ELSEVIER

Medical Image Analysis 7 (2003) 95–108

MEDICAL IMAGE ANALYSIS

www.elsevier.com/locate/media

Detection and segmentation of drusen deposits on human retina: Potential in the diagnosis of age-related macular degeneration

K. Rapantzikos, M. Zervakis*, K. Balas

Department of Electronic Computer Engineering, Digital Image and Signal Processing Laboratory, Technical University of Crete, GR-73100 Crete, Greece

Received 27 July 2001; received in revised form 1 January 2002; accepted 21 May 2002

Abstract

Assessment of the risk for the development of age-related macular degeneration requires reliable detection and quantitative mapping of retinal abnormalities that are considered as precursors of the disease. Typical signs for the latter are the so-called drusen that appear as abnormal white-yellow deposits on the retina. Segmentation of these features using conventional image analysis methods is quite complicated mainly due to the non-uniform illumination and the variability of the pigmentation of the background tissue. This paper presents a novel segmentation algorithm for the automatic detection and mapping of drusen in retina images acquired with the aid of a digital Fundus camera. We employ a modified adaptive histogram equalization, namely the multilevel histogram equalization (MLE) scheme, for enhancing local intensity structures. For the detection of drusen in retina images, we develop a novel segmentation technique, the histogram-based adaptive local thresholding (HALT), which extracts the useful information from an image without being affected by the presence of other structures. We provide experimental results from the application of our technique to real images, where certain abnormalities (drusen) have slightly different characteristics from the background. The performance of the algorithm is established through statistical analysis of the results. This analysis indicates that the proposed drusen detector gives reliable detection accuracy in both position and mass size.

Keywords: Age-related macular degeneration; Drusen; Segmentation algorithm

1. Introduction

Age-related macular degeneration (AMD) is the leading cause of irreversible vision loss among the elderly in developed countries. Many studies have confirmed that the presence of the so-called drusen, identified as gray-yellow deposits that build up in or around the macula of the retina, represents a significant risk factor for the development of visual loss from AMD (International ARM Epidemiological Study Group, 1995; Smiddy and Fine, 1984; Bressler et al., 1990; Holz et al., 1994). Drusen are deposited by-products of rod and cone metabolism located just beneath the retinal pigment epithelial (RPE) cell layer (Friberg, 1992). It is believed that they may signal the presence of an altered pathophysiology of the retinal pigment epithelium and consequently they may be a marker for the degree of diffused RPE dysfunction in patients with AMD (Friberg, 2000). The existing strong indications for the correlation between AMD and drusen development characteristics suggest that the clinical assessment of the latter might have predictive value in determining if and when a patient will suffer visual loss from (AMD). Additionally, it could facilitate the development of efficient, fast and accurate clinical tests for the evaluation of the effectiveness of different treatment modalities.

Routinely, drusen characteristics are evaluated by inspecting the retina with the aid of an optical imaging apparatus known as Fundus camera. In some cases and in

*Corresponding author. Tel.: +30-82-103-7206; fax: +30-82-103-7202.

E-mail address: michalis@dilos.systems.tuc.gr (M. Zervakis).

1361-8415/02/$ – see front matter
PII: S1361-8415(02)00093-2

order to assist the evaluation of features of diagnostic importance, slides or digital images of the retina are submitted to Medical Centers, where specialized professionals assess the drusen characteristics. In other clinical studies this assessment is performed with the aid of comparisons with standard photographs or with templates (Smiddy and Fine, 1984; Bressler et al., 1990; Olk et al., 1999; The Choroidal Neovascularization Prevent Trial Research Group, 1998; Klein et al., 1997). While the use of such procedures provides important data toward the standardization of the diagnostic procedure, their precision is relatively low. Klein et al. (1991) show that the agreement rate between different retina image readers in classifying features such us drusen total area and size reaches only 67%.

Besides the subjectivity and the lack of reproducibility, visual assessment is not efficient in analyzing and classifying complex morphological patterns. Drusen vary in size from a few microns in diameter to large confluent complexes, which may extend to hundreds or even thousands of microns (International ARM Epidemiological Study Group, 1995). Moreover, their color appearance varies notably even within the same eye, depending on the amount of the deposited byproducts beneath the RPE in each spatial location. Their color appearance is also affected by the color of the overlaid RPE, which also varies as a function of the location within the same eye, while it is strongly affected by several factors such as blood vasculature, race, etc. The appearance of the retinal features is also affected by the non-uniform transfer function of the illumination-imaging optics of the Fundus camera. These variables affect randomly the perceived contrast between drusen and background, which makes the attempt for the automatic drusen extraction a demanding image analysis task.

The problem of automated, unsupervised drusen detection has received considerable attention over the last decade by various research groups (Kirkpatrick et al., 1995; Morgan et al., 1994; Peli and Lahav, 1986). However, acceptable performance has not yet been achieved mainly due to the inability to compensate satisfactorily for factors that result in poor contrast. A recent interesting work by Shin et al. (1999) proposes an automated, supervised Fundus image analysis technique. Shin et al. are facing the problems of retinal images (non-uniform illumination, poor contrast) in two algorithmic steps, namely preprocessing and segmentation. Their segmentation scheme analyzes each pixel of the preprocessed image as part of a square area varying in size from 20 to 100 pixels. A skewness value greater than -0.5 signifies the presence of drusen. The main drawback of this technique is the requirement of close supervision by experts to achieve adequate accuracy and robustness.

Motivated by the work of Shin et al. (1999), we attempt to expand histogram-based operators and improve the accuracy of drusen detection moving towards an unsupervised tool for detection and mapping of AMD symptoms. We examine enhancement techniques and propose a robust multilevel scheme that can effectively operate without supervision. In the segmentation step, we thoroughly analyze the local histograms' shape by employing more descriptors than the skewness alone, so as to derive robust and accurate thresholding results through the so-called HALT operator. Thus, the purpose of this paper is twofold. First, to develop novel enhancement and segmentation methods that overcome the inefficiencies of other techniques in detecting vaguely defined structures. Second, to propose, analyze and test a complete system for the detection of drusen. Due to the obvious difficulties in detecting drusen, the main focus of the paper is on the HALT operator and its efficiency in recovering drusen hidden in the background. The paper proceeds as follows. Section 2 reviews conventional adaptive contrast enhancement and segmentation algorithms and establishes a novel scheme for image enhancement. Section 3 introduces the HALT operator for drusen detection as the main contribution of this paper and considers step by step the application of the proposed AMD detection algorithm. Section 4 presents the experimental results on four representative images of macular degeneration and the paper concludes in Section 5.

2. Processing tools and methods

2.1. Image enhancement: MLE operator

Drusen are roughly distinguished visually from their background by means of their brightness, morphology and yellowish color. However, the color by itself does not convey consistent information for discrimination. Thus, in order to evaluate the contribution of color to the characterization of the symptoms, an experimental analysis was initially performed considering the RGB (red–green–blue), HSI (hue–saturation–intensity), CMYK (cyan–magenta–yellow–black), G/R and R/B (green/red, red/blue bands) and CIElab color spaces. Several references for color processing can be found that use different color bands for enhancement purposes. We studied drusen visibility in various color spaces and concluded that the gain in visual improvement is less than or almost the same as that of the green band of the RGB space. More specifically, the red band provides information for reflectance in the image and therefore is strongly affected by the non-uniform illumination, whereas the blue band contains almost no useful information for drusen. The green band is more informative and less affected from the overall variation of illumination. Our empirical observations also agree with the selection of the green channel in (Shin et al., 1999), as the channel with the maximum contrast. Another issue related to illumination concerns the normalization of surfaces to light exposure and reflection. When irregular surfaces are

illuminated, the amount of light reflected back to the camera from each region is a function of its orientation with respect to the source of light and the camera. The shape irregularity of the retina produces variable shading across the field of view when illuminated with a bright source, as in the Fundus camera (Shin et al., 1999). For illumination compensation we employ a simple technique, such as homomorphic filtering (Russ, 1999). We prefer such a simple technique to a more complicated finite element analysis of light distribution, due to its computational efficiency and the fact that the subsequent processing adapts to and tolerates gradual intensity variations not completely eliminated by homomorphic filtering.

Following the illumination compensation, the next processing step aims at enhancing the contrast of the retina's image. Towards this direction, histogram equalization and its adaptive versions are quite promising, since they can spread out the modes of a histogram. We use histogram equalization as the core of our enhancement method. Global histogram techniques, like contrast stretching and histogram equalization are widely used to achieve contrast enhancement. Although they are simple to implement, global schemes are affected by the overall distribution in the image and they only stretch illumination differences that are widely spread within the image. Actually, they only separate strong concentrations in the histogram distribution of the image as demonstrated in Fig. 1(b). Such techniques are more effective in our case when applied in small windows as local transforms after the non-uniform illumination compensation. We utilize such enhancement operators in a hierarchical form as to stretch local histogram distributions and enhance the contrast of the image taking under consideration both global and local variations.

In order to standardize the enhancement of retina images and overcome the need for selecting different parameters for each image considered, we develop the multilevel histogram equalization (MLE) technique based on sequential application of histogram equalization. In fact, MLE is a multilevel (hierarchical) scheme that progresses from the entire image to smaller regions defined via windows. 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. A potential problem could arise using windows that are small enough to fit inside a drusen's region. Similar to adaptive histogram modification algorithms, it can produce non-desirable misleading contrast variations within a drusen, as shown in Fig. 1(c). This problem is only experienced when using small windows and forms the opposite drawback (over-enhancement) from that of global techniques. To avoid such effects, we limit the size of windows considered up to the expected size of any drusen in the image.

Considering all constraints, the MLE enhancement algorithm proceeds as follows. The first stage of equalization uses a window equal to the size of the image (global). The second stage splits the image into nine non-overlapping windows and applies the same operation to each sub-block of the previous result. At any stage i, a window w^i is segmented and labeled to estimate the mean size of the drusen involved. This window is further processed by nine smaller non-overlapping windows if and only if it involves smaller drusen. More specifically, the algorithm proceeds to the $i+1$ stage for a specific window w^i if the size of the largest label in w^i is smaller than $1/9$th the size of w^i.

An example of the MLE operation is presented in Fig. 2. The first 'pass' is responsible for enhancing the brightest parts of the image, including small, bright drusen and central parts of larger drusen (Fig. 2(c)). 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(d), contributes in generating more distance between those 'hidden' anomalies and their surrounding areas. In our application we always proceed to the second stage of equalization. Nevertheless, due to the relatively large drusen experienced in all images tested, further window splitting and enhancement is not necessary. Fig. 2(e) and (f) demonstrate the additional enhancement achieved by the 2nd stage of MLE.

2.2. Threshold-based drusen detection: histogram properties

Segmenting the drusen in the enhanced image is an intriguing task. Parts of the drusen are difficult to distinguish from the background because of brightness similarities; especially when encountering drusen near to vessels. In order to efficiently address the problem of region segmentation, two general approaches are widely

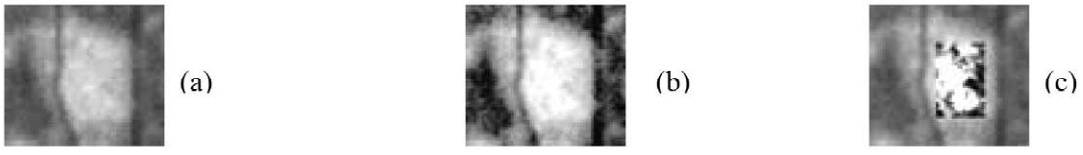

Fig. 1. (a) Original window containing one relative large drusen; (b) histogram equalization using the entire window; (c) histogram equalization using a smaller window inside the drusen's area.

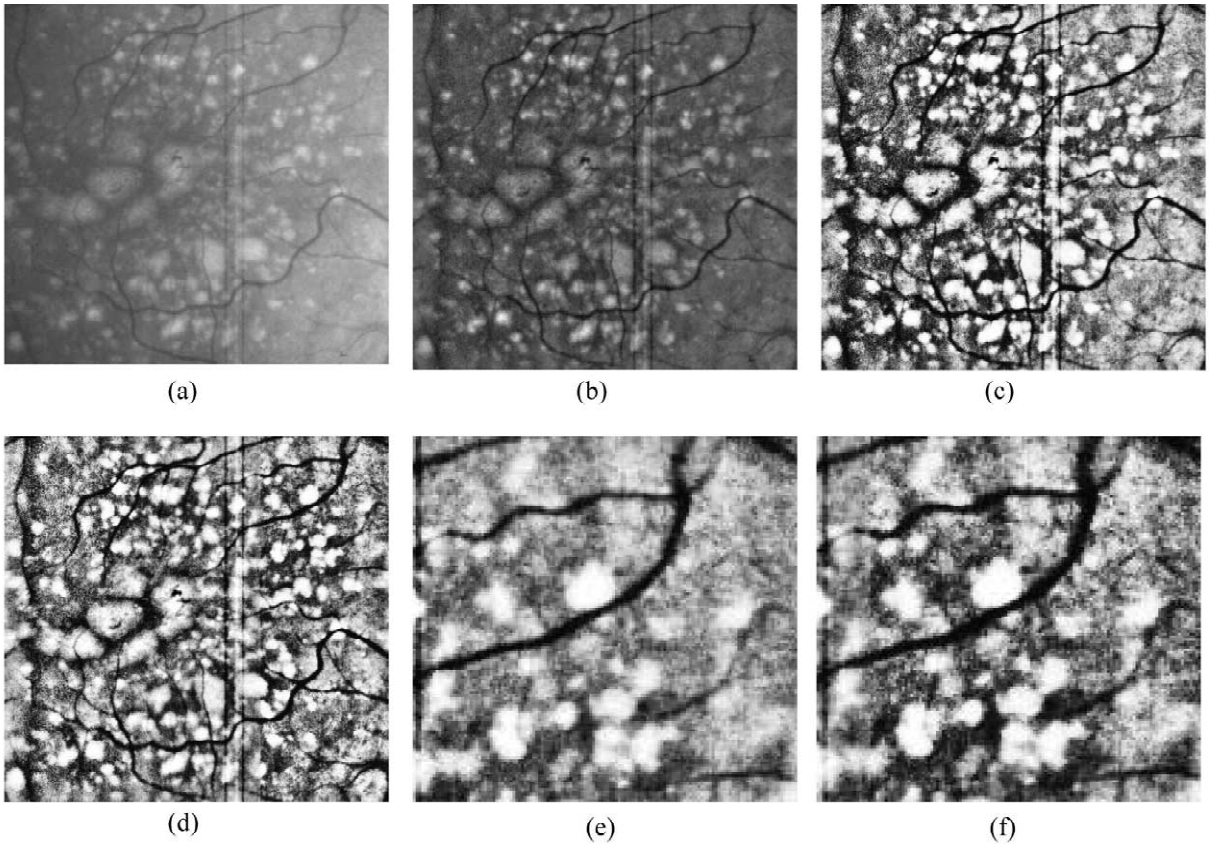

Fig. 2. (a) Original image; (b) image after non-uniform illumination compensation; (c) 1st level of histogram equalization (global) applied to entire image; (d) 2nd level of histogram equalization applied to regions of previous result; (e) enlarged section of upper left corner in (c); (f) enlarged section of upper left corner in (d).

used, namely the stochastic classification (Pappas, 1992) of pixels into object classes and the histogram thresholding for clustering similar intensity pixels into compact objects. In this work we adopt the second approach, i.e. histogram-based thresholding, and particularly focus on the analysis of local histogram that reveals small included objects. This analysis leads to the definition of the HALT operator, presented in detail in the next section. The HALT operator as a region-based threshold selection scheme fits well with the region-based MLE approach developed for enhancement. The observed value of a random variable can provide significant information with respect to the stochastic distribution of the variable. In fact, under the ergodicity assumption, statistical measures of the distribution can be accurately inferred from the sample measures as the size of the observed data increases (Papoulis, 1984). The statistics that can be easily computed from the observations (data set) of a random variable fall into several categories, with the most important being the following:

(a) Measures of central tendency—statistics that describe the center of the data, including the mean, the median and the mode.
(b) Measures of spread—statistics that describe the dispersion of the data, including the variance, standard deviation, range, and inter-quartile range.
(c) Measures of shape—statistics that compare the shape of the data to that of a normal distribution, including the skewness and kurtosis.

The mean, median and mode are normally close to each other. These three statistics measure the center of the data in somewhat different ways. The mean is the value that minimizes the mean square distance from the data set. The median is the value that minimizes the absolute distance and the mode is the one that minimizes the H_∞ (norm) distance from the data set. The case that these statistics are very close to each other indicates that the data is probably symmetric. As the distribution of the data becomes skewed, the sample mean moves away from the sample median in the direction of the longer tail. Generally, the sample mean is affected by extreme values (outliers), while the median tends to stay with the main body of the data (Huber, 1983; Hampel et al., 1986; Pitas and Venetsanopoulos, 1990). Thus, the median is often located in between the mode and the mean of a distribution.

A measure of the center of a distribution is much more useful if there is a corresponding measure of *dispersion* that indicates how the distribution is spread out with respect to the center (Otsu, 1979; Sahoo et al., 1988; Boukharouba et al., 1985). For the median, a natural measure of dispersion can be obtained from the lower and

upper *quartiles*. The lower quartile is the value that integrates $1/4$ of the data set and the upper quartile is the value at $3/4$ of the data set. For the mean value, these limits are usually obtained at 10 and 90% of the cumulative distribution.

In this paper we propose to distinguish between background and drusen regions by means of automatic segmentation performed through thresholding of the local histogram. Our goal is to separate the drusen, without being affected by intensity variations caused by vessels, noise and uncompensated non-uniform illumination. If we zoom into each local intensity area, we observe different shapes of the histogram for each of these regions and different relative distributions of the drusen and the background. Thus, in order to determine an efficient threshold for each neighborhood, we study local histogram in terms of its central tendency, symmetry and shape tendency. More specifically, we consider symmetry of a distribution via two quotients. The first quotient, namely the $|mean - median|$ difference, is a first measure of symmetry based on local statistics, as indicated before. The second quotient, namely the $|mode - mean|$ difference, is chosen as a measure of histogram's main lobe spread. If both of them are small, smaller than $1/3\,\sigma_b$, then the distribution is considered symmetric. Otherwise, the distribution is labeled asymmetric.

Subsequently, the skewness in conjunction with the kurtosis are used as measures of the histogram's tendency. Let b and $P(b)$ denote the variable (intensity) and its distribution (histogram), with σ_b and $\bar{b}$ representing its standard deviation and mean, respectively. Skewness is defined as $S_S = 1/\sigma_b^3 \sum_{b=0}^{L-1}(b-\bar{b})^3 P(b)$ and kurtosis as $S_K = 1/\sigma_b^4 \left[\sum_{b=0}^{L-1}(b-\bar{b})^4 P(b)\right] - 3$. A distribution is skewed if one of its tails is longer than the other. Positive skew indicates a long tail in the positive direction, while negative skew indicates a long tail in the negative direction. Kurtosis is based on the size of a distribution's tails. Distributions with relatively small tails (sharp-peaked, $S_K > 0$) are called 'leptokurtic'; those with large tails (flat-topped, widely spread, $S_K < 0$) are called 'platykurtic'. A distribution with the same kurtosis as the normal distribution ($S_K = 0$) is called 'mesokurtic'. These measures can increase the confidence with which drusen (outliers on the assumed normal distribution of the background) are detected and they are used in the definition of the HALT operator, which is analyzed and tested within the context of the next section.

3. Methodology for drusen detection: HALT operator

This section introduces a complete algorithm for automatic segmentation and detection of drusen, based on the previous analysis. The algorithmic steps are shown in Fig. 3 and explained in the following. The homomorphic filter is applied at the front end of the algorithm to compensate for illumination irregularities. The second step involves the enhancement operation that is responsible for stretching intensity differences characterizing drusen and background. The proposed MLE approach succeeds in enhancing most drusen, being insensitive to small brightness variations that are caused, for example, from remaining non-uniform illumination and noise. Small and bright drusen are extracted successfully, whereas large and spread drusen that tend to be darker near the edges are also identified.

As a result of the previous operators, sharp abnormalities in intensity (candidate drusen) are strongly enhanced and differentiated from the background. Such intense drusen can be readily detected by thresholding techniques. 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 global thresholding technique (Otsu, 1979; Sahoo et al., 1988). A single threshold, however, cannot identify small intensity differences that often discriminate vague abnormalities hidden in bright background areas. Thus, we propose to use a two-stage histogram thresholding approach. The first stage applies the global Otsu threshold to provide an initial segmentation map. This level of thresholding cannot

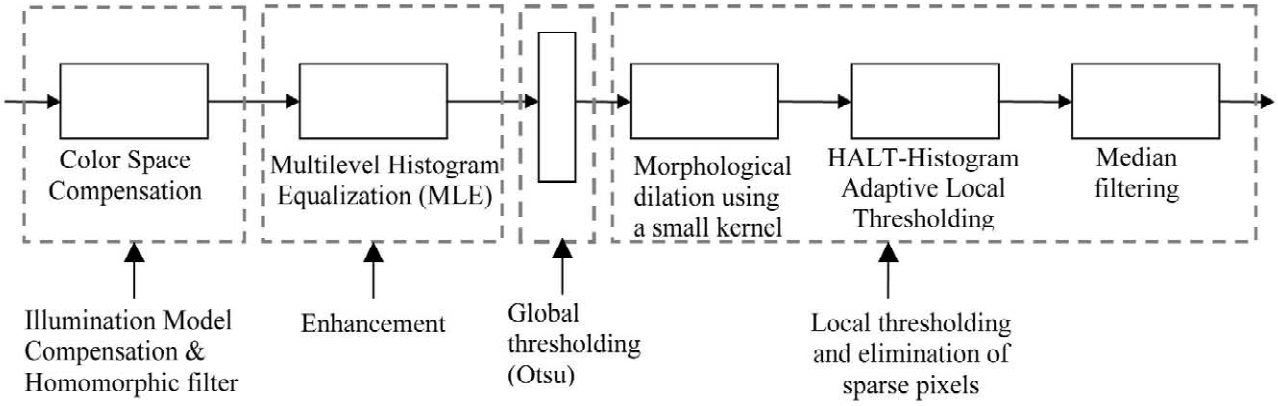

Fig. 3. Proposed algorithm for the detection of anomalies in the retina of the human eye.

discriminate vague abnormalities hidden in the local regions of background. It only detects and preserves regions of evident abnormalities that are crisply separated from their background, as well as background regions mixed with vague abnormalities. The second stage of thresholding that refines the segmentation map operates on a local level defining a different threshold for each local region of interest. For this particular stage, a novel local thresholding operator (HALT) is designed and analyzed in this section. The HALT operator checks the local histogram for general symmetry or asymmetry and uses shape tendency indicators for assessing regions as drusen or actual background.

A morphological dilation operator precedes the HALT operator. The morphological dilation (disk shaped structuring element, 3 pixels in diameter) expands the regions that are not removed by global thresholding. If this expansion occurs in background areas, there is no serious effect, since the following segmentation step is capable of removing these expanded regions completely. The main advantage of dilation is appreciated in areas that contain only one or two large drusen without background, where direct application of any local threshold would completely eliminate the drusen area. The morphological expansion reconstructs some of the background and recovers different intensity areas in the local histogram. In other words, it forces better distinction between bright areas and their darker surroundings at the corresponding local histogram. To achieve this operation, the dilation operator is embedded into the global thresholding operator, such that the overall global thresholding mask is obtained by dilating the initial threshold mask of the Otsu operator. Thus, the dilation is applied on the binary map of thresholded drusen areas to expand their extent.

The HALT operator applies different thresholds to regions of the image, depending on the properties of the corresponding histogram. As in the case of the MLE operator, the image is split into nine non-overlapping windows, where the HALT operator is applied. If needed, each window is further split into nine sub-windows, in order to refine the segmentation. Within each window, the HALT operator checks the statistics of local histogram and assigns the appropriate threshold. The background is composed of a noise process superimposed on a deterministic smoothly varying terrain. A symmetric Gaussian distribution efficiently 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. Thus, in case of a window in pure background regions, it is expected that by thresholding the distribution at its cumulative 90% level and preserving only values above this 90% threshold, we preserve only isolated pixels randomly distributed along the spatial extent of the image. These pixels are easily removed by median filtering. A symmetric distribution, however, may not always characterize background alone but can also characterize certain combinations of drusen and background distributions. This case requires thorough examination. By similar means, a positively skewed distribution indicates the presence of drusen influencing the higher part of the intensity distribution. Otsu's threshold is most appropriate in this case; if there is strong evidence that drusen is the cause of this positive influence. Negatively skewed distributions are most likely to describe areas of background, since drusen abnormalities affect the higher end of the histogram (bias towards bright values). So, by setting 90% as a threshold would also remove such regions.

Organizing these potential distributions, the HALT operator first classifies the local histogram into two distinct cases, depending on its symmetry properties, as described by the two symmetry quotients in Section 2.2. Subsequently, the appropriate threshold is determined according to the measures of spread and shape, as follows.

3.1. Histogram totally or almost symmetric (Table 1)

- A totally symmetric gray level distribution signifies areas that are mainly occupied by background regions. However, small drusen may be present, so setting the point of 90% of the cumulative distribution as threshold would be adequate to remove background and preserve compact 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's and symptoms' gray levels are normally and equally distributed, the total histogram still appears symmetric. In this case, to avoid removal of drusen, we employ Otsu thresholding.
- In case of sharp-peaked (leptokurtic) almost symmetric histograms we observe high concentration of pixels around the mean value. These regions appear with almost uniform background. Such leptokurtic background distributions may only allow the existence of small drusen as outliers that do not alter the general uniformity of intensities. Using Otsu thresholding, that is obtaining a threshold value close to mean, would retain anomalies and a large part of the background. Alternatively, setting 90% as threshold would remove background areas and retain, if existing, small compact drusen areas.
- The case of a mesokurtic and positively skewed histogram 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 as significant as to alter drastically the background statistics. This second distribution is detected at high intensity values indicating the existence of object(s), whose

Table 1
HALT operator in symmetric distributions

Kurtosis	Skewness		
	<0	≈ 0	>0
Platykurtic	Mainly background ⇓ 90%	Possible combination of two or more distributions ⇓ Otsu	⇓ Otsu
Mesokurtic	Mainly background ⇓ 90%	Mainly background and maybe some drusen or just large drusen (one distribution) ⇓ 90%	Drusen and background are hard to distinguish ⇓ Application of HALT in smaller regions
Leptokurtic	Mainly background ⇓ 90%	Almost constant background ⇓ 90%	Can signify the case of only a small portion of drusen: segment small drusen of high intensity ⇓ 90%

intensity however interacts with that of the background. Thus, the direct segmentation of object and background 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. Thus, an additional step of local thresholding is used, which is actually the application of HALT method focused on smaller areas of the first level's region. This helps in obtaining better distinction of anomalies and background at corresponding second level histograms.

3.2. Histogram totally or almost asymmetric (Table 2)

- A positively skewed distribution of this class notifies the presence of many small or a few large drusen. In fact, bright gray levels that generally characterize anomalies dominate the histogram. Otsu thresholding is best suited to this case, since the distinction of bright and darker areas (background) is obvious.
- An asymmetric non-skewed distribution signifies the presence of drusen. This distribution results as a combination of similar distributions, characterizing background and abnormalities (drusen). Thus, Otsu thresholding is appropriate for segmenting the drusen in such regions.

The exact process for selecting the threshold in the HALT operator is outlined in Tables 1 and 2. The HALT operator is succeeded by a median filter that eliminates sparse pixels that cause false 'alarms' for presence of anomalies. In this way, the HALT preserves as drusen only those pixels that appear compactly distributed into regions. The median filter is necessary to remove sparse pixels preserved by the application of the 90% threshold in background regions.

Table 2
HALT operator in asymmetric distributions

Kurtosis	Skewness		
	<0	≈ 0	>0
Platykurtic	Mainly background ⇓ 90%	Drusen are present ⇓ Otsu	Drusen and background are almost equally distributed ⇓ Otsu
Mesokurtic	Mainly background ⇓ 90%	Drusen and background are almost equally distributed ⇓ Otsu	Mainly drusen ⇓ Otsu
Leptokurtic	Mainly background ⇓ 90%	Mostly background, less drusen ⇓ Otsu	Drusen and background ⇓ Otsu

The block-wise application of the HALT operator may produce undesirable segmentation results when a region's histogram appears to be almost symmetric (small skew and mesokurtic). Applying a 90% threshold on this region's histogram, followed by a median filter, may preserve isolated bright groups of pixels. If these small bright areas are very close to each other, then they possibly belong to the same larger drusen and must be expanded so as to capture the entire drusen region. A morphological closing with a small structuring element (disk shaped, 3 pixels in diameter) applied locally within such regions can join together the neighboring groups of pixels into a single drusen. It is emphasized here that this selective expansion process is only applied on small sub-blocks of the image that possess symmetric small skewed and mesokurtic histogram, as have been identified by the HALT operator. The proposed algorithm for drusen detection and segmentation is summarized in detail in Fig. 3.

In order to demonstrate the efficiency of the HALT approach over the localized Otsu and the Shin et al. methods for threshold selection, two representative examples are shown. One with large drusen dominating extensive areas (Fig. 4(a)) and one with few small and vaguely defined drusen (Fig. 4(b)). Both images are enhanced using multilevel histogram equalization (MLE) and then thresholded using local Otsu, Shin et al. and HALT techniques. A median filter is applied afterwards to remove isolated pixels. The results are presented in Fig. 5.

Otsu's localized thresholding scheme works well in regions dominated by drusen (brighter areas), since the distinction between them and the background is evident. This is demonstrated in Fig. 5(a), where drusen at the central part of the image are correctly distinguished from the surrounding areas. However, the algorithm is strongly affected by regions that do not contain abnormalities, like those regions at the sides of the image. Due to remaining effects of non-uniform illumination, parts of these regions are brighter and are misclassified as anomalies. Fig. 5(b) brings out another disadvantage of the local Otsu scheme. Vaguely defined drusen, which are either small or located inside bright background regions, are not segmented. The algorithm detects the most obvious drusen (two of them are easily conceived), but fails to isolate and detect 'hidden' anomalies; arrows indicate some of those.

The segmentation of Shin et al. (Fig. 5(c, d)) tends to spread and overestimate the area of drusen especially around vessels. Although the most obvious drusen of the first image are detected by this method, supervision is required in order to remove many incorrectly segmented areas. Using the same parameters in the second test image, this method produces the result of Fig. 5(d) expressing an inability to accurately isolate small and vaguely defined drusen. It is emphasized here that the Shin et al. technique can give improved results with a proper selection of its parameters for each image. This need for parameter selection specifically for each test image renders the method inappropriate for the automatic and unsupervised segmentation of drusen. On the contrary, the HALT technique removes most of the background in both cases, as shown at Fig. 5(e, f). Even the most hard-to-see drusen are segmented without loosing their actual size and shape. Some sparse false negatives generated by the existence of noise can be easily removed through simple median filtering. Notice that the parameters of our algorithm are set once and remain fixed for the ensemble of images tested.

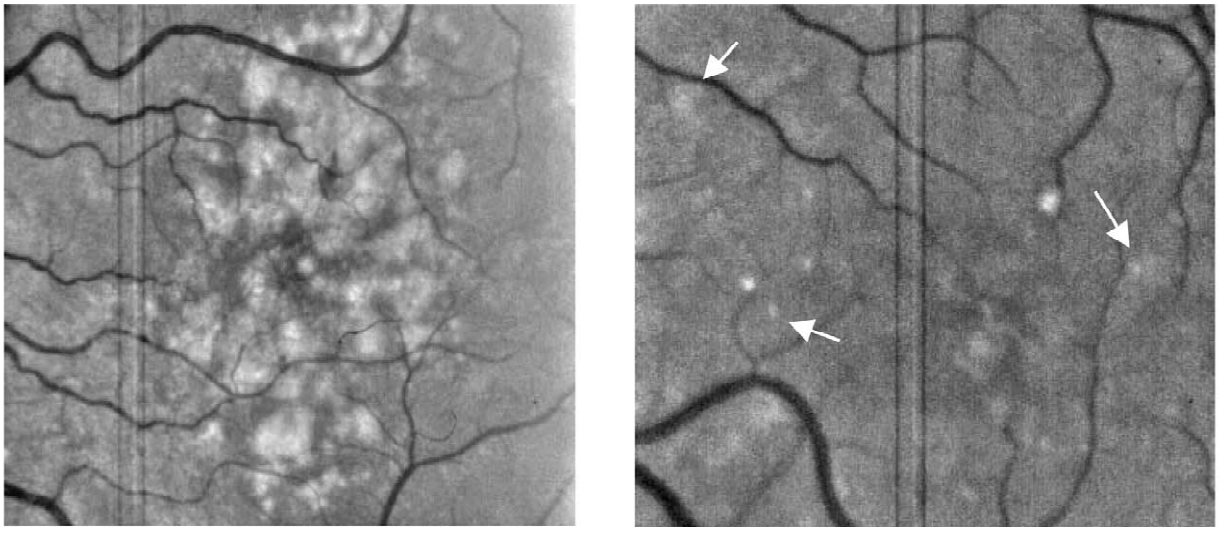

Fig. 4. (a) Image with large dense drusen; (b) image with small sparse drusen.

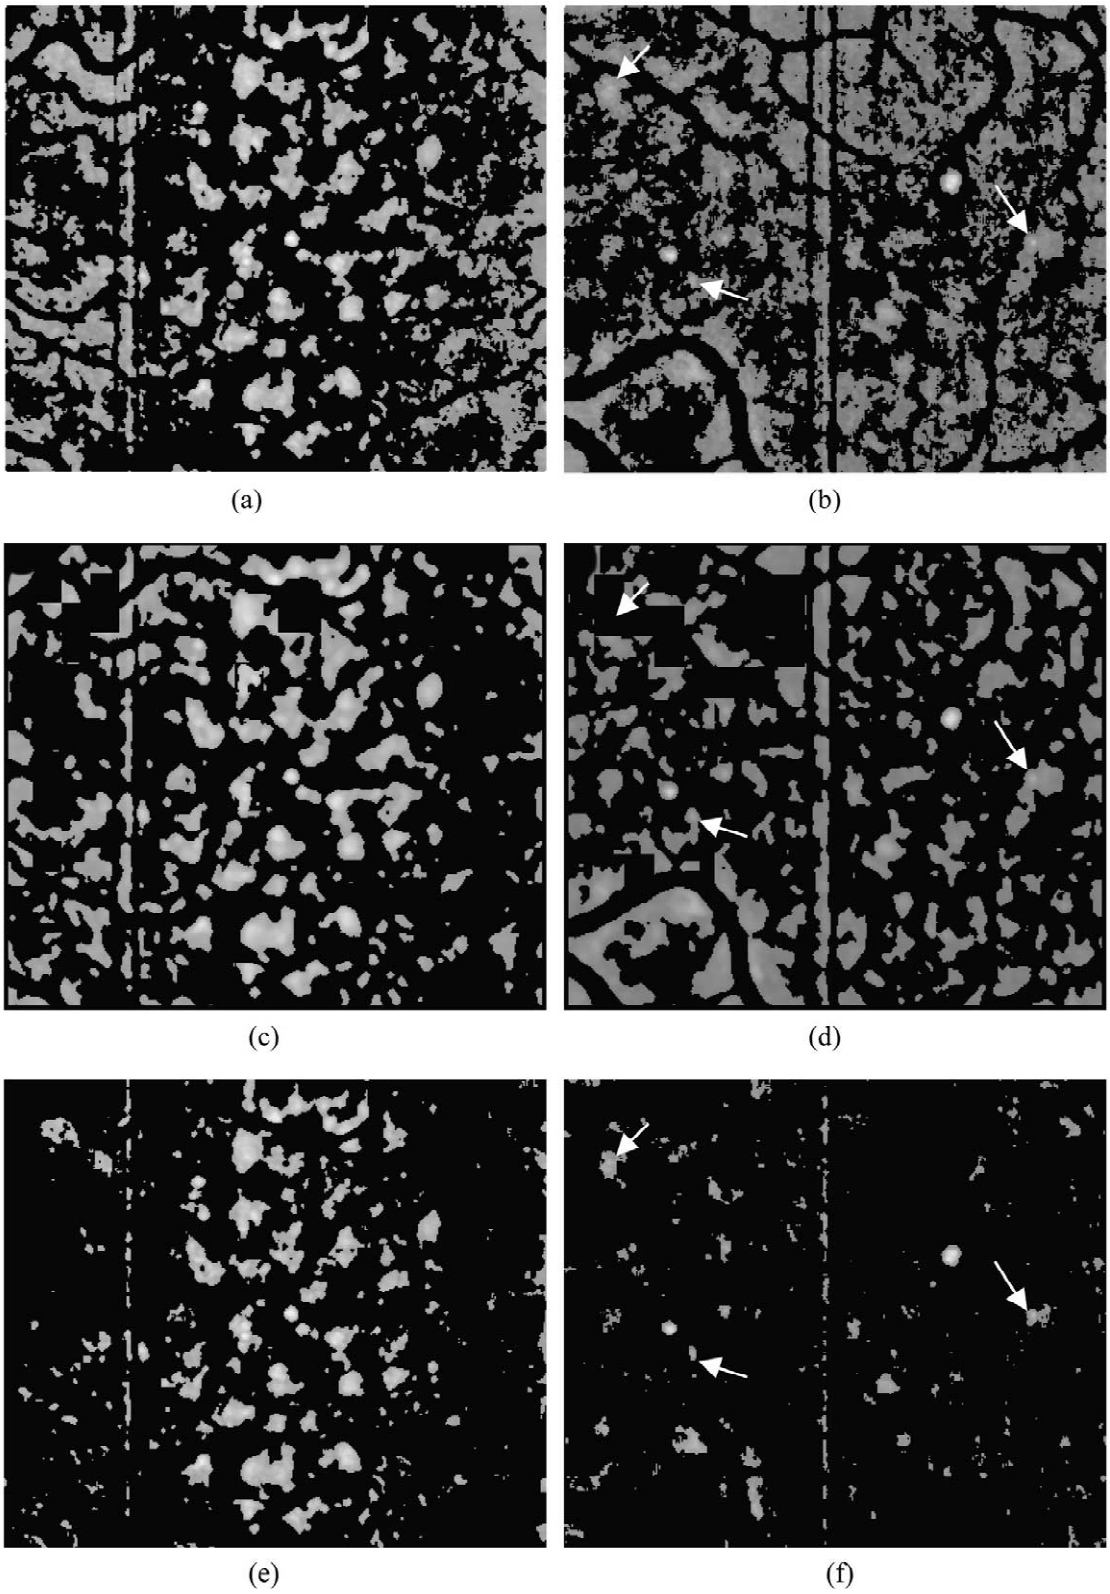

Fig. 5. Segmented images using local Otsu (top), Shin et al. (middle) and HALT (bottom) techniques.

4. Results

We tested our algorithm using a set of 23 images, acquired by the Fundus camera. Eight 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, respectively). Fig. 6 presents examples of gray-scale versions (green band) of the original color images. Drusen show up as bright blobs, but it is evident that the automatic extraction of these pathological features is difficult, since drusen vary strongly in shape and size and they tend to spread (varying brightness) around their location. Additionally, small bright regions of the background tend to create larger areas that can be mistaken as large drusen. The results of the proposed algorithm for the detection of defect regions (drusen) inside the human retina are presented in this section.

Fig. 7 demonstrates the detection potential of the proposed algorithm on three representative images from the available set of retinal images; one image with small and large drusen, a second one with large drusen, and a third one with small sparse drusen. The segmentation results are more than satisfactory since in all cases the background areas are segmented out and the symptoms are almost correctly isolated. A qualitative evaluation from experts is presented at the end of the current section. Fig. 8 presents step-by-step the results and shows the effectiveness and robustness of the proposed operators.

Another example of an image that requires expansion of some regions after the HALT operator is shown at Fig. 9. This image contains large drusen that consist of bright and darker parts. In Fig. 9(b) it is obvious that after the HALT we are left with areas that must be joined together or expanded, so as to recover missing parts of anomalies. As shown in Fig. 9(c), after the proposed morphological closing the upper areas that appeared ‘cracked’ are joined together and form a single region that covers almost entirely the actual anomaly area.

A hard-to-enhance image is shown in Fig. 10. The

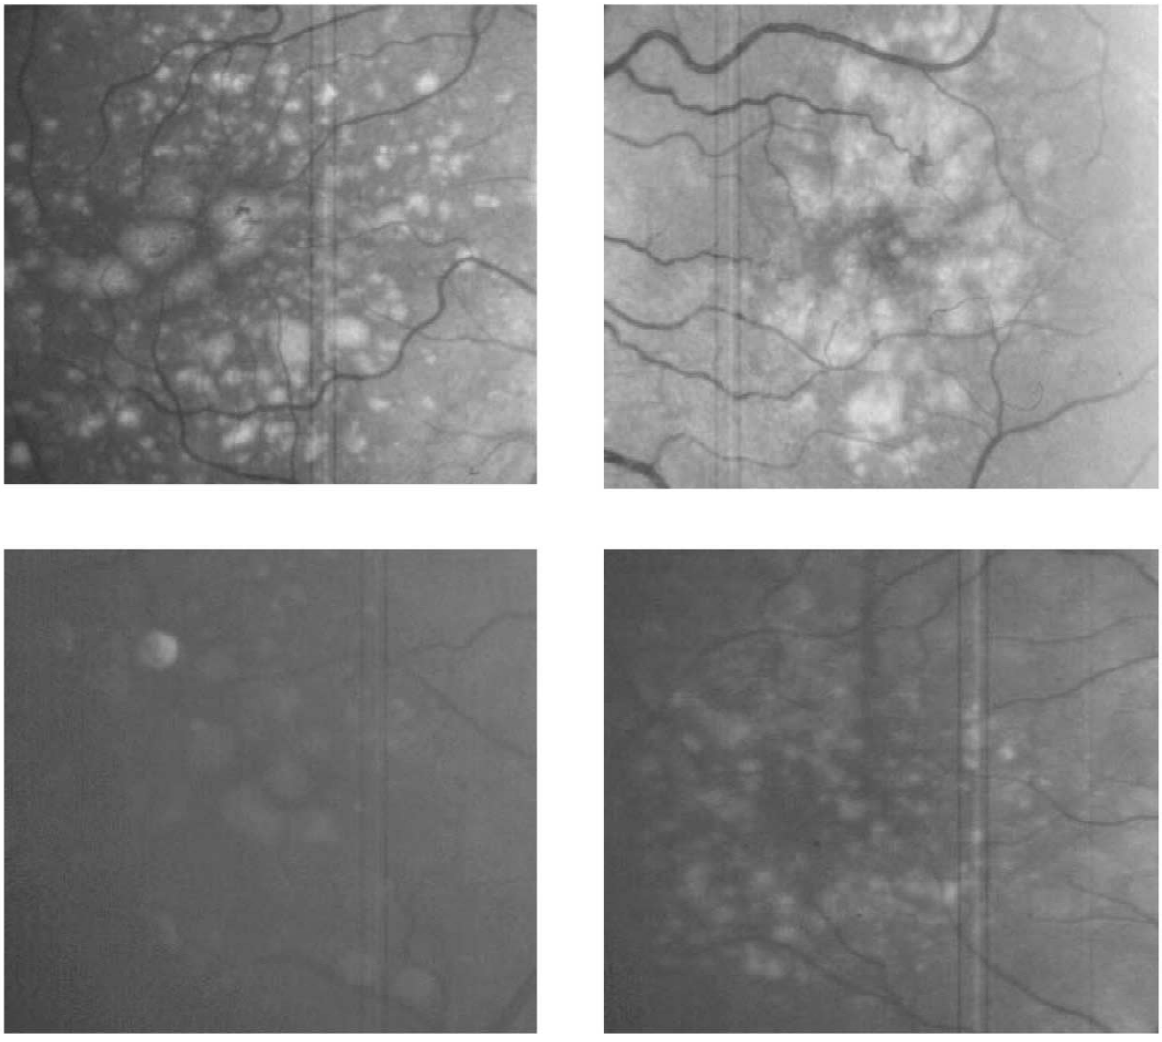

Fig. 6. Examples of test images.

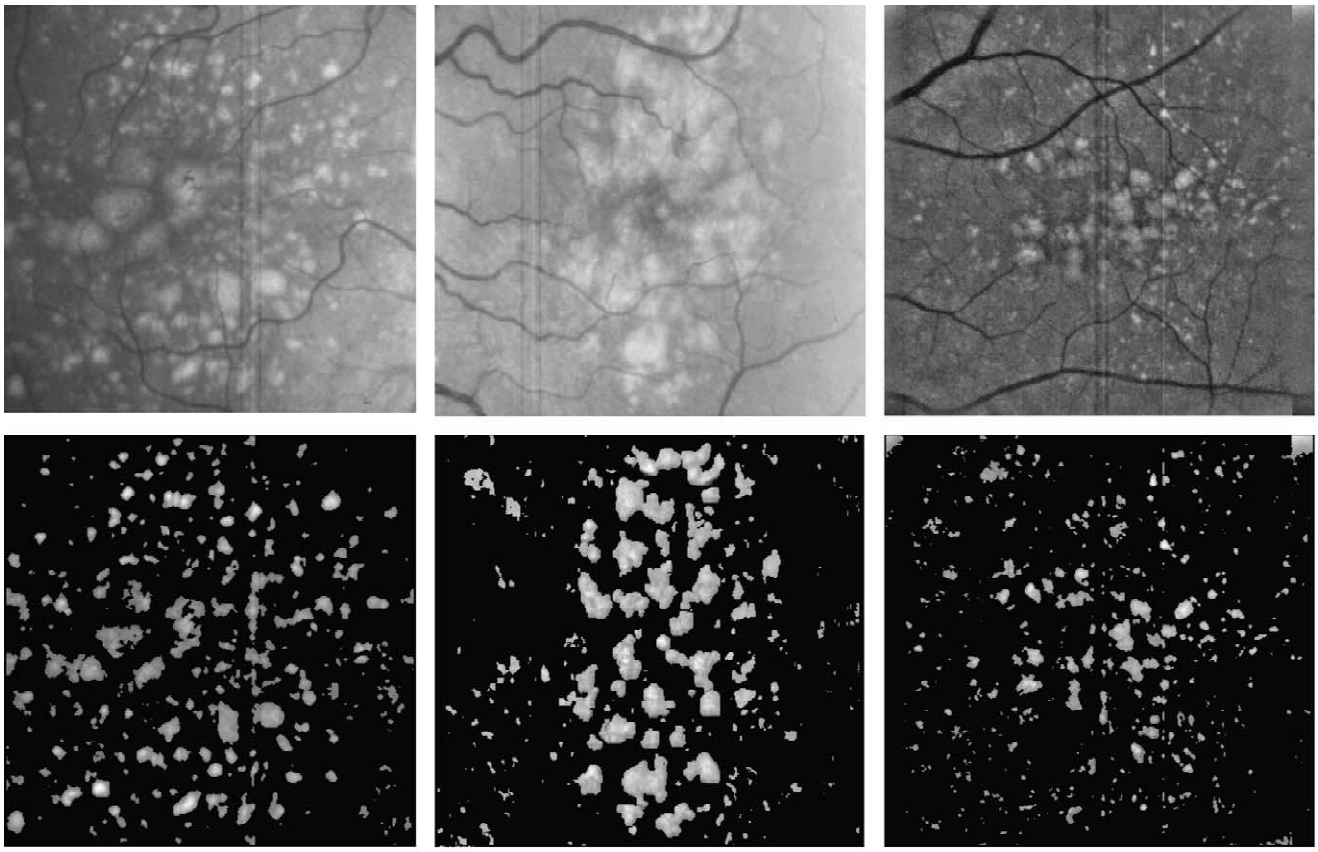

Fig. 7. Results of drusen extraction in retinas of (a)–(c) different persons. The first row shows the original images and the second row presents the segmented ones.

presence of noise is strong, as it is detected at background regions. In addition, large drusen do not differ sufficiently from the background. Except of the circular bright drusen, all others are noisy and intermixed with surrounding areas. Even in this case, our algorithm detects correctly all small drusen and loses only few parts of larger ones, which appear at the central part of Fig. 10(c).

In general, the presence of vessels and their interaction in intensity with drusen pose serious problems even in manual drusen detection. The proposed algorithm overcomes this problem and does not experience false detection, in the entire test set of images. This is due to the appropriate consideration of features in local areas that can separate drusen from vessel distributions. Overall, the proposed algorithm performs quite efficiently in the entire set of macular degeneration images tested. This set of images covers a wide range of possible drusen sizes and formations, including vague, non-canonical shaped and thin blobs.

In order to provide a statistical analysis of the algorithm's performance, we asked for experts' assistance in determining the actual drusen areas. Notice that all images reflect actual test cases without any prior information on the status and extent of AMD. Thus, for testing the algorithm's classification (drusen versus normal background) against an 'actual' state, we are based on clinical evaluations performed by the experts. Two experts have extensively studied the retinal images and all the areas that are considered drusen by the doctors have been manually segmented. Their intersections, i.e. the areas that are classified as drusen by both experts, are considered as 'true' drusen areas. Thus, our statistical tests give priority to 'correct detection' than to 'false alarm'. Statistical measures, such as the rate of true positive detection (sensitivity or TPR), false positive detection (1-specificity or FPR) and false negative detection (FNR) have been computed (Table 3), in order to establish the performance of the algorithm. As mentioned before, we tested our algorithm using a set of 23 images. Eight pairs of them were actually captured from the left and right eye of patients (indicated by a, b in Table 3). The sensitivity and specificity of the algorithm exceed 96% for almost all test cases. Only in one case the sensitivity falls around 88% due to noise (test case corresponds to Fig. 10). The FNR statistic reveals that the algorithm underestimates drusen areas in this case. The overall performance of the proposed

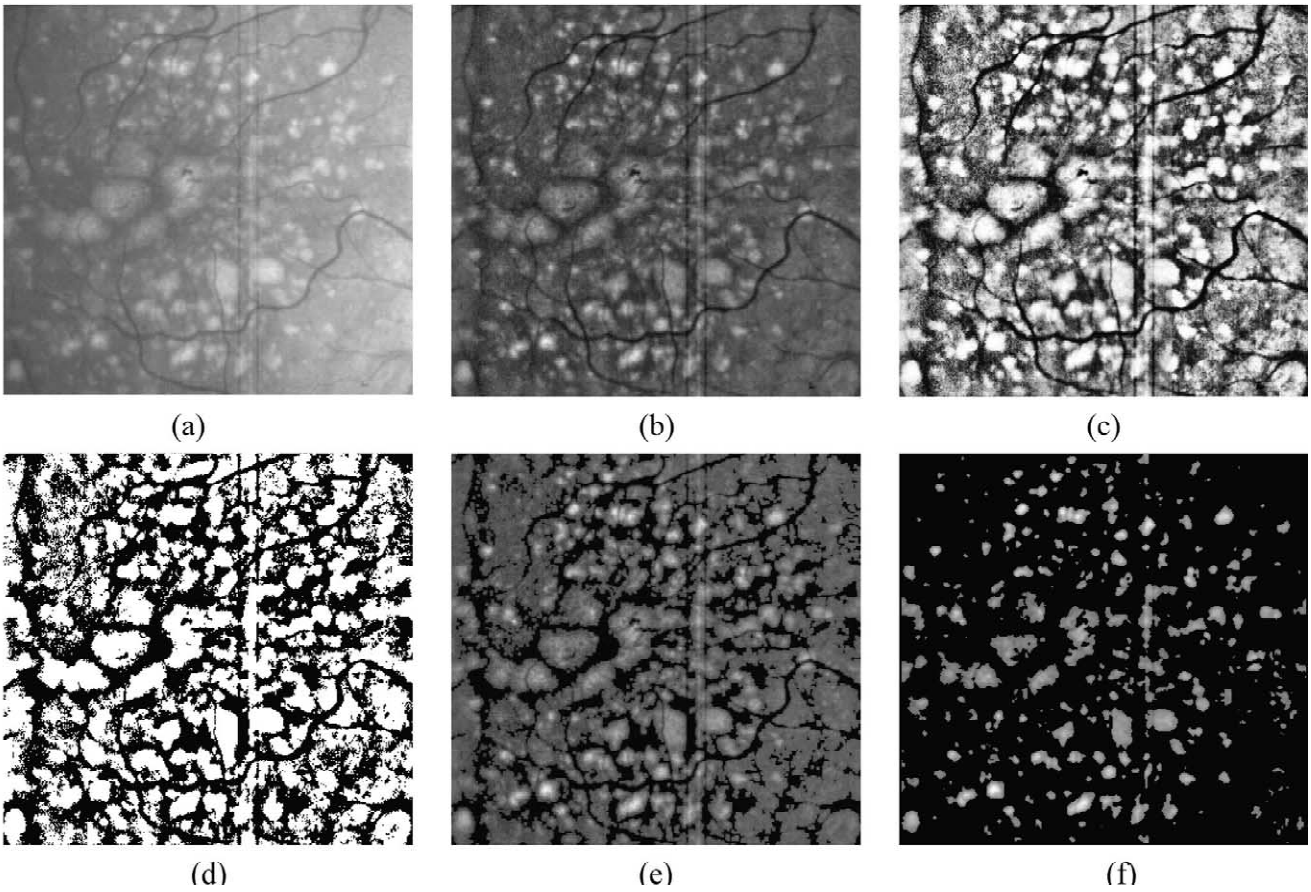

Fig. 8. Results of each step of our proposed algorithm: (a) original image; (b) non-uniform illumination correction; (c) enhancement; (d) global thresholding; (e) morphological dilation; (f) HALT and median filtering.

algorithm on the entire set of images tested is presented in the last row of Table 3.

Further demonstrating the efficiency of the proposed algorithm, the results are subtracted from the original images, so that the detected regions appear black. Parts of the drusen that are not detected should appear bright,

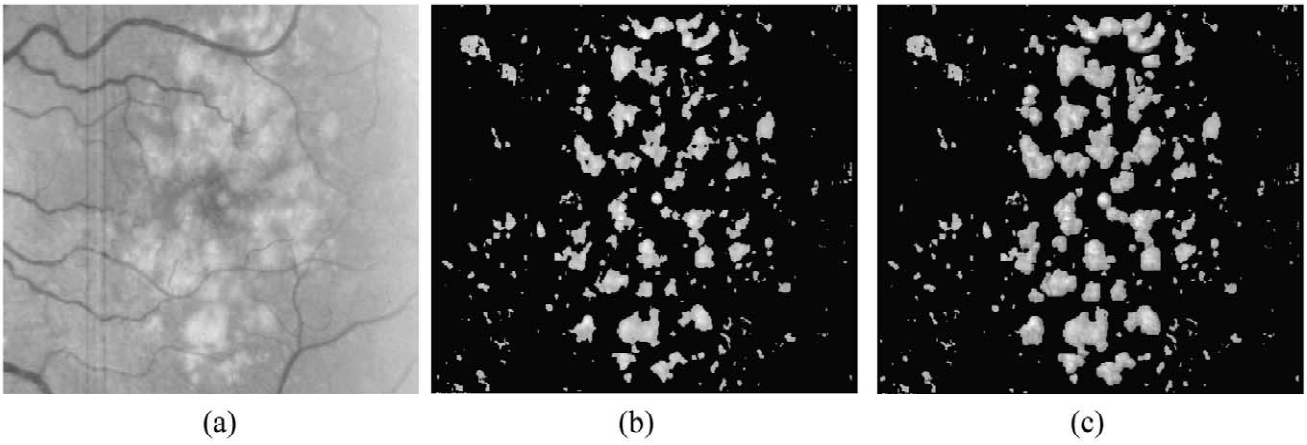

Fig. 9. (a) Original image; (b) HALT and median filtering; (c) expansion of problematic areas.

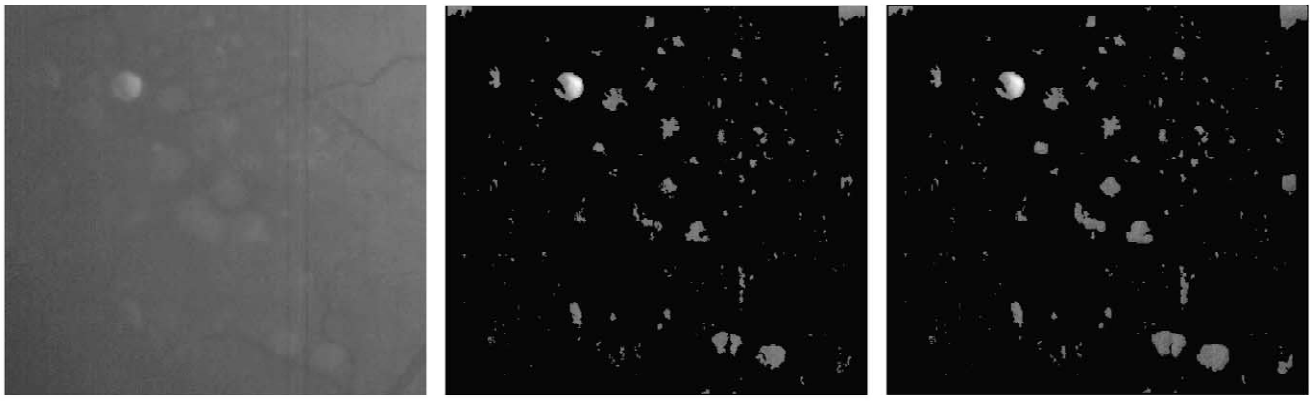

Fig. 10. (a) Original image; (b) HALT and median filtering; (c) expansion of problematic areas.

retaining their original gray level. Fig. 11 illustrates the experts' comments on representative cases. The areas inside the solid lines are underestimated (in size), while those inside dotted lines are overestimated. Overestimation is experienced mainly at the borders, due to the different lighting model from the center to the edges of the image. These false alarm areas can be easily rejected by the doctor inspecting the results and do not pose any problems to the early detection of AMD cases. In general, the drusen of interest in AMD examination are located inside or around the macula, the central part of the eye. In these areas, our proposed methodology does not produce false alarms.

Underestimation of drusen areas is a more severe problem. It should be emphasized that the underestimation of area experienced in some cases does not imply complete miss of the drusen, but only partial segmentation of it. In these cases, our methodology provides a diagnosis aid for indicating drusen presence for further examination by the doctor, who will anyway be responsible for reaching the final diagnosis.

Table 3
Statistical analysis of the results

Image	% Sensitivity (TPR)	% Specificity (1-FPR)	% FNR
1a	96.54	100	3.46
1b	100	100	0
2a	100	98.86	0
2b	100	99.75	0
3a	98.51	99.64	1.49
3b	100	99.77	0
4a	100	100	0
4b	100	99.52	0
5a	88.27	100	11.73
5b	100	100	0
6a	100	97.87	0
6b	100	98.95	0
7a	99.44	97.94	0.56
7b	99.09	99.14	0.91
8a	100	100	0
8b	100	99.78	0
9a	97.86	98.82	2.14
10a	96.15	99.19	3.85
11a	100	98.64	0
12a	97.6	100	2.4
13a	100	100	0
14a	100	96.42	0
15a	100	100	0
Overall	98.846	99.317	1.154

5. Conclusions

This paper considers histogram-based techniques for the problem of automatic AMD evaluation. The detection of anomalies in human eye's retina is a biomedical problem appropriate for image processing and automated segmentation, whose solution is intended to help the ophthalmologists 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 introduce and test a histogram-based enhancement technique (MLE), which uses histogram equalization as its core operator and a histogram-based segmentation technique (HALT) to segment areas that differ slightly from their background regions. Furthermore, we establish an unsupervised and non-parametric method for drusen extraction and consider its effectiveness through several examples. The proposed method is able to detect actual drusen in various cases tested. Even in hard-to-diagnose cases, where many small and vague drusen exist, our method succeeds in isolating them from the background. The proposed algorithm extends the work of Shin et al. (1999) towards the development of robust, unsupervised detection and reliable quantitative mapping of drusen abnormalities.

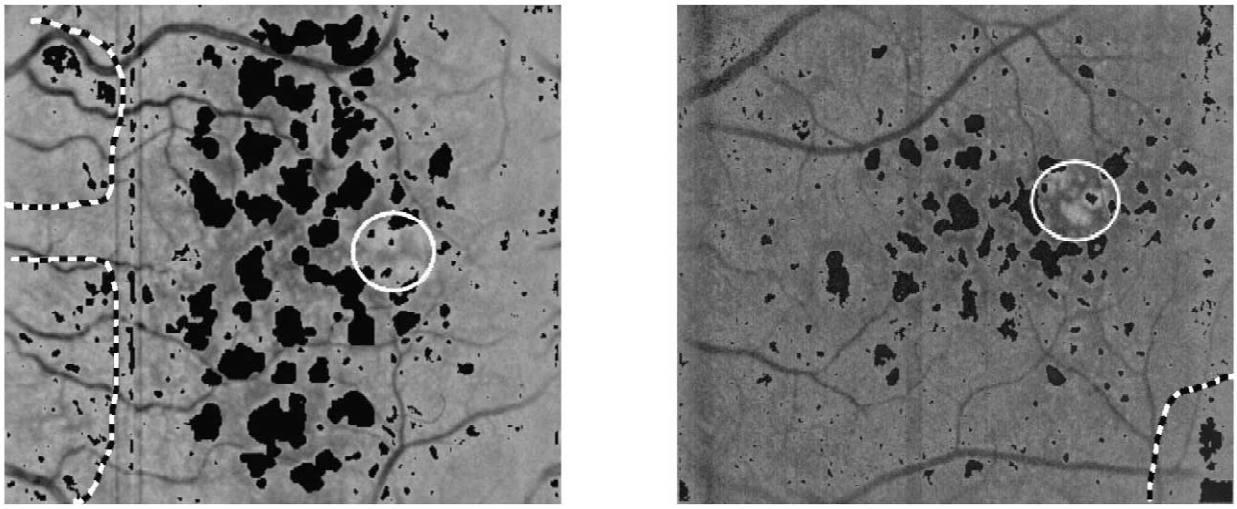

Fig. 11. Evaluation of problems judged by experts. Solid lines represent area underestimation and dotted lines represent overestimation.

References

Boukharouba, S., Rebordao, J.M., Wendel, P.L., 1985. An amplitude segmentation method based on the distribution function of an image. Comput. Vis. Graphics Image Process. 29, 47–59.

Bressler, S.B., Maguire, M.G., Bressler, N.M., Fine, S.L., 1990. The macular photocoagulation study group, relationship of drusen and abnormalities of the retinal pigment epithelium to the prognosis of neovascular macular degeneration. Arch Ophthalmol. 108, 1442–1447.

Friberg, T.R., 1992. Macular degeneration and its possible relationship to the microcirculation. In: Weinerb, R.N., Joyner, W.L., Wheeler, L.A. (Eds.), Ocular Microcirculation. Elsevier, Amsterdam, pp. 173–178.

Friberg, T.R., 2000. Discussion, surgical repair of full-thickness idiopathic macular holes associated with significant macular drusen. Ophthalmology 107, 2233–2239.

Hampel, F., Ronchetti, E., Rousseeuw, P., Stahel, W., 1986. Robust Statistics. Wiley, Chichester.

Holz, F.G., Wolfensberger, T.J., Piguet, B., Gross-Jendroska, M., Wells, J.A., Minassian, D.C., Chisholm, I.H., Bird, A.C., 1994. Bilateral macular drusen in age-related macular degeneration. Ophthalmology 101, 1522–1528.

Huber, P.S., 1983. Robust Statistics. Wiley, Chichester.

International ARM Epidemiological Study Group, 1995. A classification and grading system for age-related macular degeneration. Surv. Opthalmol. 39 (5), 368–374.

Kirkpatrick, J.N.P., Spencer, T., Manivannan, A. et al., 1995. Quantitative image analysis of macular drusen from fundus photographs and scanning laser opthalmoscope images. Eye 9, 48–55.

Klein, R., Davis, M.D., Magli, Y.L., Sgal, P., Klein, B.E.K., Hubbard, L., 1991. The Wisconsin age-related maculopathy grading system. Ophthalmology 98, 1128–1134.

Klein, R., Kelin, B.E.K., Jensen, S.C., Meuer, S.M., 1997. The five-year incidence and progression of age-related maculopathy: The beaver dam eye study. Ophthalmology 104, 7–21.

Morgan, W.H., Cooper, R.L., Constable, I.J., Eikelboom, R.H., 1994. Automated extraction and quantification of macular drusen from fundal photographs. Aust. NZ J. Ophthalmol. 22, 7–12.

Olk, R.J., Friberg, T.R., Stickney, K.L., Akduman, L., Wong, K.L., Chen, M.C., Levy, M.H., Garcia, C.A., Morse, L.S., 1999. Therapeutic benefits of infrared (810 nm) diode laser macular grid photocoagulation in prophylactic treatment of non-exudative age-related macular degeneration, two-year results of a randomized pilot study. Ophthalmology 106, 2082–2090.

Otsu, N., 1979. A threshold selection method from gray level histograms. IEEE Trans. Syst. Man Cybernet. 9 (1).

Papoulis, A., 1984. Probability, Random Variables, And Stochastic Processes. McGraw-Hill, New York.

Pappas, T., 1992. An adaptive clustering algorithm for image segmentation. IEEE Trans. Signal Process. 40 (4), 1.

Peli, E., Lahav, M., 1986. Drusen measurements from fundus photographs using computer image analysis. Ophthalmology 93, 1575–1580.

Pitas, I., Venetsanopoulos, A.N., 1990. Nonlinear Digital Filters. Kluwer, Dordrecht.

Russ, J.C., 1999. The Image Processing Handbook.

Sahoo, P.K., Soltani, S., Wong, K.C., Chen, Y.C., 1988. A Survey of Thresholding Techniques. Academic Press, New York.

Shin, D.S., Javornik, N.B., Berger, J.W., 1999. Computer-assisted, interactive fundus image processing for macular drusen quantitation. Ophthalmology 106, 6.

Smiddy, W.E., Fine, S.L., 1984. Prognosis of patients with bilateral macular drusen. Opthalmology 91, 271–277.

The Choroidal Neovascularization Prevent Trial Research Group, 1998. Laser treatment in eyes with large drusen: Short-term effects seen in a pilot randomized clinical trial. Opthalmology 105, 11–23.