal{C}) = \frac{p_i^k(\mathcal{C}) \sum_{x_j \in N_i} \sum_{\mathcal{C}' \in \Lambda} r_{ij}(\mathcal{C}, \mathcal{C}') p_j^k(\mathcal{C}')}{\sum_{\mathcal{C}' \in \Lambda} p_i^k(\mathcal{C}') \sum_{x_j \in N_j} \sum_{\mathcal{C}'' \in \Lambda} r_{ij}(\mathcal{C}', \mathcal{C}'') p_j^k(\mathcal{C}'')}$$

2.3.4 Histogram Local Adaptive Thresholding (HALT):

If our purpose was to distinguish bright and darker areas in an image, a localized Otsu thresholding scheme would work well. However, our goal is to separate those areas, as well as others hidden in the background distribution, without being affected by intensity variations, caused by non uniform illumination. Otsu's method will separate with no respect to these peculiarities of the problem.

Forced by this weakness, we propose a thresholding technique that is based on local histogram properties. After previous algorithm steps part of the background disappears. Although this facilitates the decision of a single threshold, its unique specification is still difficult because of the dense and complicated image's histogram. However, if we zoom into each local area, we observe the different shapes of histograms for each of these regions. In order to determine a "good" single threshold for each neighborhood, we can take the following facts into consideration:

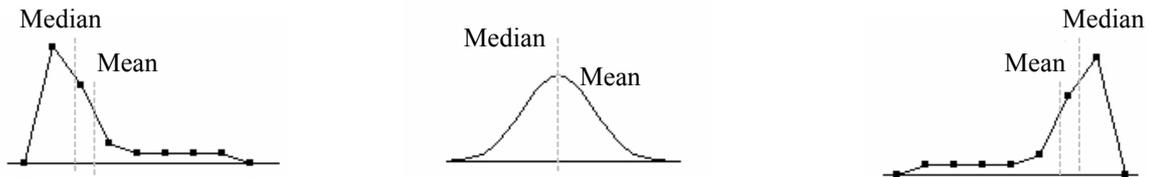
The histogram's mean is a good measure of central tendency for roughly symmetric distributions but can be misleading in skewed ones, since it can be greatly influenced by extreme values (e.g. large number of black pixels-bias). For normal distributions the mean is efficient and less subject to sample fluctuations than other measures of central tendency. Nevertheless, other statistics, such as the median, are sometimes more useful, since they are more efficient than the mean in highly skewed distributions. The peak (or mode) can be informative but it can not be used as the only measure of central tendency, since it is highly susceptible to extreme gray value differences. Another problem occurs when the histogram has more than one peaks (multimodal distribution) and it is difficult to determine the "right" one or the one closer to the distribution's center.

As a consequence, the mean, median and peak are almost equal for roughly symmetric distributions. The mean is higher than the median in positively skewed distributions and lower than the median in negatively skewed distributions.

In background or general homogeneous regions, the gray-scale distribution approaches a normal one and the histogram appears like a Gaussian function. On the contrary, when small or bigger bright spots (drusen) are present, the histogram is positively skewed denoting a non-symmetric distribution. In either case an appropriate threshold decision has to be taken.

In order to distinguish between the latter cases, two symmetry quotients are proposed. The first one utilizes the differences $|mean - median|$ and $|mean - peak|$ as a gross indication. The second quotient examines skewness in conjunction with kurtosis as a refined symmetry quotient. The $|mean - median|$ difference is a first measure of symmetry based on statistical measures. The $|mean - peak|$ difference is chosen as a measure of histogram's main lobe spread, since it is often the largest of the two differences in our application. Representative examples are shown in Fig. 2.18.

$$|mean - median|$$



$$|mean - peak|$$

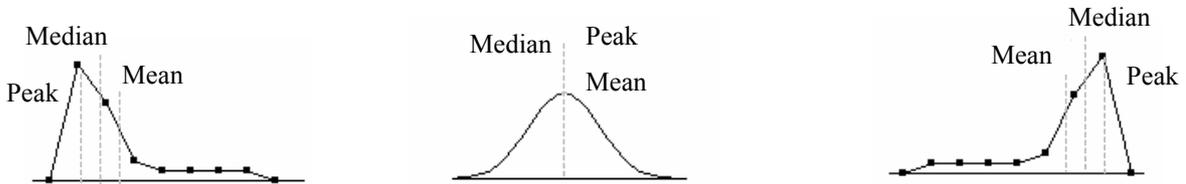


Fig. 2.18 Three typical cases of histogram shape (negatively skewed, symmetric, positively skewed) and relative positions of corresponding features.

Histogram has to be checked for general symmetry or asymmetry. The two symmetry quotients described above are used for this purpose. Thus, if both of them are small the distribution under consideration appears to be almost symmetric in contrast to the opposite case.

Generally, if a symmetric distribution is the case, then the area under consideration belongs probably to the background and setting 90% of histogram as threshold will remove it. Application of this threshold will also leave only isolated points in the sub region, which can be removed by a small median filter. On the contrary, a thresholding scheme like Otsu would be effective in separating different sub-distributions, when the histogram is asymmetric. Otsu's thresholding technique is more suitable for the non-symmetric distribution, since it can effectively detect the transition level from background to drusen. More details about the method will be given during the description of main algorithm.

2.3.5 Results of Thresholding Schemes

In order to demonstrate the efficiency of the HALT method, two representative images are chosen. One with large drusen dominating extensive areas and one with few, small and vague drusen (Fig. 2.19). Both images are enhanced using multilevel histogram equalization (MLE) and thresholded using HALT. A median filter is applied afterwards to remove isolated pixels.

Otsu's localized thresholding scheme works fine in regions that are dominated by drusen (brighter areas), since the distinction between them and the background is evident. This is demonstrated in the first image (Fig. 2.20-(a)), where drusen at the central part of it are correctly distinguished from the surrounding areas. Unfortunately, the algorithm is strongly confused by regions that do not contain any abnormality, like those located at the sides. Because of non uniform illumination, parts of these regions are brighter and are misclassified as anomalies. Second image, Fig. 2.20-(b) brings

out another disadvantage of this scheme. Vague drusen, which are either small or located inside bright background regions, are not segmented. The algorithm seems to detect the most obvious drusen (two of them are easily conceived), but fails to detect “hidden” anomalies; some of those are indicated by the arrows (Fig. 2.19 –(b)). On the contrary, the HALT technique removes most of the background in both cases, as shown in Fig. 2.21. Even the most hard-to-see drusen are segmented without losing their actual size and shape. Some false negatives generated by the existence of noise can be easily removed at a following detection step.

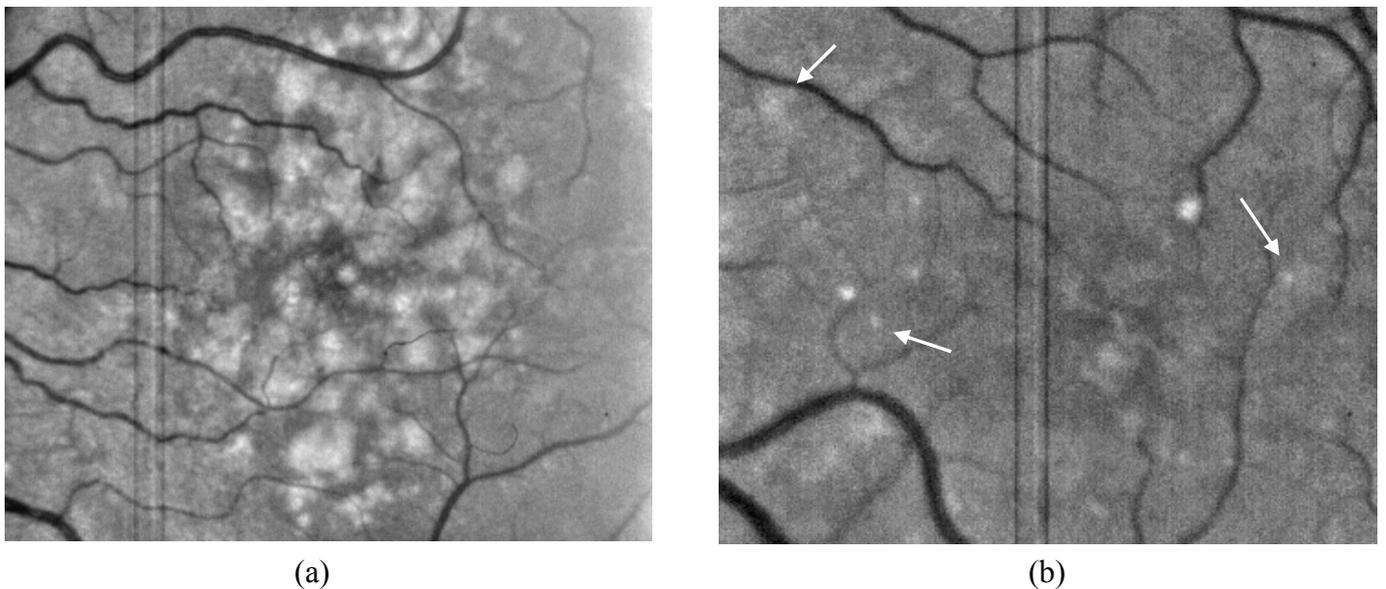


Fig. 2.19 (a) Image with large dense drusen; (b) image with small sparse drusen

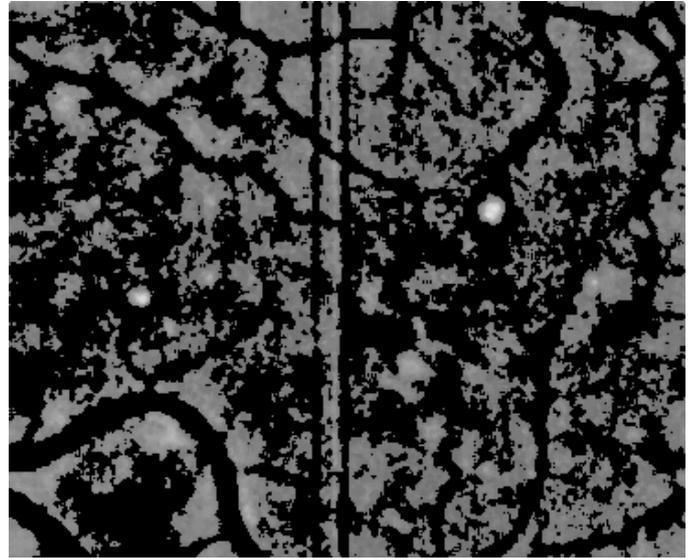
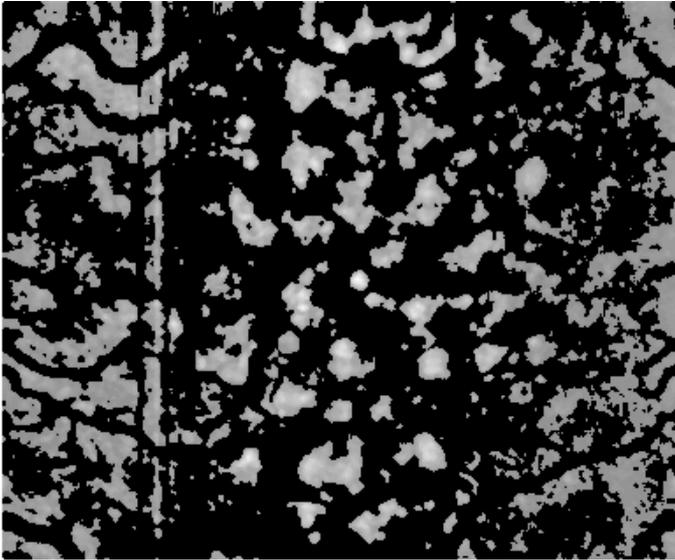


Fig. 2.20 Images after MHE enhancement, local thresholding using Otsu's method at each block and median filtering to eliminate sparse pixels

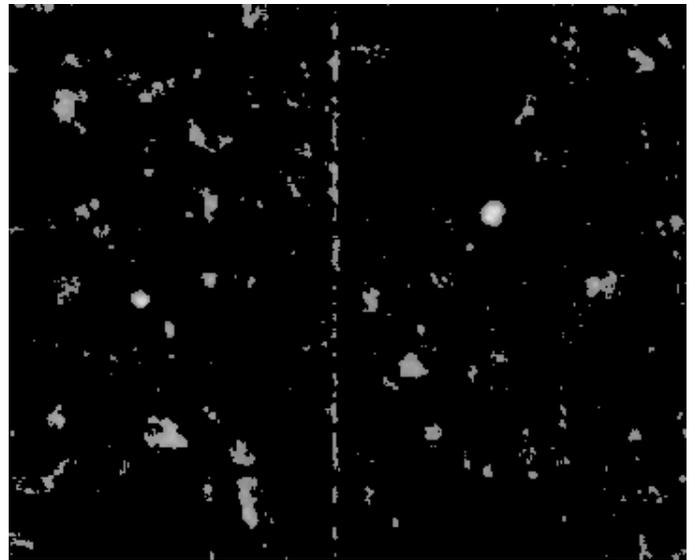
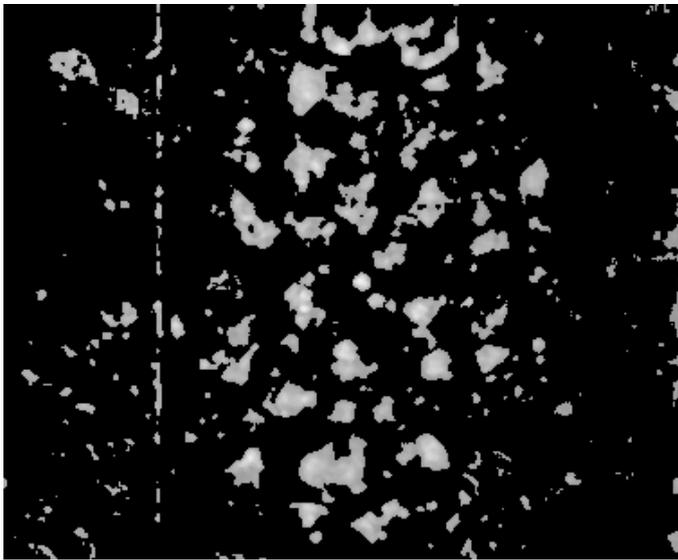


Fig. 2.21 Images after MHE enhancement, local thresholding using HALT method and median filtering to eliminate sparse pixels

3. MAIN ALGORITHM FOR DRUSEN DETECTION

The main algorithm we used is summarized in Fig. 3.1. Each block stands for a different algorithm operation.

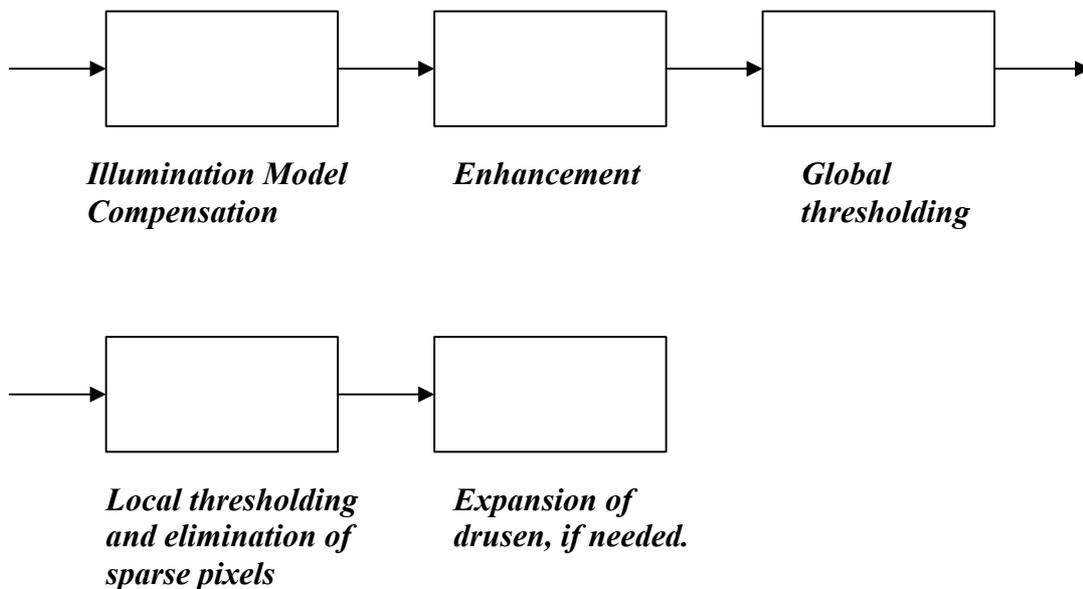


Fig. 3.1 Algorithm for anomalies' detection in human eye's retina

3.1 Illumination Model Compensation

The first problem we encountered is the non-uniform illumination caused by the varying curvature of the retina's surface. Thus, the first block in our algorithm is directed towards illumination model compensation, which is achieved by border averaging and homomorphic filtering. This part of the algorithm is actually the only one that requires parameter specification, so as to obtain the best possible illumination correction. These parameters are case dependent, implying that for our class of images constant parameters are used, while different values would fit better to e.g. mammogram application. Nevertheless, the non-uniform illumination effect cannot be fully cancelled. The following steps of the algorithm have to take in consideration that not all bright parts of the image (drusen are brighter than anything else) are defected areas.

3.2 Enhancement

The second step in our algorithm is the enhancement operation that is responsible for generating more distance between gray levels of abnormalities and background. We

emphasize here that the background also includes the vessels, which are of any interest to our application. The proposed MLE approach succeeds in enhancing most drusen, being insensitive to small brightness variations, caused e.g. from non-uniform illumination. Small and bright drusen are extracted correctly, something that holds also true for large and spread drusen that tend to be darker near the edges.

3.3 First Level Thresholding

As a result of the previous operators, sharp abnormalities in intensity (candidate drusen) are strongly enhanced and separated from background. They, actually, look brighter than everything else in the image does. A global threshold is capable of removing darker parts that belong to the drusen's surrounding areas (background). For this purpose we employ Otsu's [5] thresholding technique that is unsupervised, non-parametric and generally considered as an efficient method [7]. In addition to these desirable characteristics, its implementation is straightforward and fast. The thresholding operation is used for eliminating unnecessary regions of the image. Thus, values above the threshold are set to zero. Another scheme to achieve the same result first derives a binary image with 1's and 0's at values above and below the threshold, respectively, and then multiplies the input image with this binary template. In our analysis we adopt this scheme to preserve the structures of interest in binary form. Nevertheless, this first level of thresholding cannot discriminate vague abnormalities hidden in the local regions of background.

This step, in general, will detect and preserve regions of evident abnormalities that are crisply separated from their background and background regions mixed with vague abnormalities. In order to refine the segmentation, the algorithm proceeds with local thresholding operations.

3.4 Local Thresholding and Elimination of Sparse Pixels

In order to extract the actual drusen and get rid of misleading areas a novel local thresholding operator is designed and tested. The HALT technique provides quite a powerful unsupervised tool for separating objects of interest from background areas. At first, the image is divided in small areas and a threshold for each area is selected, according to its histogram shape. This process is continued to a second stage of local thresholding inside each block area, if required.

It is emphasized here that local thresholding considers only the histogram of pixels that pass the global thresholding steps. In essence it is only applied to pixels that are not set to zero by the first level thresholding in 3.3.

As shown in Fig. 3.4 the HALT operator is preceded by a morphological dilation and succeeded by a median filter. The morphological dilation expands the regions that are not removed by global thresholding. If this expansion occurs in background areas, there is not much effect, since the following segmentation is capable of removing these expanded regions completely. The main advantage of dilation comes into sight in those areas that contain only one or two large drusen and no background. In this case any threshold application would result in loss of drusen' area. The expansion

results in the recovering of some anomalies' background. In other words, it forces better distinction between bright areas (anomalies) and darker (surroundings) at the corresponding histogram. Morphological filtering compensates this critical problem. Mathematical morphology consists of two basic operators, namely *dilation* and *erosion*. If A and B are sets in Z^2 , with components $a=(a1, a2)$ and $b=(b1, b2)$ then morphological dilation, [1], is defined as

$$A \oplus B = \left\{ x \mid (\hat{B})_x \cap A \neq \emptyset \right\}.$$

Set B is referred as the *structuring element* and previous operator results usually to an expansion of set A . After the application of HALT, a median filter is applied. The gray level of each pixel is replaced by the median of the gray levels in a predefined neighborhood of that pixel. This operator eliminates sparse pixels, causing false "alarms" for presence of anomalies, and preserves only those pixels that appear compactly distributed into regions.

The background is composed of a noise process superimposed on a deterministic smoothly varying ramp image with uniform histogram. A symmetric Gaussian distribution characterizes this overall background process. Using the ergodicity assumption, any realization of the stochastic process or any acquired image from this process is also characterized by this Gaussian distribution. It is expected therefore that by thresholding the distribution at its 90% and preserving only values above this 90% threshold, we leave only isolated pixels randomly distributed along the spatial extent of the image that can be easily removed by median filtering.

Negatively skewed distributions are likely to describe areas of background. So, setting 90% as threshold would remove them. That holds for both cases of symmetric and non-symmetric distributions.

A more analytic consideration of symmetry and histogram analysis in the HALT algorithm is given in Fig. 3.2-3. Recall that crude symmetry indicators are used as first level discriminants and the categorization is refined through histogram-shape indices.

A. Histogram is totally or almost symmetric (Table)

- A totally symmetric gray level distribution signifies areas that are mainly occupied by background regions; uniform large regions, which surround anomalies. However, small drusen may be present, so setting 90% as threshold would be adequate to remove background and leave whole or part of the anomalies.
- The class of platykurtic distributions may be misleading. Generally, symmetric distributions signify background areas. Nevertheless, the platykurtic feature signifies interaction of distributions that jointly preserve symmetry. For example, if background areas' and anomalies' gray levels are normally and equally distributed, histogram will still appear symmetric. So, in order to avoid removal of drusen, we use Otsu method for thresholding.

- In case of sharp-peaked (leptokurtic) almost symmetric histograms we observe high concentration of pixels around the mean value. These sections appear with almost uniform background. Leptokurtic distributions allow the existence of only small drusen as outliers that do not alter the general uniformity of the intensities. Using Otsu thresholding, that is obtaining a threshold value close to mean, would retain anomalies and big part of the background. So, setting 90% as threshold would remove background areas and retain, if existing, small drusen.
- The case of mesokurtic and positively skewed histograms requires particular attention. The mesokurtic characteristic most likely arises from the background distribution. The positive skewness indicates interaction with another distribution, which is observable but not significant one to alter drastically the background statistics. This distribution is detected at high intensity values indicating the existence of object(s), whose intensity however interacts with that of the background. Thus, their direct segmentation may be inefficient. Using Otsu's threshold may leave large areas of the background, whereas using the 90% threshold may delete a good portion of the object's structure. So, an additional step of local thresholding is used, which is actually the application of HALT method focused on smaller areas of first level's region. This helps in obtaining better distinction of anomalies and background at corresponding histograms.

B. Histogram is totally or almost asymmetric (Table)

- A positively skewed distribution of this class notifies the presence of many small or large drusen. In fact, bright gray levels that generally characterize anomalies dominate the histogram. Otsu technique is best suited to this case, since the distinction of bright and darker areas (background) is obvious.
- In general, an asymmetric distribution signifies the presence of drusen. The asymmetric platykurtic and mesokurtic distributions can result as combinations of similar distributions, characterizing background and abnormalities (drusen). The leptokurtic distribution describes an area dominated by background (with highly concentrated values) and less drusen.

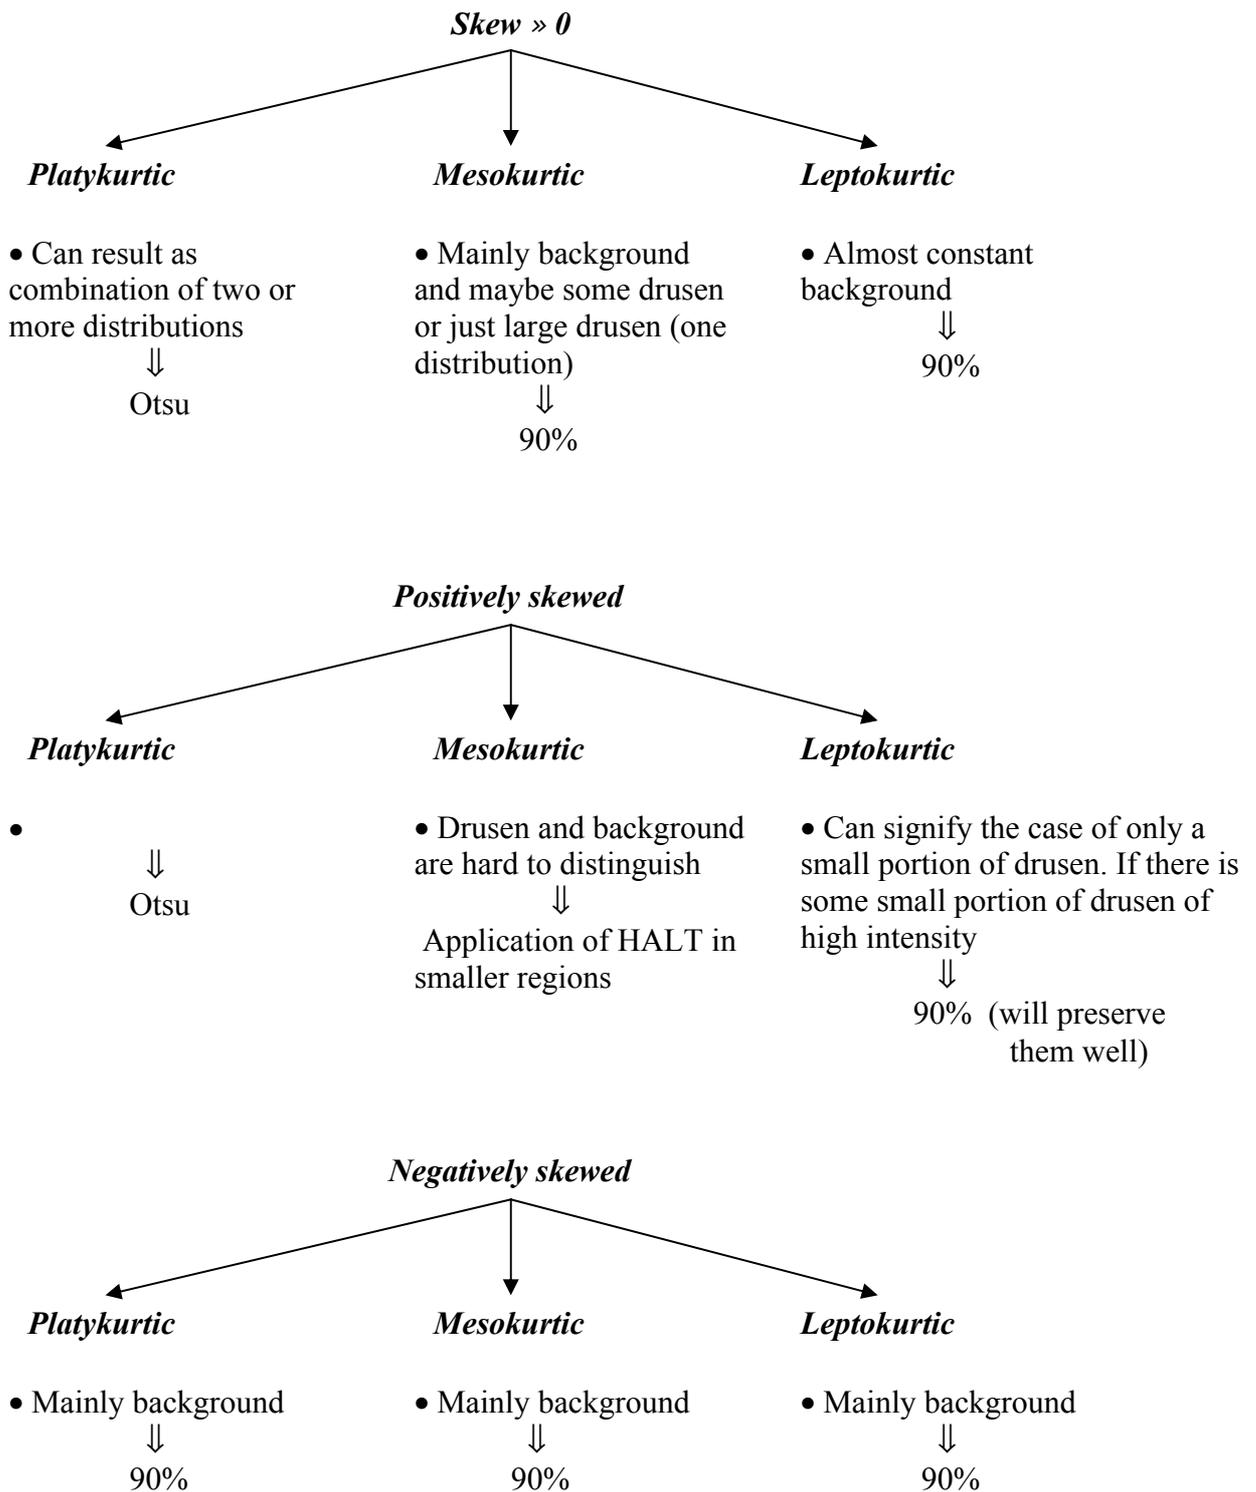


Fig. 3.2 Decision making of HALT in case of a symmetric distribution

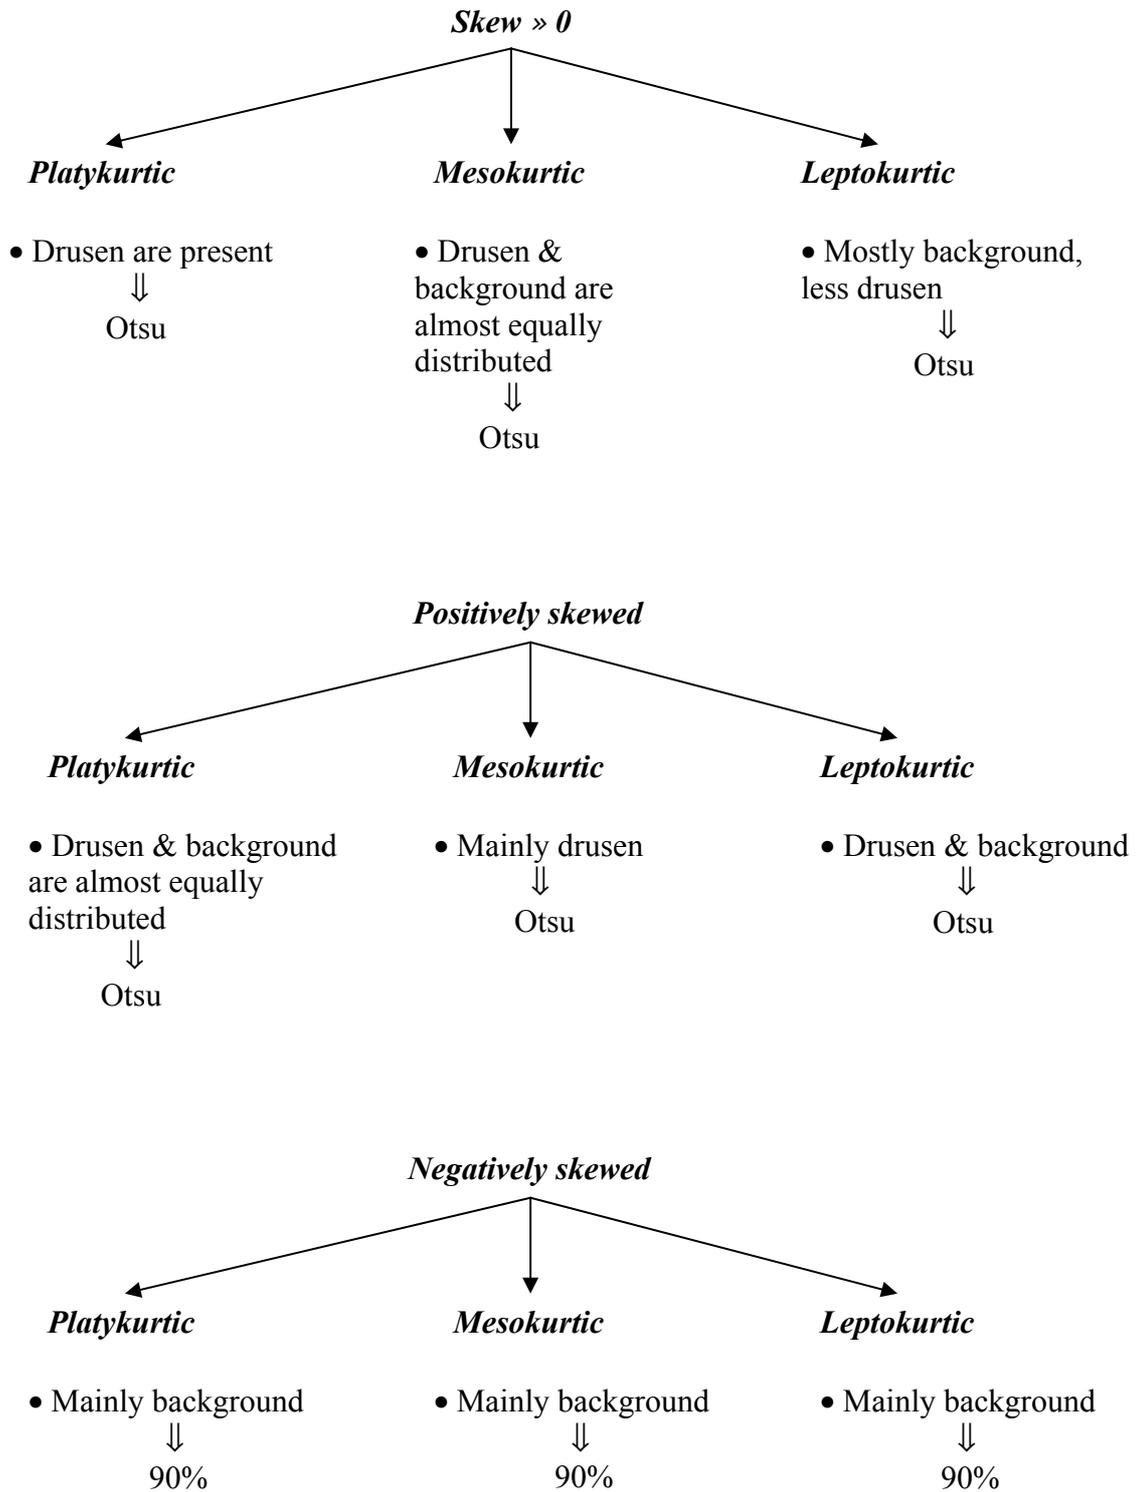


Fig. 3.3 Decision making of HALT in case of an asymmetric distribution

3.5 Expansion

If the median filter preserves large drusen, especially in symmetric mesokurtic, skewed regions, then it is likely that the region is dominated by large drusen which appear homogeneous, just as the background. This is the case of a histogram that looks totally symmetric (small skew-mesokurtic). In such cases, we need to check if these remaining compact regions need to be extended further to capture the entire blob region.

This expansion is achieved by mathematical morphology. The second basic operator of morphology is the *erosion*, [1], which is defined as

$$A \ominus B = \left\{ x \mid (B)_x \subseteq A \right\}.$$

A morphological closing operator is actually a dilation followed by erosion and is defined as

$$A \bullet B = (A \oplus B) \ominus B.$$

A morphological closing with a small structuring element (disk shaped) will dilate all remaining regions and as a consequence it will join neighboring groups of pixels that possibly belong to a larger defective area. The following erosion will restore the size of those areas, but will retain previous joined regions. These regions generated by previously separated groups of pixels must be further expanded, in order to obtain more or less their actual size. A small dilation is needed to achieve this goal.

3.6 Overall Algorithm

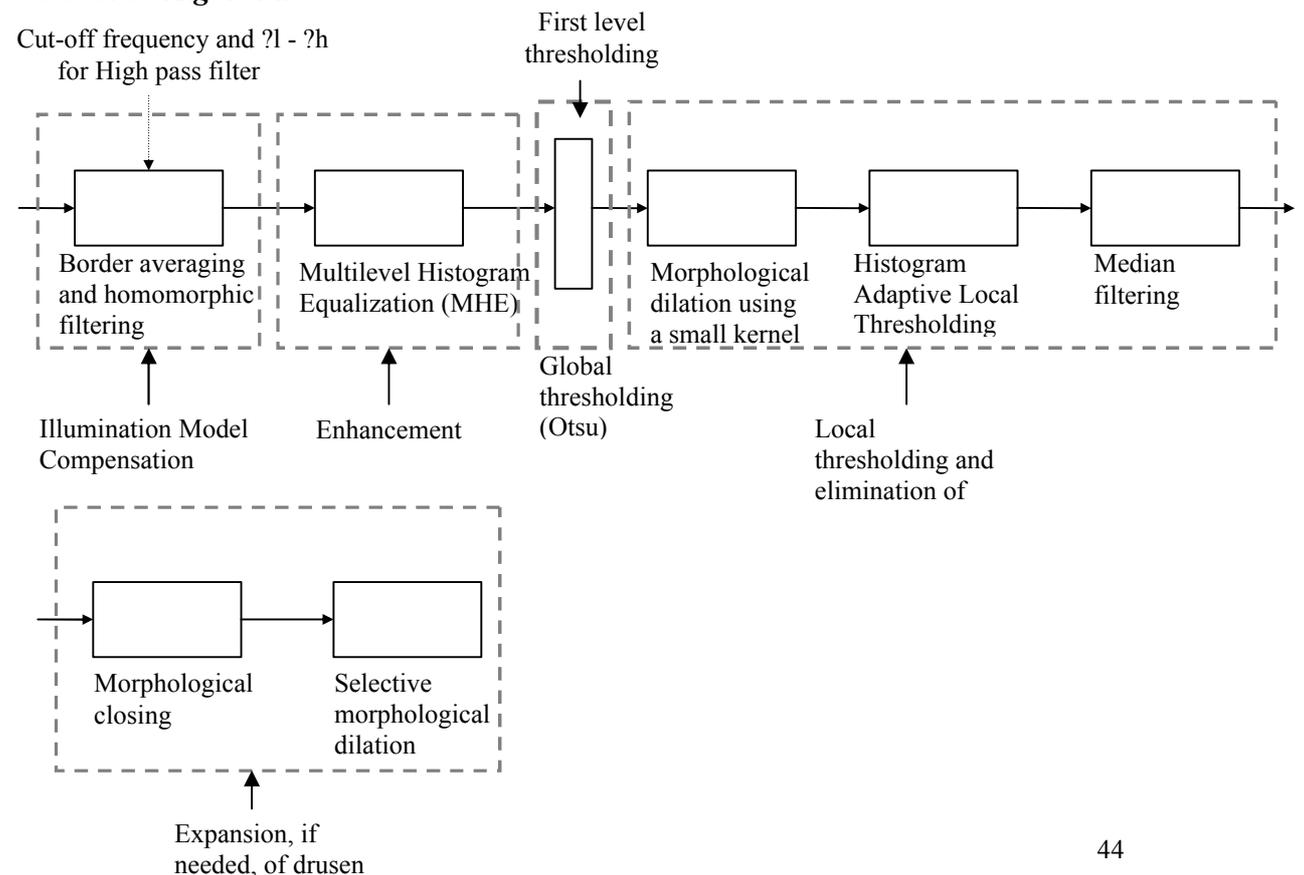


Fig. 3.4 Algorithm for anomalies detection in human eye's retina

4. EXAMPLES

4.1 Illumination Model Compensation

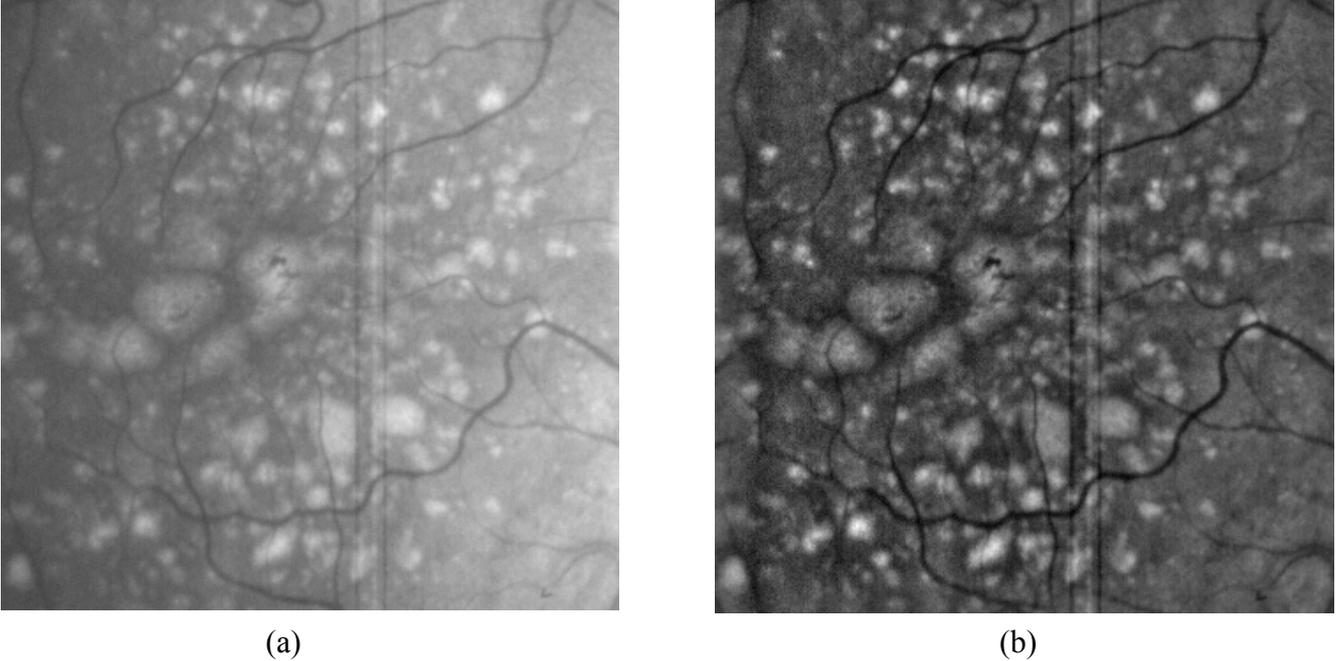


Fig. 4.1 (a) Original image with non-uniform illumination, (b), image after non-uniform illumination correction using homomorphic filtering

The original image, in Fig. 4.1-(a), shows an eye's retina in which an overall variation in brightness is obvious. Left part of the image looks darker than the right one and some anomalies, especially in the bright areas, look similar to the background. After averaging the borders and applying homomorphic filtering, as shown in Fig. 4.1-(b), drusen are better defined and separated from the background. Diseased areas become brighter and surroundings become darker, as can be seen from the large drusen at the central part. In addition to the non-uniform illumination correction signifies an increase in contrast is obvious. The distinction among neighboring drusen is improved and vessels become dark enough, so as not to misjudge parts of them as actual anomalies.

As mentioned in 2.3.4 and 3.4, the crude symmetry indicators are used by HALT as first level discriminants and the categorization of the histogram is refined through histogram-shape indices. The numerical limits of symmetry quotients are derived after careful observations of images' local histograms and are consolidated as

- a) If $|mean - median| \leq 4$ and $|mean - peak| \leq 5 \Rightarrow$ symmetry
- b) If $-0.5 < skew < 0.5 \Rightarrow$ histogram is not skewed
If $skew \geq 0.5 \Rightarrow$ histogram is positively skewed
If $skew \leq -0.5 \Rightarrow$ histogram is negatively skewed
- c) If $-2 < kurtosis < 2 \Rightarrow$ distribution is mesokurtic
If $kurtosis \geq 2 \Rightarrow$ distribution is leptokurtic
If $kurtosis \leq -2 \Rightarrow$ distribution is platykurtic

4.2 Enhancement and Global Thresholding

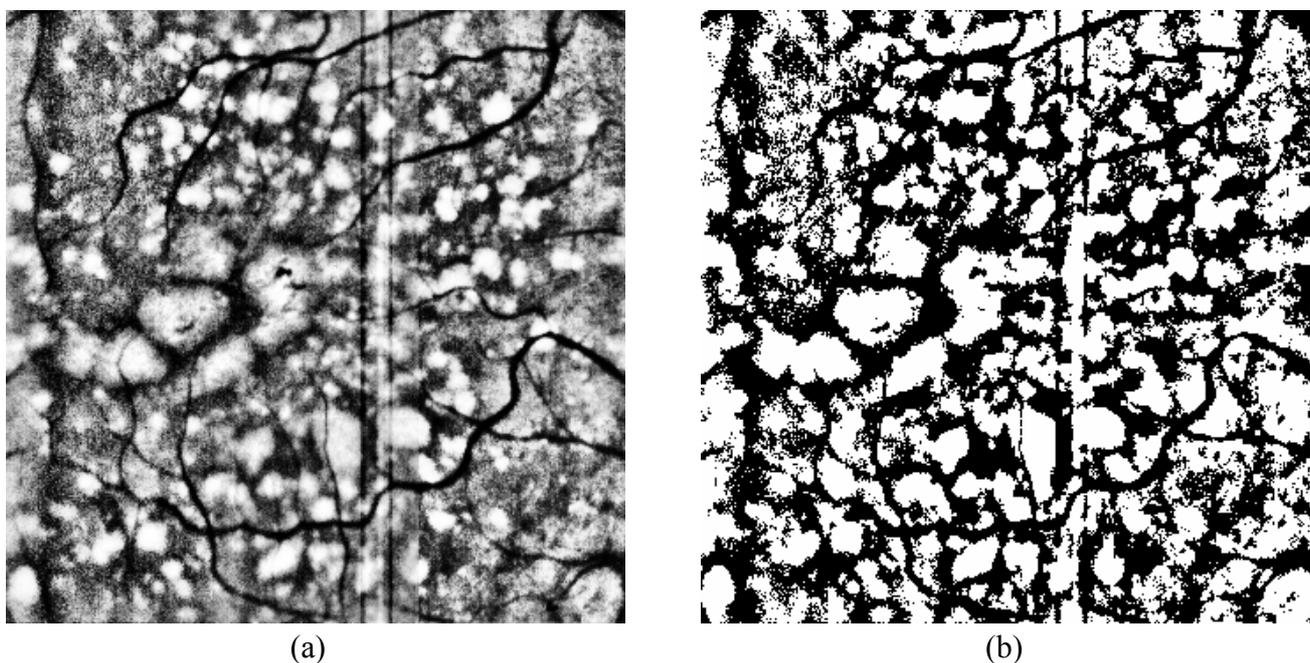
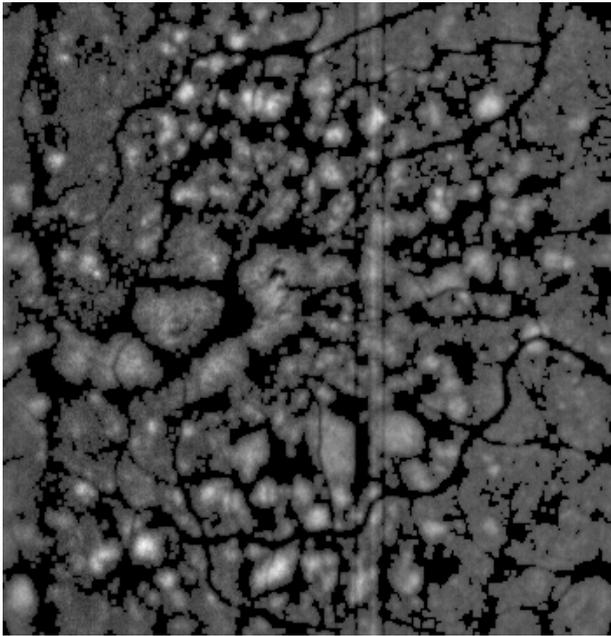


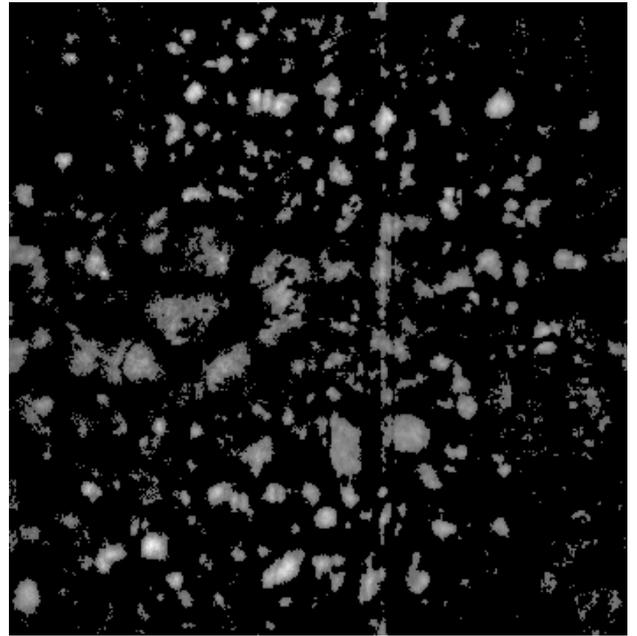
Fig. 4.2 (a) Image after multilevel histogram equalization; and (b) after global thresholding

Although, we experienced an increase in contrast after applying homomorphic filtering, further enhancement is needed, so as to achieve a correct and accurate segmentation. MHE enhances large and small drusen and makes them look brighter than other areas. As shown in Fig. 4.2-(a), diseased areas surrounded by uniform background are almost thoroughly defined and separated. Also, drusen that were vague become more visible. The disadvantage of histogram equalization and, consequently, of MHE, is the further enhancement of existing noise; as shown at the upper left-hand corner of Fig. 4.2-(a), several groups of pixels, not belonging to drusen, are strongly enhanced. These pixels must be removed during following algorithm steps or else they will lead to false estimation of diseased areas. Global thresholding, illustrated in Fig. 4.2-(b), succeeds in removing small parts of the background areas that are mainly located around small drusen. It fails in removing noisy areas, but it's still a useful step for getting rid of many background regions.

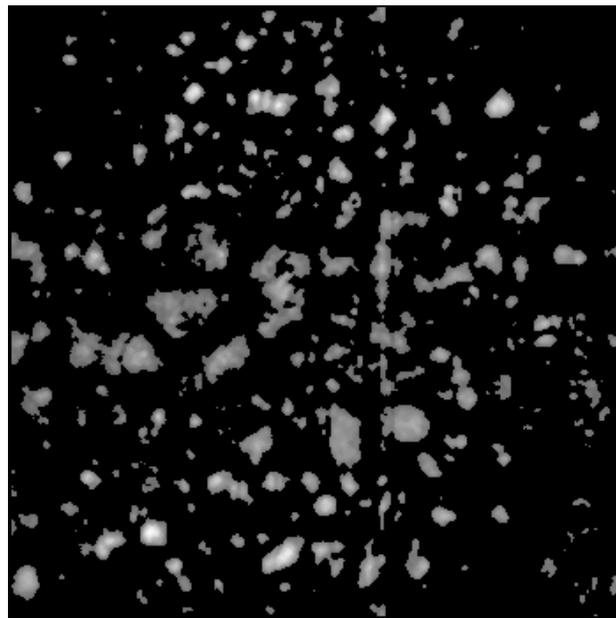
4.3 Morphological Dilation, HALT and Median Filtering



(a)



(b)



(c)

Fig. 4.3 (a) Image after morphological dilation; (b) image after HALT; (c) image after median filtering

Large drusen at the central part of the image are already good separated, so a thresholding operation in a small bounding area will fail to detect the whole object under consideration. It's obvious that dilation, Fig. 4.3-(a), results in a small expansion of these drusen coming into view by a retrieval of surrounding background areas. As a consequence HALT is successful in detecting a good threshold in the resulting area's histogram and separating almost the whole area of large drusen (Fig.

4.3-(b)). Smaller anomalies are also good segmented and it seems that false negatives exist only because of sparse 8-connected pixels. This problem is solved by median filtering, which removes these isolated pixels as shown in Fig. 4.3-(c), without affecting large connected regions. This image does not need expansion, since the criteria in section 3.4 –(A) are not met anywhere in the image.

4.4 Other Examples

Another example of an image that requires expansion of some regions after HALT is illustrated in Fig. 4.4-(a). This image contains large drusen that consist of bright and darker parts. Correct segmentation is hard to achieve, since setting whichever

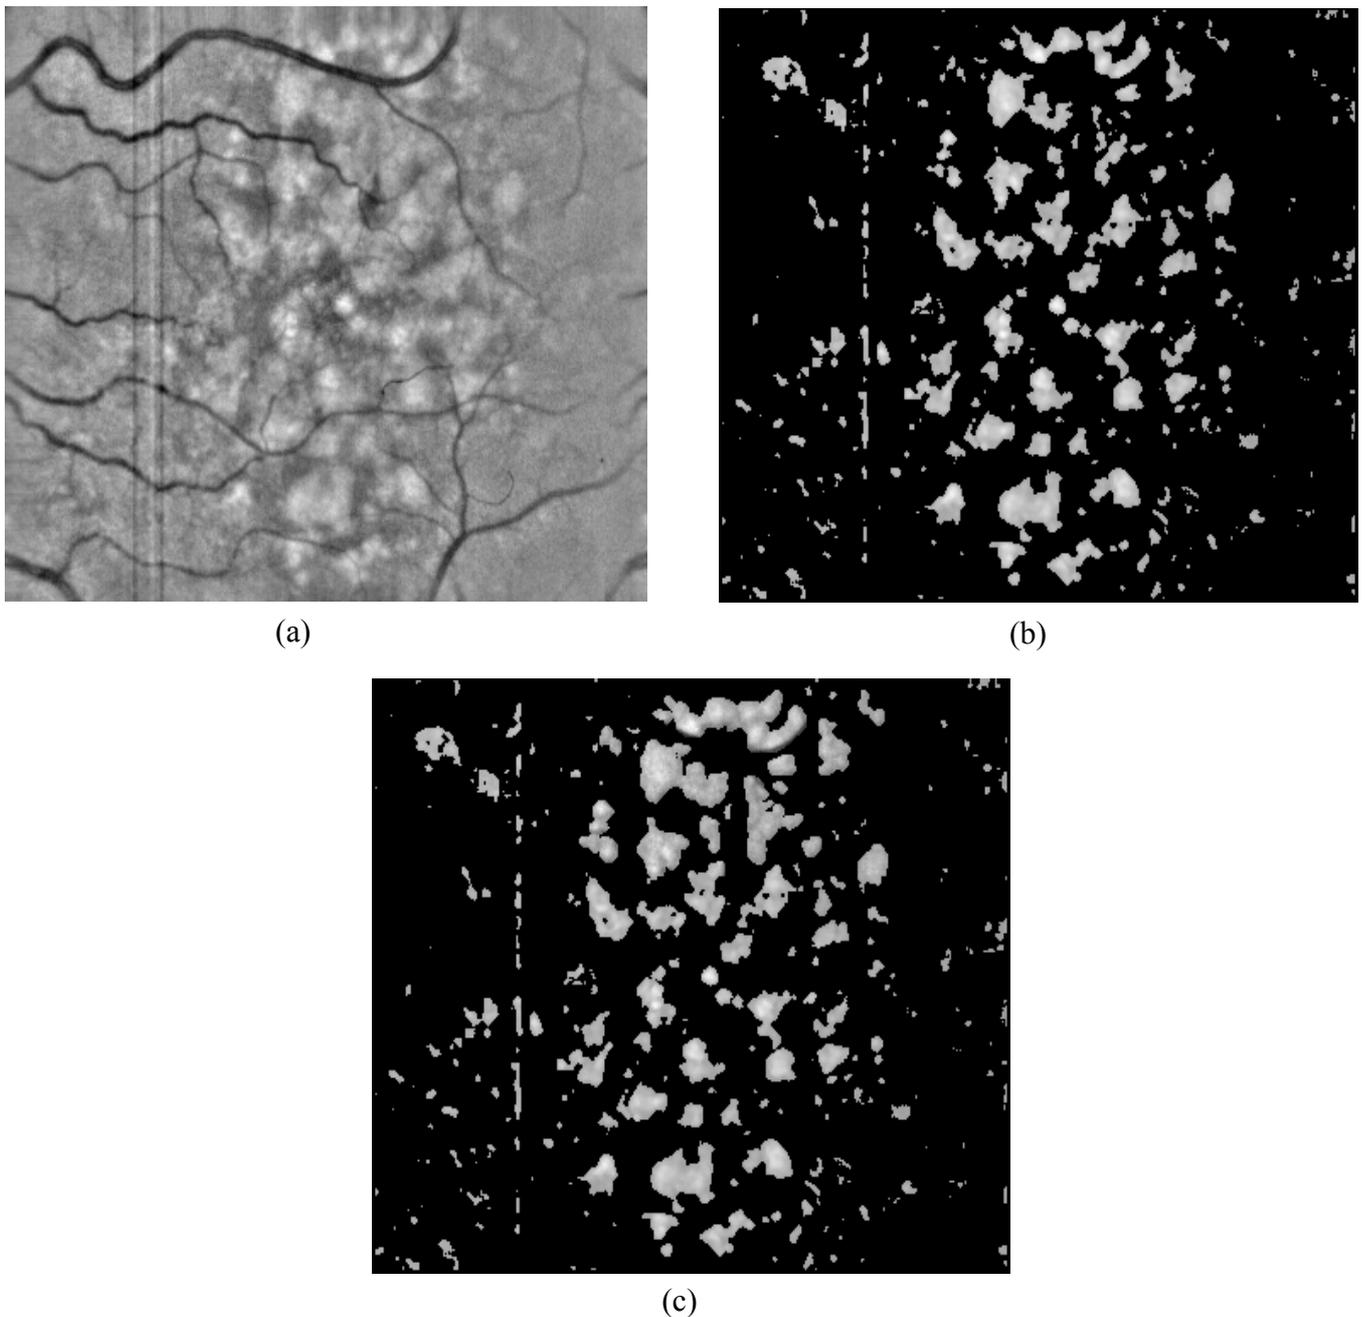
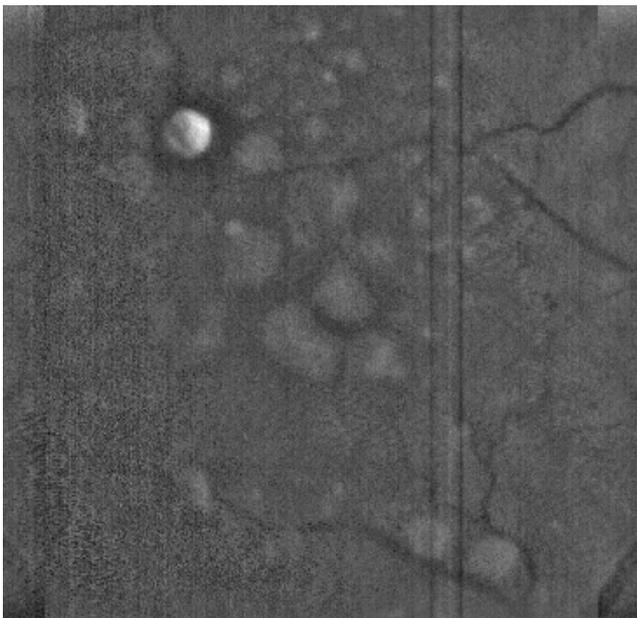


Fig. 4.4 (a) Original image; (b) image after HALT and elimination of sparse pixels; (c) image after expansion

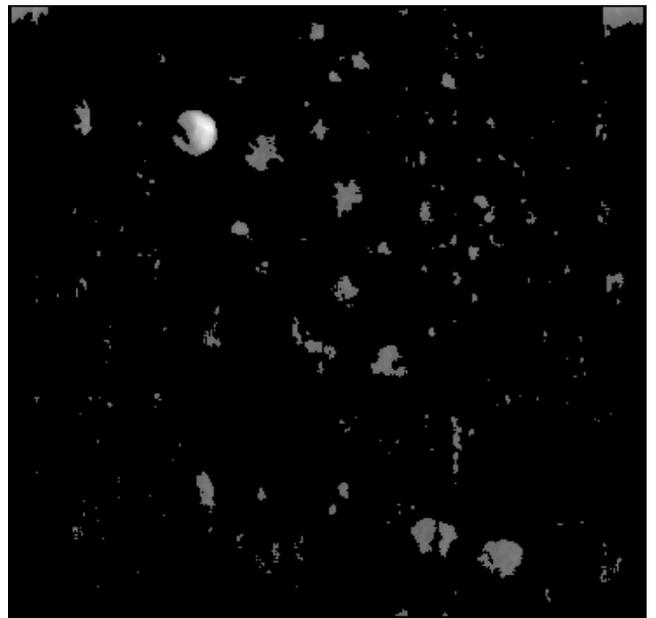
threshold inside an area that mainly consists of one drusen and almost no background will remove part of the anomalies.

In Fig. 4.4-(b) it is obvious that after HALT we are left with areas that must be joined or expanded, so as to recover missing parts of anomalies. As shown in Fig. 4.4-(c), after morphological dilation and closing, the upper areas that appeared “cracked” are joined together and form a single region that covers almost entirely the actual anomaly’s area.

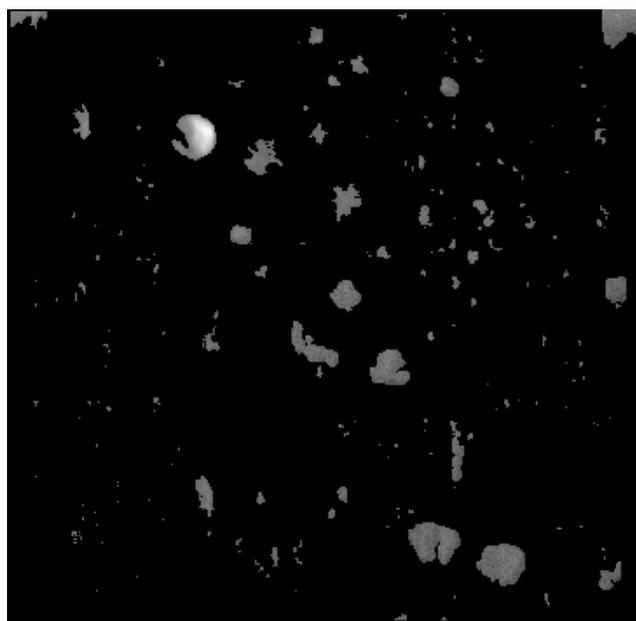
A hard to enhance image is shown in Fig. 4.5-(a). Presence of noise is strong, as it can be seen from the supposed background regions. Although, these areas had to be almost uniform, they appear noisy and non-homogeneous. In addition to that, large drusen do not differ sufficiently from the background.



(a)



(b)



(c)

Fig. 4.5 (a) Original image; (b) image after HALT and elimination of sparse pixels; (c) image after expansion

Except of the circular bright drusen, all others are noisy and intermixed with surrounding areas. Even at this case, our algorithm detects correctly all small drusen and loses few parts of bigger ones, as shown at the central part of Fig. 4.5-(c),(d).

In order to further demonstrate the efficiency of the proposed algorithm the results are subtracted from the original images, so that detected regions appear black. If parts of the drusen are not detected, they will appear bright, retaining their original gray level. As shown in Fig 4.6 for few representative images, almost all significant drusen have been detected.

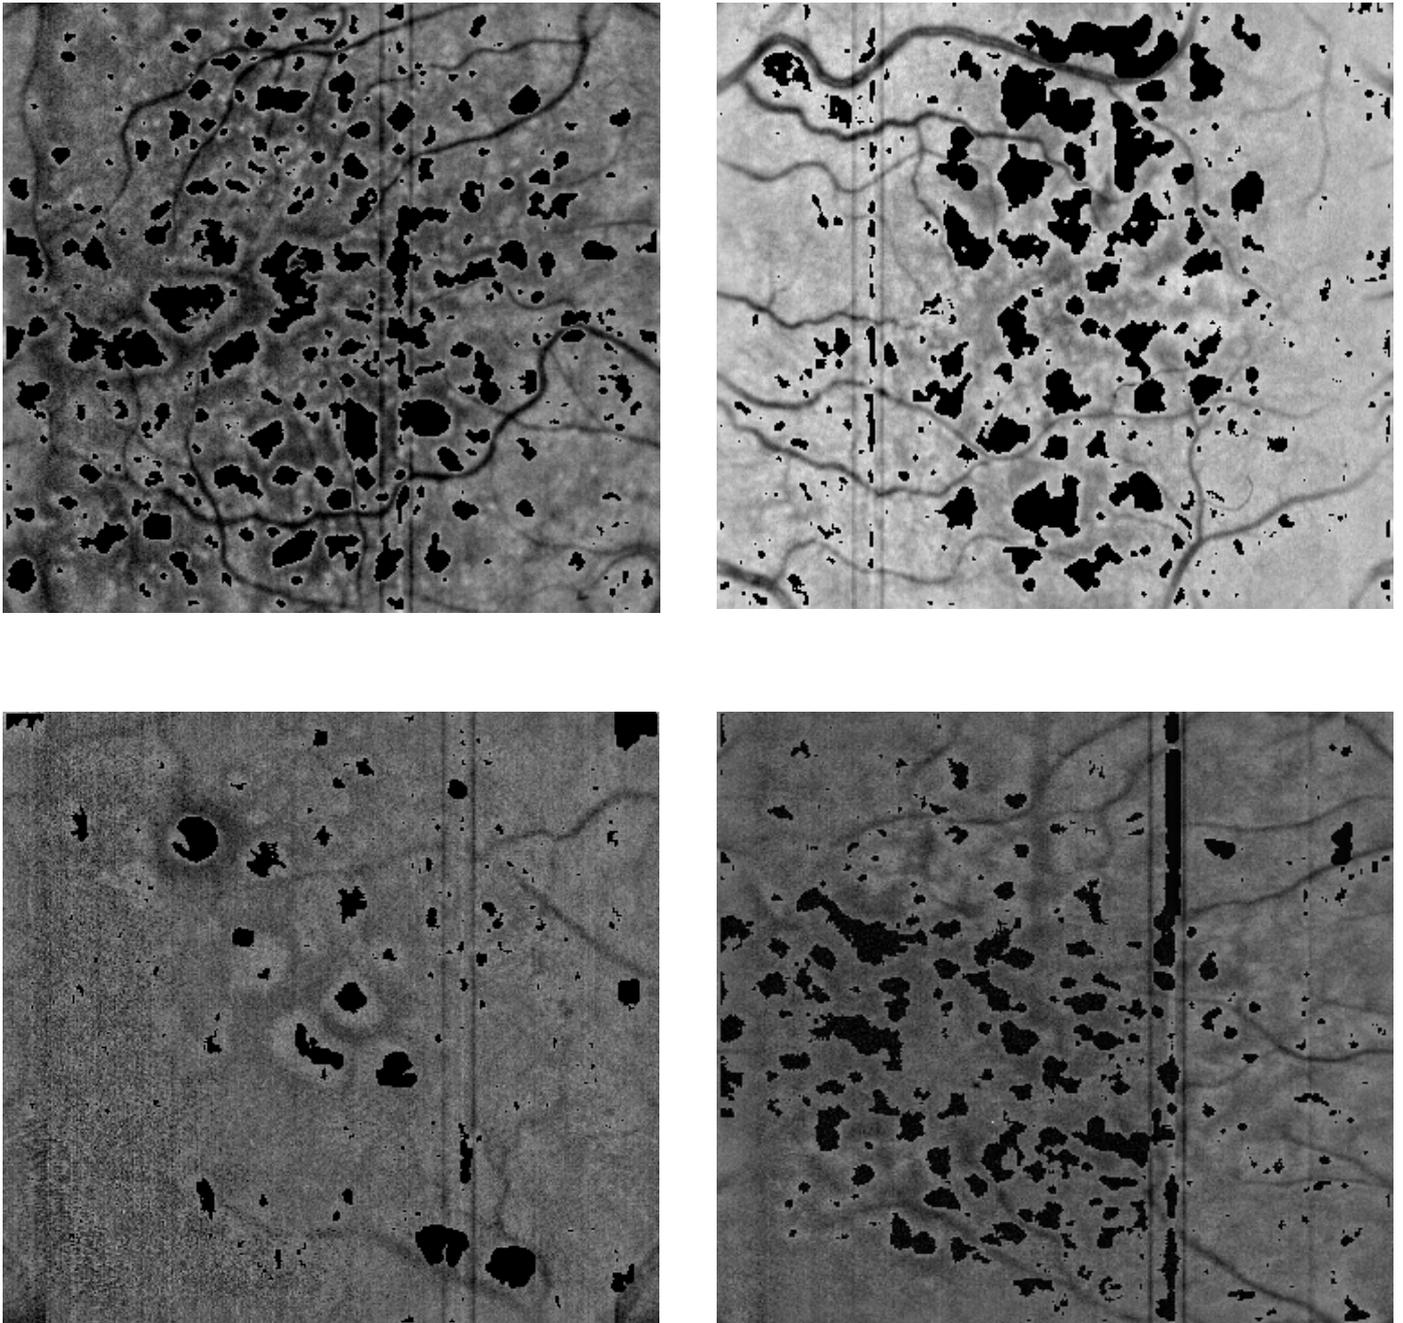


Fig. 4.6 Detected regions appear black, while not detected regions retain their actual gray level

4.5 General conclusion on the application

We tested our algorithm using a set of 23 images. 8 pairs of them were actually captured from the left and right eye of patients. We focused in the central part of the retina by defining a rectangle at the right or left side of the optical nerve (right or left eye). The proposed method was able to detect actual drusen in all cases. Even in hard-to-diagnose cases, where many small and vague drusen exist, our method succeeded in isolating them from background.

Initially, we experienced problems with the presence of vessels and their interaction with drusen. After applying the proposed algorithm we eliminate this problem and we don't experience false detection, due to vessels, in the entire test set of images. Large drusen, covering usually the central part of the eye, are hard-to-segment. Nevertheless, most of the images, containing such drusen, are correctly segmented and the drusen are accurately detected. Generally, our test set of images covers a wide range of possible drusen sizes and formations, including vague, non-canonical shaped and thin blobs. Thus, we faced the most common problems of macular degeneration and solved them in an efficient way.

Some other techniques we developed and tested for the same problem are presented in Appendix B. These methods, however, perform well only in specific images from the test set and require the tedious selection of parameters for their efficient operation.

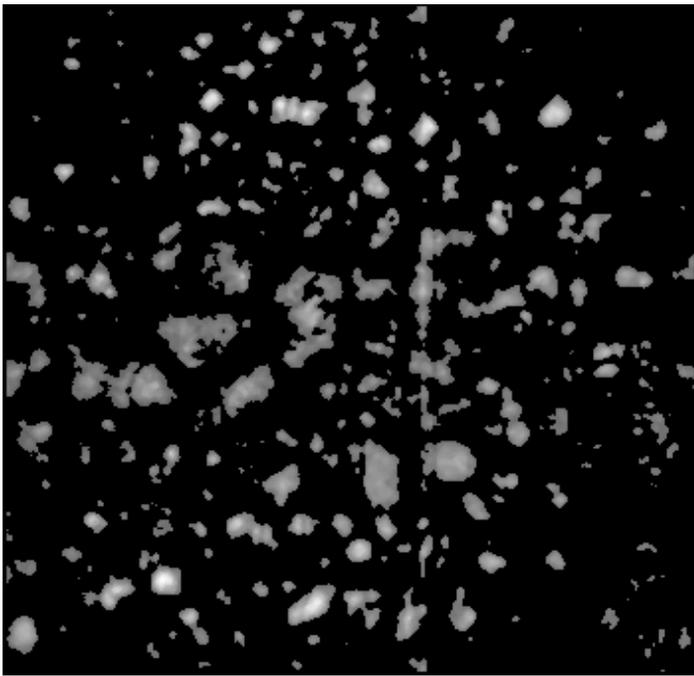
5. CONCLUSION AND FUTURE WORK

The detection of anomalies in human eye's retina is a biomedical problem, appropriate for image processing and automated pattern recognition, whose solution is intended to help the doctors in their decision making process. Use of the proposed detector may reduce false negatives and give reliable detection accuracy in both position and mass size.

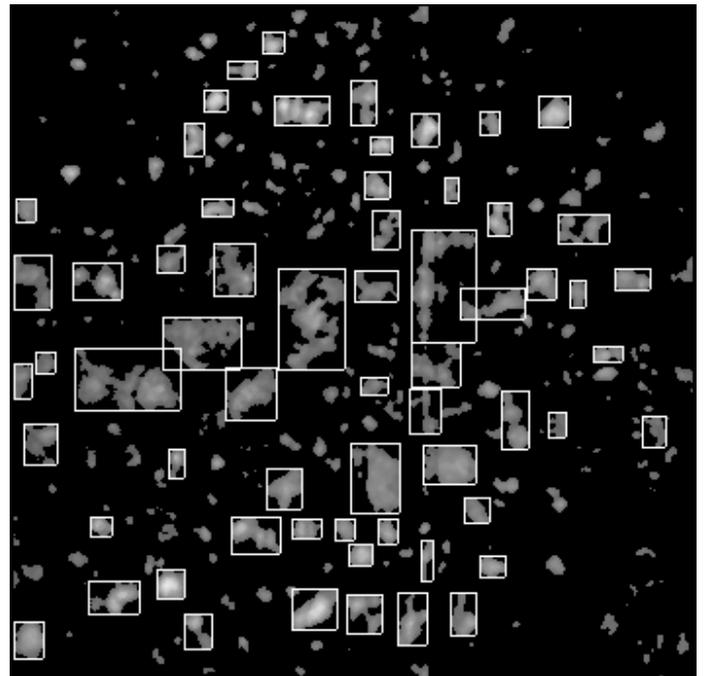
We started this project with very simple ideas such as the single thresholding technique and calculation of curvature or orientation of objects inside the image. However, as we got more involved in this project, we faced a lot of critical problems that must be solved to obtain good results, such as the non-uniform illumination correction, contrast enhancement, feature extraction, image segmentation and so on. Using existing techniques found in bibliography solved these problems. The ineffectiveness of such methods in dealing with the entire set of images led to the design of new operators. The most critical and challenging part was the final segmentation, achieved by the HALT algorithm. This is a novel method for segmenting areas that differ slightly from their background regions. It is fully unsupervised and non-parametric. Furthermore, it is easily adjusted to other segmentation problems, where anomalies or defects are not brighter than the background.

The proposed method was able to detect actual drusen in all cases. Even in hard-to-diagnose cases, where many small and vague drusen exist, our method succeeded in isolating them from the background. A significant factor that affects the overall performance of other approaches is the presence of noise, which makes surfaces look rough and renders the segmentation process difficult. Although, it is not a common case, since the presence of noise is rare in such images (only one in our test set), our method provides adequate results even in the case of noise contamination.

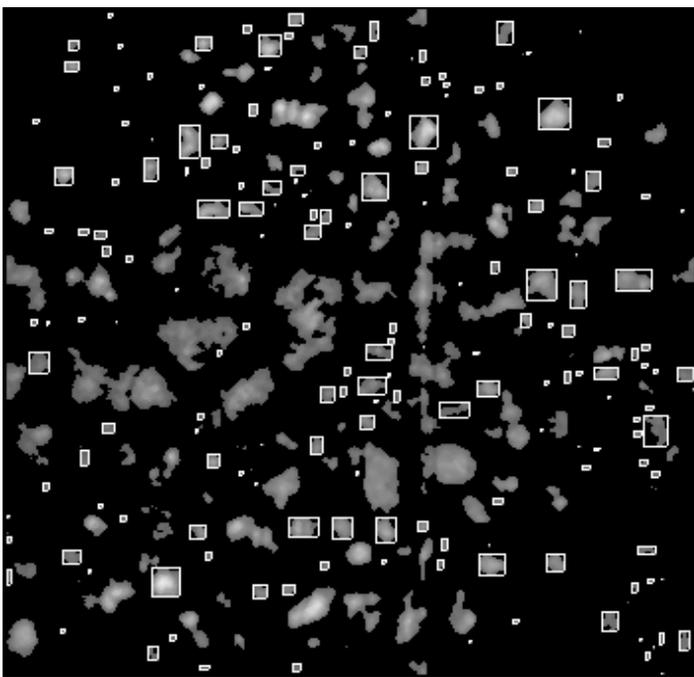
Further investigations could be focused towards increasing the mass-location accuracy and reducing the number of false diagnoses. The latter objective requires the correct characterization of anomalies according to their shapes or masses. We tried different features at the detection phase (see Appendix C), but we need an expert's advice so as to use the correct combination of them. Some examples are shown in Fig. 5.1. It is obvious that the features used in the example are not capable of producing a robust result. The shape and size of anomalies are irregular and, therefore, cannot be described by features addressing only circularity or mass size.



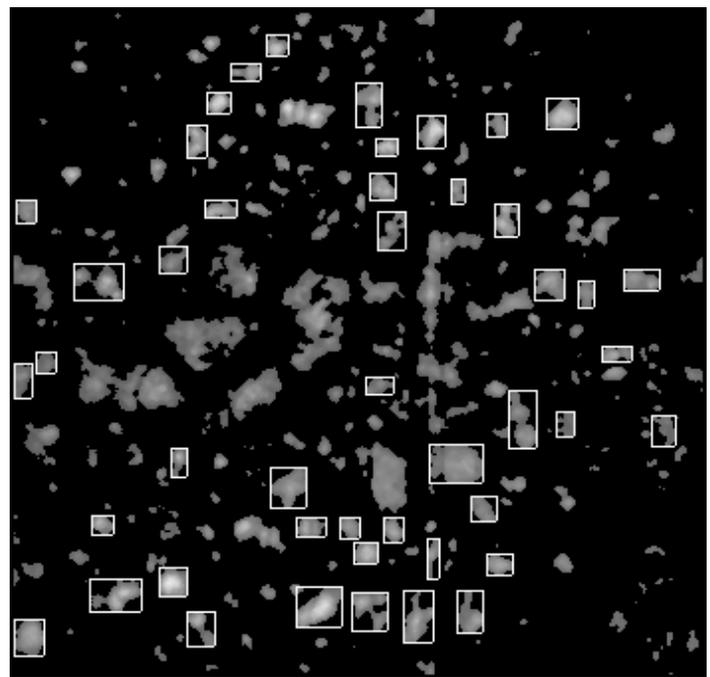
(a)



(b)



(c)



(d)

Fig. 5.1 (a) result of proposed method; (b) selection of anomalies according to area; (c) selection of anomalies according to compactness; (d) selection of anomalies according to area + compactness

The use of a neural network classifier can also be investigated to provide fast and accurate results regarding the presence and shape of anomalies. The fast detector of

masses in computer-aided mammographies presented in [11], uses a RBFNN (Radial Basis Function Neural Network) as a classifier (Fig. 5.2) in conjunction with three decision criteria to avoid false diagnoses. The classifier's inputs are the first histogram moments, namely mean, variance, and the higher-order statistics, namely skewness and kurtosis (see Appendix A).

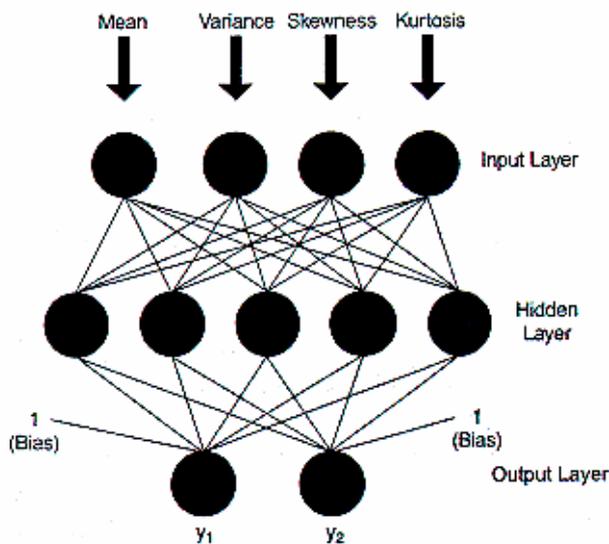


Fig. 5.2 The RBFNN classifier

The cork quality classification system presented in [9], uses a Fuzzy-Neural network methodology for classifying corks. The utility of fuzzy sets ([38],[39]) lies in their ability to model the uncertain or ambiguous data to often encountered in cork classification problem. Finally, an MLP (Multi Layer Perceptron) classifier is used with a back-propagation training algorithm, which incorporates concepts from fuzzy sets at the training stage.

Such NN and FS approaches could be also used in our problem in a more robust feature classification set-up, rather than the pixel classification scheme developed in this thesis.

APPENDIX A HISTOGRAM PROPERTIES

Features for Histogram Analysis

The histogram of an image with gray levels in $[0, L-1]$ is a discrete function $p(r_k) = n_k/n$, where r_k is the k th gray level, n_k is the number of pixels in the image with that gray level, n is the total number of pixels in the image, and $k = 0, 1, 2, \dots, L-1$.

Under the ergodicity assumption $p(r_k)$ provides an estimate of the probability of occurrence of gray level r_k . A plot of this function for all values of k can provide a crude description of image's appearance. Four simple examples are shown in Fig. A.1.

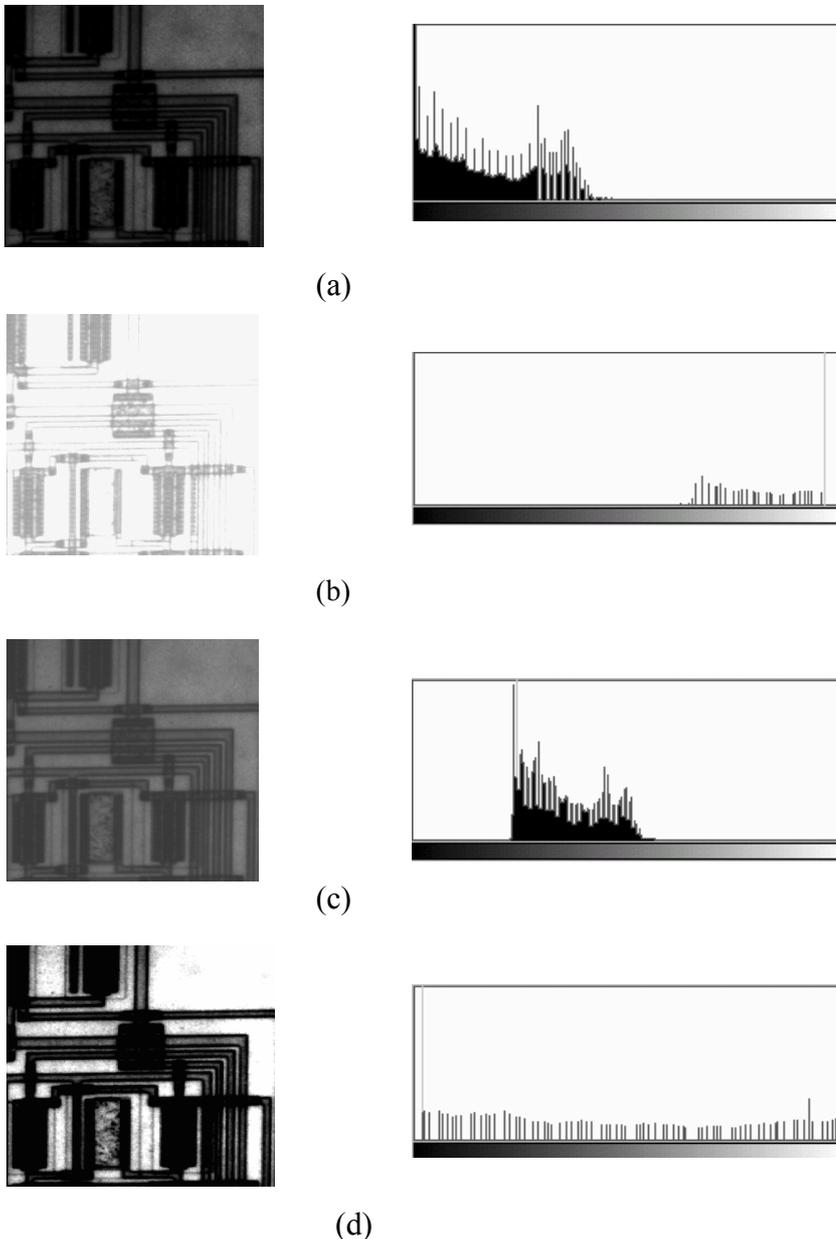


Fig. A.1 (a) dark image; (b) bright image; (c) low contrast image; (d) high contrast image

It's obvious that the spread of values and the shape of the gray level distribution provide useful information about the content of an image. In order to quantify this information, several features can be extracted. The most often used features are presented in the following.

· **Mean**

$$S_M \equiv \bar{b} = \sum_{b=0}^{L-1} bP(b)$$

A measure for average brightness of an image

· **Standard Deviation**

$$S_D \equiv s_b = \left[\sum_{b=0}^{L-1} (b - \bar{b})^2 P(b) \right]^{1/2}$$

A measure of spread about the mean. The standard deviation since it is calculated from “squared deviations” (deviation from the mean value) is more sensitive to a few extreme observations than is the mean.

For a normal distribution (which is generally the case), S_D provides useful information about the gray level distribution:

- About 68% of data lies within one standard deviation of the mean.
- About 95% of data lies within two standard deviations of the mean.
- Almost all data lies within three standard deviations of the mean.

· **Median**

Median is that pixel amplitude for which one half of the pixels are equal or smaller in amplitude and one half is greater in amplitude.

· **Skewness**

$$S_S = \frac{1}{s_b^3} \sum_{b=0}^{L-1} (b - \bar{b})^3 P(b)$$

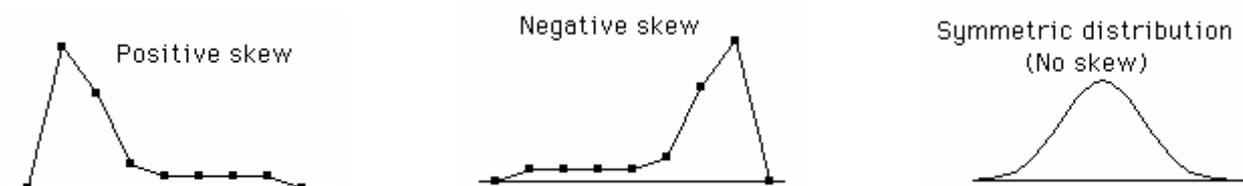


Fig. A.2 Histogram shapes according to their skewness

A distribution is skewed if one of its tails is longer than the other, as illustrated in Fig. A.2. Positive skew means a long tail in the positive direction. Negative skew means a long tail in the negative direction. Zero skew means symmetric distribution.

· **Kurtosis**

$$S_K = \frac{1}{s_b^4} \sum_{b=0}^{L-1} (b - \bar{b})^4 P(b) - 3$$

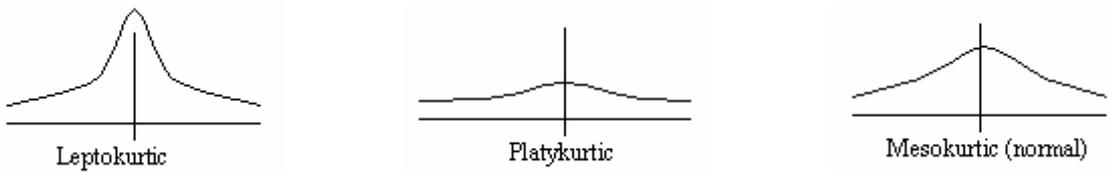


Fig. A.3 Histogram shapes according to their kurtosis

Kurtosis is based on the size of a distribution’s tails, as shown in Fig. A.3. Distributions with relatively large tails (sharp-peaked) are called “leptokurtic”; those with small tails (flat topped) are called “platykurtic”. A distribution with the same kurtosis as the normal distribution is called “mesokurtic”. In order to distinguish between the three cases the following rule is valid

- $S_K = 0$ mesokurtic
- $S_K > 0$ leptokurtic
- $S_K < 0$ platykurtic

· **Energy**

$$S_N = \sum_{b=0}^{L-1} [P(b)]^2$$

· **Entropy**

$$S_E = - \sum_{b=0}^{L-1} P(b) \log_2 [P(b)]$$

Raising S_E characterizes flat-topped distributions, while decreasing S_E characterizes sharp-peaked distributions.

· **Peak**

Peak is the pixel amplitude corresponding to the most commonly occurring pixel amplitude in the area.

APPENDIX B OTHER DETECTION APPROACHES

Several other approaches towards the detection and localization of drusen are presented in this appendix.

B.1 Detection Based on Gradient Information

The schematic diagram of this approach is shown in Fig. B.1.

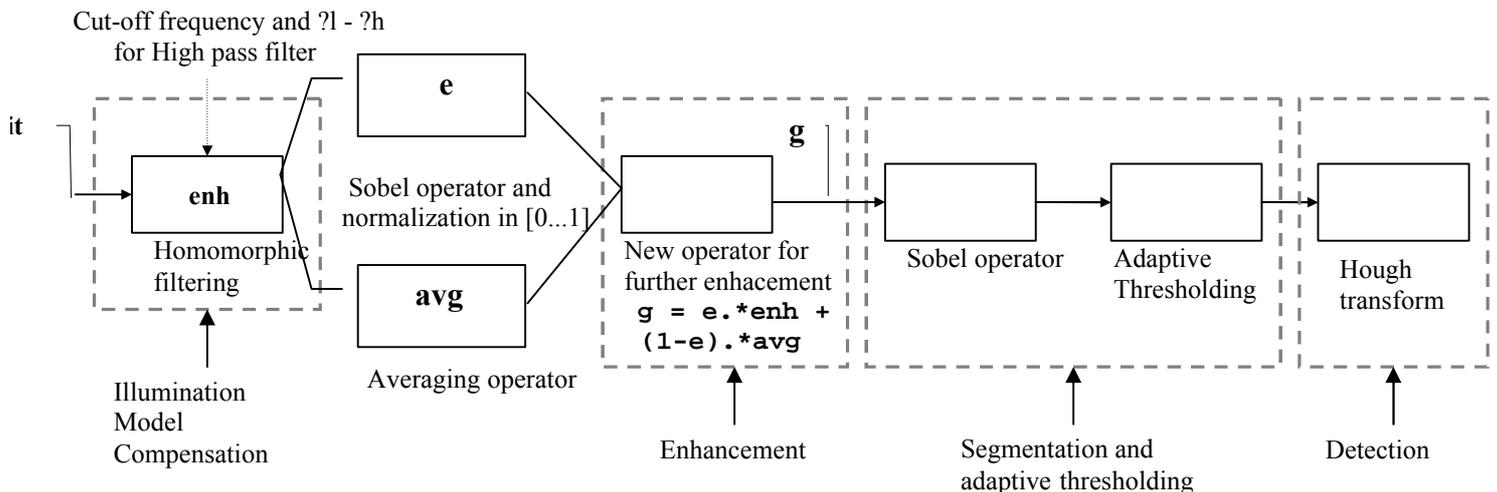


Fig. B.1 Algorithm

B.1.1 Illumination model compensation

Illumination model compensation is performed by border averaging and homomorphic filtering, as described in 4.1.

B.1.2 Enhancement

The background areas are supposed to be uniform in intensity, so the presence of edges signifies areas of drusen or vessels. Although, we are not interested in vessels, it would be useful to have a gross discrimination of background and non-background regions. Thus, we use a simple operator that enhances possible edges and suppresses background, which is based on input from a Sobel and an averaging operator. Sobel is responsible for providing information on edges' location, while the averaging mask generates a smooth estimate of background, as illustrated in Fig. B.2. The enhancement operator is

$$g = e.*enh + (1-e).*avg, \text{ where}$$

- .* : point-to-point operator
- enh: result of Sobel operator (edges)
- avg: result of averaging operator (estimate of background)

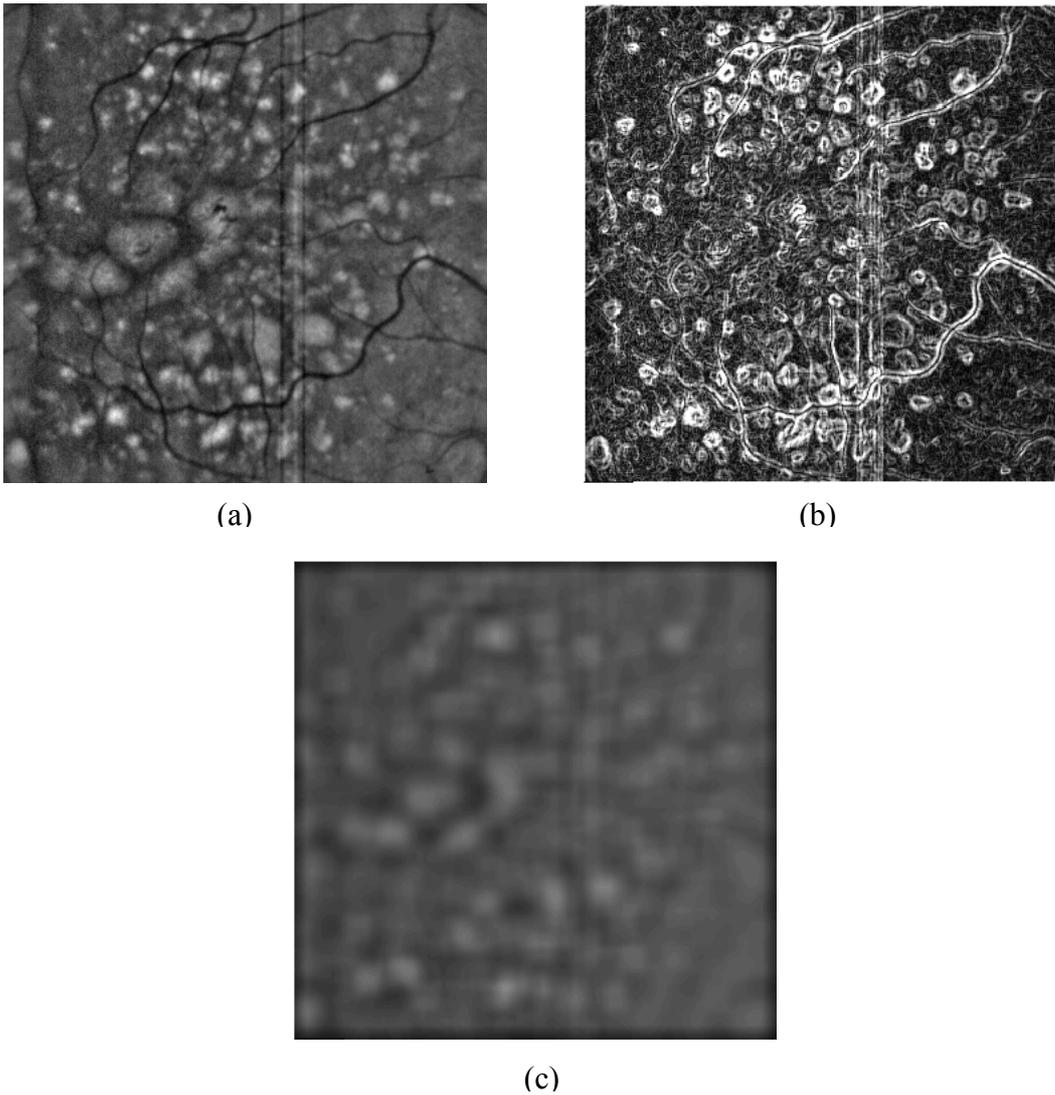


Fig. B.2 (a) original image; (b) application of Sobel operator; (c) application of average mask

B.1.3 Segmentation and adaptive thresholding

The result of previous operations is an image with a more uniform background than before and emphasized edges, as illustrated in Fig. B.3-(a). Another Sobel operator is applied that detects less edges, since the actual background is more homogenous, as shown in Fig. B.3-(b). False negatives are mostly generated from non-uniform areas, located inside large blobs, and pixels belonging to vessels.

Correct segmentation is a difficult task, since many drusen-candidate edges exist. A adaptive thresholding scheme is employed to separate actual drusen and remove most of the remained vessels. It requires two parameters: a local threshold value and a global one. The algorithm is described below:

- Find local maxima within specified areas (e.g. 10x10)

- Select maxima higher than local threshold
- Multiply them with global threshold

B.1.4 Detection

Drusen are generally circular or ellipsoid or have a non-canonical shape. Thus, the detection cannot be based on a certain shape such as circle or ellipse. Nevertheless, we tried to test the standard Hough transform with a small modification on resulting images.

Following parameters are used for circle detection using Hough Transform:

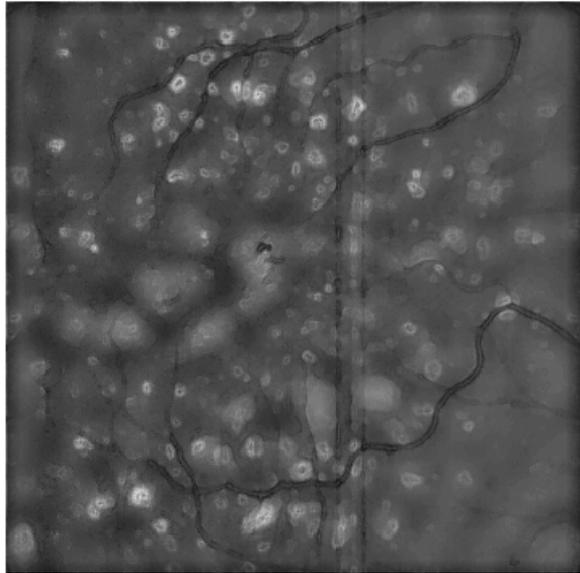
1. Size of square window to search for local maxima
2. Minimum number of updates for a maximum to be used. The minimum number is determined from UpdateLimit and the radius corresponding to the current band:

$$\text{updateLimit} * 2.0 * \text{Pi} * \text{radius}$$

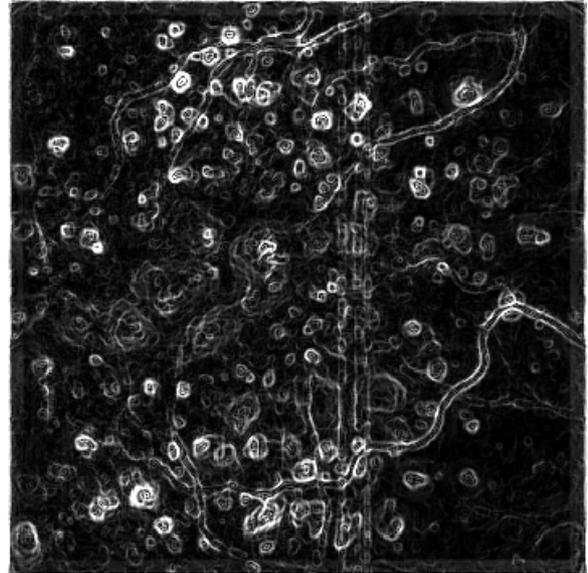
3. Min & Max Radius of Circles to be detected
4. Size of square window to search for global Maxima within all bands of Hough transform

B.1.5 Conclusion

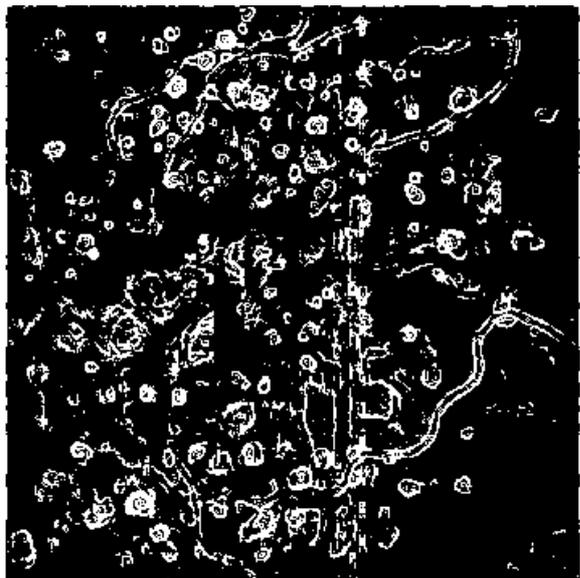
The major disadvantage of this approach is the inability to enhance the borders of large drusen. These borders are usually smooth and do not generate strong edges after the Sobel operation. As a consequence, the adaptive thresholding scheme is not capable of segmenting large drusen although it does work well for smaller drusen. This scheme also fails in generating closed contours (Fig. B.3–(c)), because of brightness variation in edges detected by Sobel (Fig. B.3–(b)). The results are improved by using the Hough transform (Fig. B.3–(d)) for detecting the actual anomalies. After regulating the required parameters, the transform becomes elastic in open or broken contours and succeeds in detecting many anomalies. However, the parameter regulation is a hard-to-perform and the algorithm's results are not always satisfactory. Many small drusen are not detected and false negatives are present, because of false contours generated during previous steps.



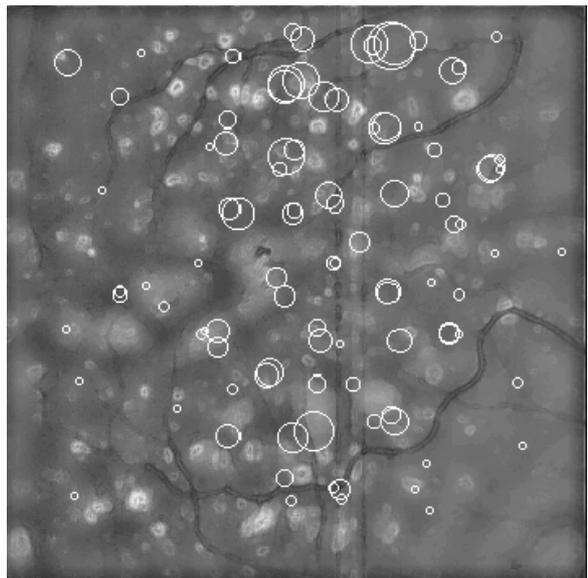
(a)



(b)



(c)



(d)

Fig. B.3 (a) Result of applied operator; (b) image after Sobel operator; (c) adaptive thresholding; (d) Hough Transform

B.2 Otsu thresholding scheme

During our experimentation with different thresholding schemes we tried the following approach based on Otsu thresholds.

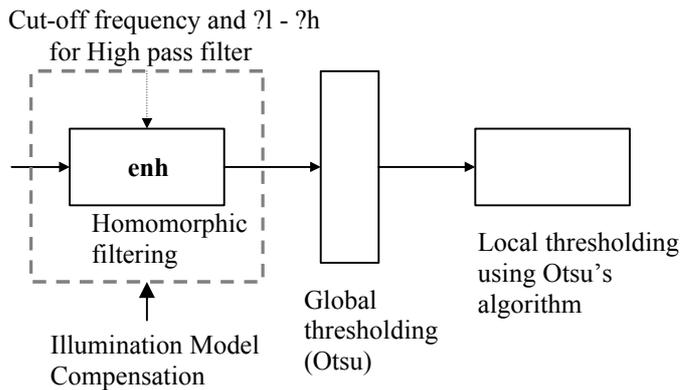


Fig. B.4 Otsu thresholding scheme

Homomorphic filtering is an inseparable step of our case-study. Global Otsu thresholding is applied to remove parts of the background that surrounds abnormalities. A more refined segmentation is needed, in order to separate drusen and detect them correctly. We apply the Otsu thresholding scheme in small windows. Our purpose is to segment anomalies without being misled by remaining bright areas of the background, which owe their existence to global non-uniform illumination.

The result is almost satisfactory for images with many and large drusen (Fig. B.5 – (a),(b)). Such drusen are bright and easy to segment when processing image in small blocks. However, as shown in Fig. B.5-(c),(d), when few and vague drusen are present the result is disappointing. Many background areas are mistaken as macular degeneration symptoms, since they look a bit brighter than usual.

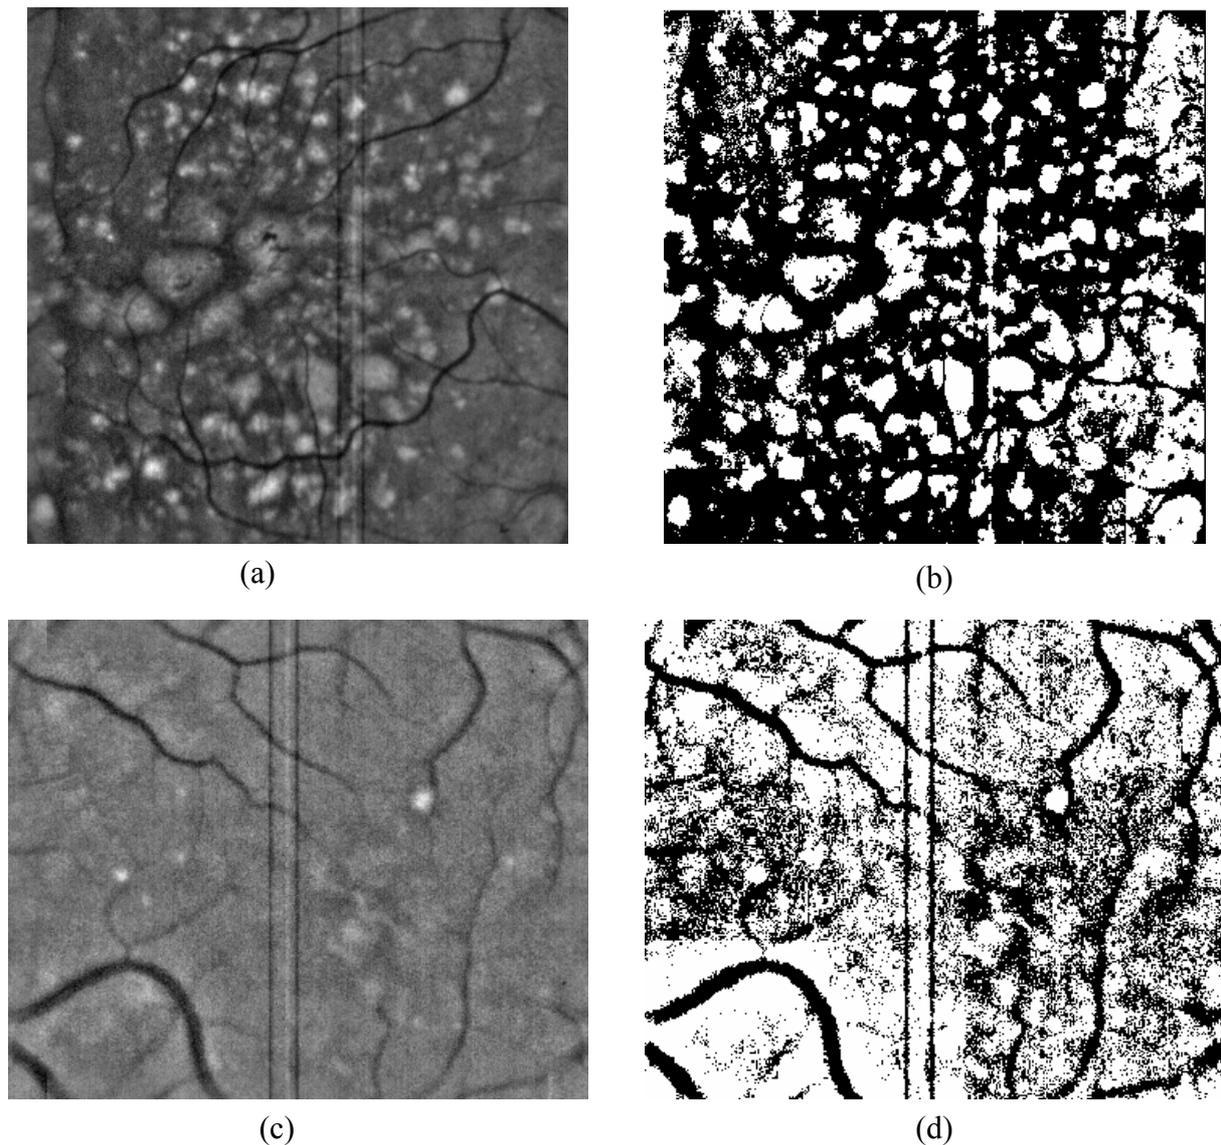


Fig. B.5 Images before and after application of Otsu thresholding scheme

B.3 Template matching and Curvature estimation

An advantage of examining the biomedical images at hand is the knowledge acquired for structures of certain size and orientation that must be detected. The existence of certain shapes to be localized makes template-matching tempting. In addition, linearity or circularity of structures is useful for distinguishing objects of interest. In our case, many drusen look circular as opposed to vessels that look linear.

In this approach we attempted to extract drusen or vessels based on their shape. Directional filters, like Kirsch and Nevatia-Babu [2], were used in order to achieve this goal. Unfortunately, brightness variation of the background and non-canonical shapes of objects prevented the deviation of good results. For the same reason another

idea based on curvature estimation failed. This technique is briefly described in the following.

We use the Sobel operator to approximate curvature. After running two orthogonal directional derivatives (vertical-horizontal masks) two vectors are generated – call them DX and DY. We derive line magnitude $m = \text{sqrt}(DX^2 + DY^2)$, and the orientation image $o = a \tan(DY/DX)$. Subsequently we use these measures in the following way, to estimate where strong (thresholdable) lines can be identified. We compute the variance of the orientation image and use the Sobel magnitude and the Sobel orientation-variance images. The variance image is directly correlated to curvature and the magnitude image indicates where line strength is located. If we use a threshold on both images, the areas of high curvature (high variance) AND areas of high magnitude will fit our criteria of drusen detection as opposed to vessel detection.

APPENDIX C NUMERICAL FEATURES FOR CLASSIFICATION OF DRUSEN

Finding numerical descriptors of shape is a difficult task, since there are dozens of possible size parameters that can be calculated. A shape descriptor is actually a combination of these parameters that finally form a dimensionless expression. Although common combinations are few, there exist an inconsistency in naming conventions. By using the conventions given in [3], we provide the formulas of few shape descriptors and definitions of some parameters.

$$\text{Formfactor} = \frac{4p \cdot \text{Area}}{\text{Perimeter}^2}$$

$$\text{Roundness} = \frac{4 \cdot \text{Area}}{p \cdot \text{MaxDiameter}^2}$$

$$\text{AspectRatio} = \frac{\text{MaxDiameter}}{\text{MinDiameter}}$$

$$\text{Convexity} = \frac{\text{ConvexPerimeter}}{\text{Perimeter}}$$

$$\text{Solidity} = \frac{\text{Area}}{\text{ConvexArea}}$$

$$\text{Compactness} = \frac{\sqrt{\left(\frac{4}{p}\right) \text{Area}}}{\text{MaxDiameter}}$$

$$\text{Extent} = \frac{\text{NetArea}}{\text{BoundingRectangle}}$$

Area:

Object's area in pixels

Perimeter:

Actual perimeter of object

Convex Perimeter:

Perimeter of polygonal approximation of object's boundary

Convex Area:

Area include in previous polygon

Max-Min Diameter:

Maximum (minimum) chord inside convex polygon

Net Area:

Object's area without counting internal "holes"

REFERENCES

- [1]. Gonzalez R.C., and Woods R.E. [1993]. "Digital Image Processing"
- [2]. Pratt W.K. [1991]. "Digital Image Processing"
- [3]. Russ J.K. [1999]. "The Image Processing Handbook"
- [4]. Haralick R.M and Shapiro L.G. "Image Segmentation Techniques", Computer Vision, Graphics and Image Processing, vol. 29, no.1, pp. 100-132, January 1985.
- [5]. Otsu Nobuyuki ."A Threshold Selection Method from gray level Histograms", IEEE Transactions on Systems, Man and Cybernetics, vol. SMC-9, no.1, January 1979
- [6]. C.K. Chow and T. Kaneko. "Automatic Boundary Detection of the left Ventricle from Cineangiograms", Comput. Biomed. Res., vol. 5, pp. 388-410, 1972.
- [7]. Sahoo P.K., Soltani S. and Wong K.C.. "A Survey of Thresholding Techniques"
- [8]. Boukharouba S., Rebordao J.M. and Wendel P.L.. "An Amplitude Segmentation Method based on the Distribution Function of an Image", Computer Vision Graphics, and Image Processing, vol. 29, pp. 47-59,1985.
- [9]. Joongho Chang, Gunhee Han, Hose M. Valverde, Norman C. Grisworld, J. Francisco Duque-Carillo, Edgar Sanchez-Sinencio. "Cork Quality Classification System using a Unified Image Processing and Fuzzy-Neural Network Methodology", IEEE Transactions on Neural Networks, vol. 8, no. 4,July 1997
- [10]. Xu L., Jackowski M., Goshtasby A., Roseman D., Bines S., Yu C., Dhawan A., Huntley A. "Segmentation of Skin Cancer Images", Image and Vision Computing, vol. 17, pp. 65-74,1999.
- [11]. Christoyianni Ioanna, Dermatas Evagelos, Kokkinakis George. "Fast Detection of Masses in Computer-Aided Mammography", IEEE Signal Processing Magazine, January 2000
- [12]. Rene G. Aarnick, Sayan Dev Pathak, Jean J.M.C.H. de la Rosette, Frans M.J. Debryune, Yongmin Kim, Hessel Wijkstra. "Edge Detection in Prostatic Ultrasound images using integrated Edge Maps", Ultrasonics, vol. 36, pp. 635-642, 1998.
- [13]. Aarnick R.G., "Automated Contour Detection; Design and Clinical Applications in Urology", Ph. D. Thesis, University of Nijmegen, The Netherlands, 1996
- [14]. Omer Demirkaya, Robert M. Cothren, D. Geoffrey Vince and J. Fredrick Cornhill. "Automated Identification of stained cells in Tissue Sections Using Digital

- Image Analysis",
Analytical and Quantitative Cytology and Histology, vol. 21, no.2, April 1999
- [15]. Riccardo Poli, Guido Valli. "An Algorithm for real-time Vessel Enhancement and Detection",
Computer Method. and Programs in Biomedicine, vol. 52, pp. 1-22, 1997.
- [16]. J.F. Canny. "Finding Edges and Lines in Images",
Technical Report 720, Artificial Intelligence Laboratory -Massachusetts Institute of Technology, June 1983
- [17]. Ahmed S. Abutaleb. "Automatic Thresholding of gray-level Pictures using Two Dimensional Entropy",
Computer Vision, Graphics and Image Processing, vol. 47, pp. 22-32, 1989.
- [18]. E. David Jansing, Thomas A. Albert, Darrel L. Chenoweth. "Two Dimensional Entropic Segmentation",
Pattern Recognition Letters, vol. 20, pp. 329-336, 1999.
- [19]. A. D. Brink. "Thresholding of Digital Images using Two Dimensional Entropies",
Pattern Recognition, vol. 25, no.8, pp. 803-808, 1992.
- [20]. Andrew K. C. Wong, P. K. Sahoo. "A gray Level Threshold Selection Method Based on Maximum Entropy Principle",
IEEE Transactions on Systems, Man and Cybernetics, vol. 19, no.4, July/August 1989
- [21]. Kittler J., Illingworth J. "Minimum Error Thresholding", Pattern Recognition, vol.19, no.1, pp. 41-47, 1986,.
- [22]. Mardia K.V., Hainsworth T.J. "A Spatial Thresholding Method for Image Segmentation",
IEEE Transactions on Pattern Analysis and Machine intelligence, vol. 10, no.6, pp. 919-927 Nov 1988.
- [23]. J, Illingworth, J. Kittler. "Survey of the Hough Transform",
Computer Vision, Graphics and Image Processing, vol. 44, pp. 87-116, 1988.
- [24]. R.O. Duda, P.E. Hart. "Use of the Hough Transformation to Detect Lines and Curves in Pictures",
Communications of the Association of Computing Machinery, vol. 15, pp. 11-15, 1972
- [25]. T.J. Atherton, D.J. Kerbyson. "Size Invariant Circle Detection", Image and Vision Computing, vol. 17, pp. 795-803, 1999.
- [26]. Peng- Yeng Yin. "A new Circle/Ellipse Detector using Genetic Algorithms",
Pattern Recognition Letters, vol. 20, pp. 731-740, 1999.

- [27]. Dimitrios Ioannou, Walter Huda, Andrew F. Laine. "Circle Recognition through 2D Hough Transform and Radius Histogramming", Image and Vision Computing, vol. 17, pp. 15-26,1999.
- [28]. J.S. Lee. "Digital Image Enhancement and Noise Filtering by Use of Local Statistics".
- [29]. G.Deng, L. W. Cahill and G.R. Tobin."The Study of Logarithmic Image Processing Model and Its Application to Image Enhancement" .
- [30]. R. Southwell. "Relaxation Methods in Engineering Science", A Treatise on Approximate Computation, Oxford Univ. Press, 1940.
- [31]. R. Southwell. "Relaxation Methods in Theoretical Physics", Oxford Univ. Press (Clarendon), London, 1946.
- [32]. A. Rosenfeld and R.C. Smith. "Thresholding using Relaxation", IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. P AMI-3, pp. 598-606, 1981
- [33]. G. Fekete, J.O. Eklundh and A. Rosenfeld. "Relaxation: Evaluation and Applications ", IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. PAMI-3, pp. 460-469, 1981
- [34]. S. Zucker, R. Hummel and A. Rosenfeld. "An Application of Relaxation Labelling to Line and Curve Enhancement", IEEE Transactions on Computing, vol. C-26, pp. 394-403, 1977
- [35]. T. Pavlidis. Structural Pattern Recognition, Springer-Verlag, New York, 1977.
- [36]. S. Peleg. " A new Probabilistic Relaxation Scheme", IEEE Transactions on Pattern Analysis and Machine Intelligence.
- [37]. R. Gordon and R.M. Rangayyan. "Feature Enhancement of Film Mammograms using Fixed and Adaptive Neighborhoods", Applied Optics, vol. 23, pp. 560-564, 1984
- [38]. H. J. Zimmerman. "Fuzzy Set Theory and its Applications", Boston, MA: Kluwer, 1991
- [39]. R. J. Marks. "Fuzzy Logic Technology and Applications", IEEE Technology Update Series- New York: IEEE Technical Activity Board, 1994