Drusen Detection in a Retinal Image Using Multi-level Analysis

Lee Brandon and Adam Hoover

Electrical and Computer Engineering Department
Clemson University
{lbrando,ahoover}@clemson.edu
<http://www.parl.clemson.edu/stare/>

Abstract. This paper concerns a method to automatically detect drusen in a retinal image without human supervision or interaction. We use a multi-level approach, beginning with classification at the pixel level and proceeding to the region level, area level, and then image level. This allows the lowest levels of classification to be tuned to detect even the faintest and most difficult to discern drusen, relying upon the higher levels of classification to use an ever broadening context to refine the segmentation. We test our methods on a set of 119 images containing all types of drusen as well as images containing no drusen or other potentially confusing lesions. Our overall correct detection rate is 87%.

1 Introduction

Drusen are deposits of cellular waste that form beneath the retina. They are the primary manifestation of age-related macular degeneration, the leading cause of late age blindness. Limited drusen formation is considered normal in older patients but excessive formation can lead to serious problems. Retinal detachment and cell death can result from drusen blocking blood flow to nearby cells. Nearby areas of the retina can also be stretched and distorted, hurting visual acuity [1].

A common method to screen for drusen is through retinal imaging. This paper concerns a method to automatically detect drusen in a retinal image without human supervision or interaction. This method could be applied to screening, reducing costs by eliminating the need for an eyecare professional to examine the large quantity of images produced for screening. This method could also be applied to treatment evaluation, by providing a quantified measurement of drusen presence that is objective and repeatable.

In an image of the retina, drusen appear as yellowish, cloudy blobs. Drusen exhibit no specific size or shape, and can appear with a variety of characteristics. Drusen are often classified as either hard or soft. Hard drusen tend to be smaller and more sharply defined than soft drusen and are generally less harmful. Soft drusen may be accompanied by other symptoms such as new vessel formation or fluid build-up in the macula. Drusen can form in groups or clusters, for example radially around the fovea or near the optic nerve head. Drusen can have a similar appearance to other lesions, such as cotton wool spots. Some faint drusen can

Fig. 1. The image on the left shows a large amount of soft, confluent drusen, while the image on the right shows a relatively smaller amount of fine hard drusen towards the right side of the image

also appear similar to normal features of the retina, such as the background pattern caused by the choroidal vessels. Figure 1 shows two examples of images¹ containing drusen.

There are very few works published specifically on automated drusen detection or segmentation in retinal imagery. Sbeh et. al. [2] approached the problem using a new morphology operator to detect the brightest points (peaks) within individual drusen. Rapantzikos and Zervakis [3] approached the problem using an adaptive local histogram to identify an appropriate local threshold for segmenting each druse. The results shown in each of these papers are promising, but their methods were applied only to image sub-areas that contained drusen. The specific areas in each image were selected manually prior to processing.

In this paper, we introduce a technique to detect and segment (quantify) drusen in a full retinal image. We report our results on images containing a variety of types of drusen, as well as on images containing no drusen and on images containing lesions that could be confused with drusen. To our knowledge, we are the first to report such results.

2 Methods

We use a multi-level approach to drusen detection, beginning with classification at the pixel level and proceeding to the region level, area level, and then image level. At the pixel level, many retinal features have a similar appearance to drusen. Therefore a pixel-level classifier will over-segment, producing a large number of false positives. At the region level the segmentation can be pruned by passing only those regions that appear druse-like. The area and image level clas-

¹ Contrast enhanced for better viewing

sification steps further prune the segmentation, successively taking into account larger contextual properties.

For our final result we produce a pixel-level segmentation of drusen, and one of five classifications for the image: no drusen, fine few drusen, fine many drusen, large few drusen, or large many drusen. The image-level classification provides a rough indication of the amount and type of drusen present in the image that may be used for screening or diagnostic functions.

2.1 Preprocessing

The preprocessing phase consists of two steps: intensity based equalization, and wavelet based equalization. For our methods we use only the green band of a color retinal image since most of the drusen detail lies in that band.

The intensity based equalization algorithm shifts the local average intensity toward 128 using a sliding window approach. This has the effect of making the image background uniform without sacrificing local detail. The illumination equalization algorithm is fully described in [4].

The wavelet-based equalization step was developed because some images are sharper and than others. This caused background variations in sharp images to have stronger positive responses to the pixel classifier than actual drusen in the smoother images. Since the pixel level classifier uses wavelet based analysis, we sought to somehow normalize the images in the wavelet space.

The algorithm works by transforming the image into a wavelet space using an orthonormal wavelet basis. In our case we use a high order coiflet [5]. Next, each dilation (frequency) band's wavelet coefficients are shifted up or down so that the average coefficient value over the entire band is equal to some predetermined constant. The image is then inversed transformed and the preprocessing phase is complete.

The value of the predetermined constant c_n for each dilation band is determined by Equation 1 where b determines base dilation band energy, m determines the rate of change, and n is dilation band number. For our image set, we used an m value of 4 for the low frequency half of the dilation band and 3 for the high frequency half. This produces output images that are relatively smooth at low frequencies and sharp at high frequencies. Figure 2 illustrates these effects.

$$c_n = b * m^{-n} \quad (1)$$

A similar algorithm has been used to enhance radiological images [6]. Here however, our purpose is different. We are trying to equalize detail levels rather than enhance them. Any enhancement that aids human viewing is a secondary benefit.

2.2 Pixel Level Classification

For pixel level classification we apply the idea of a “wavelet signature”. The idea is that a pixel within a druse should exhibit a characteristic response to

Fig. 2. Wavelet-based equalization with varying m . Using $m = 4$ produces a smoothing effect. Using $m = 3$ produces a sharpening effect. Using $m = 4$ for low frequencies and $m = 3$ for high frequencies sharpens detail while smoothing background variations

a wavelet transform regardless of the size, shape, or texture of the druse. The orthonormal wavelet bases used in the preprocessing phase are not as useful here since they do not provide the resolution needed and are not as flexible in shape. Since we do not need our transform to be non-destructive we can use a discrete approximation of any continuous time wavelet. Our approach uses discretized approximations of the Mexican hat wavelet [5].

We first compute four one dimensional wavelet responses centered around the pixel for a given dilation: one horizontal, one vertical, one diagonal, and one backwards diagonal. Each one dimensional response is computed by convolving a 1D slice of the image with the mexican hat wavelet. The product of the four one dimensional responses is then used as a measure of the similarity about the pixel and at the given dilation to the Mexican hat wavelet. When these values are calculated for a set of discrete dilations we obtain a feature vector hereafter referred to as the pixel's "wavelet signature". This vector, as expected, has a nice smooth peak for any pixel near the center of a druse or other blob-like feature. On the other hand, non-druse pixels tend to produce a wavelet signature that is either not smooth or shows a peak in the wrong dilation range.

We form the wavelet signature for classification using multiple of 10 dilations from 10 to 100 pixels. We then feed the wavelet signature into a feed forward neural network which outputs the pixel classification. Figure 3(b) shows an example result from this step.

2.3 Region Level Classification

For region level classification, we use the output of the pixel classifier as a labelling on the preprocessed image. Groups of connected, labelled pixels form regions. We can now measure properties of each region in the preprocessed image to see if it qualifies as a druse. We define the following properties for individual regions: region area, a ; region average intensity, i_{ave} ; region maximum intensity, i_{max} ; average border intensity, b_{ave} ; maximum border intensity, b_{max} ; minimum

border intensity, b_{min} ; average gradient, g_{ave} ; gradient standard deviation, g_{σ} ; horizontal density, μ_h ; and vertical density, μ_v .

Gradients are calculated using a 5x5 Sobel operator. The border properties are computed using pixels on the perimeter of the region.

$$\mu_h = \frac{\sum_{k=0}^a |r_k - x_c| p_k}{i_{ave} a} \quad \mu_v = \frac{\sum_{k=0}^a |c_k - y_c| p_k}{i_{ave} a} \quad (2)$$

where (x_c, y_c) is the centroid of the region, r_k and c_k are the row and column of the k th pixel in the region, and p_k is k th pixel's intensity. The density measures are low when the region is symmetric, compact, and has its intensity concentrated in the center. When these criteria are not met the density measures increase.

The classification rules are listed in Table 1. If a region does not pass all of these rules it is eliminated. The specific values used in each test were selected through repeated experimentation.

2.4 Area Level Classification

Area level classification is done on the preprocessed image using the region level output as a labelling. Area classification uses a 100 by 100 pixel window to define an individual area. For computational efficiency, the window is centered on each region in the image rather than convolved through the entire image. Each area classification label is only applied to the region in the center of the area.

Area level classification is done in two passes: a filtering pass and a classification pass. The filtering pass is used to eliminate obvious false positives so that the area level statistics are as accurate as possible for the classification pass.

The classification rules for area filtering are listed in Table 2. If an area does not pass all of these rules, the region at its center is eliminated. The region properties introduced in the previous step are used once again, but we use the average and standard deviation of each property over all regions in the area. The

Table 1. Rules for region level analysis. A region will be eliminated as a potential druse if it does not pass all of these tests.

	Equation	Description
1	$a \geq 4$	Eliminate stray pixel responses
2	$\mu_h < 3.5, \mu_v < 3.5,$ $\mu_h + \mu_v < 5.5$	Regions must be relatively round with brightness concentrated toward the center
3	$g_{ave}/g_{\sigma} > 1.8$	Gradient must be relatively uniform around entire region
4	$b_{min} > 50, b_{max} < 160$	Eliminate false positives induced by very strong nearby non-drusen objects (vessels, lesions, optic nerve etc...)
5	$30 < g_{ave} < 190$	Drusen gradients fall within a certain range. High gradients often indicate lesions and low gradients often indicate nothing is really there.
6	if $a < 30$ then $i_{ave} - b_{ave} > 6$	Eliminate background noise induced false positives.

standard deviation of the background intensity, σ_{bg} , and the average background gradient, g_{bg} , are also used. We define the area background as all unlabelled pixels in the area. As with the region classification rules, specific values used in each test for both area classification passes were selected through repeated experimentation.

The classification pass looks at the properties of each area and attempts to give a specific classification for the region at the center. This classification will be none, fine few, fine many, large few, or large many. The rules for the classification pass are listed in Table 3.

2.5 Image Level Classification

The last step in the process is image level classification. In this step we use the count and average size of each region class from the area level classification step to give an overall image classification. This takes place in two steps. First we classify large and fine drusen seperately as none, few, or many. We then give an

Table 2. Rules for area level filtering. This pass eliminates obvious false positives at the area level so that area level statistics will be more accurate in the classification pass.

	Equation	Description
1	$\sigma_{BG} < 16.5$	High variance of non-labelled pixels indicates positive responses were caused by a non-drusen feature such as a lesion
2	$\sigma_a < 1.5 * \bar{a}$	We expect clustered drusen to have similar size
3	$\bar{g}_\sigma < 0.5 * \bar{g}_{ave}$	A stricter version of the region level gradient variation rule applied to the averages across an area rather than individual regions
4	$\bar{\mu}_h < 2.6, \bar{\mu}_v < 2.6$	A stricter version of the region level density rule applied to averages for the area
5	$\bar{i}_{max} - \bar{b}_{max} > 6$	Most regions in the area are significantly brighter than their surroundings
6	$150 > \bar{g}_{ave} > 35$	A stricter version of the region level average gradient rule applied to averages for the area

Table 3. Rules for area-level classification.

Classification	Rules
Large Many	$regions \geq 8, 20 < \bar{a} < 115$ $\sigma_{bg} < 15, g_{bg} < 60$
Fine Many	$regions \geq 8, \bar{a} \leq 20$ $\sigma_{bg} < 15, g_{bg} < 60$
Large Few	$regions < 8, 20 < \bar{a} < 115$ $\sigma_{bg} < 10, g_{bg} < 35, \bar{g}_{ave}/\bar{g}_\sigma > 2.5, \bar{i}_{ave} - \bar{b}_{ave} > 7$
Fine Few	$regions < 8, \bar{a} \leq 20$ $\sigma_{bg} < 10, g_{bg} < 35, \bar{g}_{ave}/\bar{g}_\sigma > 2.5, \bar{g}_{ave} > 70.0, \bar{i}_{ave} - \bar{b}_{ave} > 9,$
False Positive	All others.

overall image classification as none, fine few, fine many, large few, or large many. The rules for image level classification are listed in Table 4.

3 Results

The algorithm was run on a set of 119 images 62 of which contained drusen, 38 of which were of healthy retinas, and 19 of which contained nondrusen abnormalities. The results were evaluated based on two criteria: percent correct (drusen versus non-drusen) and percent correct (exact classification). Both statistics are based on comparison to a ground truth provided by an ophthalmologist. The full results are available at www.parl.clemson.edu/stare/drusen/. A results summary is shown in Table 5.

Table 4. Rules for image level classification. The rules are evaluated from the top down and the first rule passed gives the overall image classification.

Classification	Requirements
Large Many	At least 15 regions classified as large many. More regions classified as large many than false positive.
Fine Many	At least 20 regions classified as fine many. More regions classified as large many than false positive.
Large Few	At least 2 regions classified as most likely large few.
Fine Few	At least 2 regions classified as most likely fine few.
N	All others.

Table 5. Results (percentage of 119 images correctly classified)

Catagory	Correct (D vs. ND)	Correct (Exact Classification)
Normal	89%	89%
Large Many	100%	85%
Large Few	85%	54%
Fine Many	90%	40%
Fine Few	79%	47%
Abnormal	74%	74%
Overall	87%	71%

Fig. 3. Results for second image from Figure 1 at each stage of processing

4 Conclusion

This paper presented a method to automatically detect drusen in a retinal image without human supervision or interaction. We used a multi-level approach, beginning with classification at the pixel level and proceeding to the region level, area level, and then image level. This allowed the lowest levels of classification to be tuned to detect even the faintest and most difficult to discern drusen, relying upon the higher levels of classification to use an ever broadening context to refine the segmentation. We tested our methods on a set of 119 images containing all types of drusen as well as images containing no drusen or other potentially confusing lesions. Our overall correct detection rate is 87%.

References

1. A. Abdelsalam, L. Del Priore and M. Zarbin, "Drusen in Age-Related Macular Degeneration: Pathogenesis, Natural Course, and Laser Photocoagulation-Induced Regression" *Survey of Ophthalmology* Vol. 44, Num 1, August 1999.
2. Zakaria Ben Sbeh and Laurent D. Cohen, "A New Approach of Geodesic Reconstruction for Drusen Segmentation in Eye Fundus Images" *IEEE Transactions on Medical Imaging*, Vol. 20, Num 12, December 2001.
3. K. Rapantzikos and M. Zervakis, "Nonlinear enhancement and segmentation algorithm for the detection of age-related macular degeneration (AMD) in human eye's retina" *Image Processing, 2001. Proceedings. 2001 International Conference on*, Volume: 3, 2001 Page(s): 1055–1058 vol.3
4. A. Hoover and M. Goldbaum, "Locating the Optic Nerve in a Retinal Image Using the Fuzzy Convergence of the Blood Vessels", *IEEE Transactions on Medical Imaging*, to appear.
5. I. Daubechies, *Ten Lectures on Wavelets*, SIAM, Philadelphia, 1992.
6. M. Vidaurrezaga, L.A. Diago and A. Cruz, "Contrast Enhancement with Wavelet Transform in Radiological Images" *EMBS International Conference*, July 23–28, 2000.