

AUTOMATED DRUSEN DETECTION IN A RETINAL
IMAGE USING MULTI-LEVEL ANALYSIS

A Thesis
Presented to
the Graduate School of
Clemson University

In Partial Fulfillment
of the Requirements for the Degree
Master of Science
Electrical Engineering

by
Lee Brandon
May 2003

Advisor: Adam Hoover

ACKNOWLEDGMENTS

I would like to thank my advisor, Dr. Hoover, for introducing me to the field of machine vision and giving me the opportunity to work on this project. This research would not have been possible without his guidance, encouragement, and countless hours of brainstorming.

I would also like to thank Dr. Hoover and Dr. Schalkoff for their classroom instruction which has provided the foundation of my knowledge in this field, and thank Dr. Schalkoff and Dr. Walker for taking the time to serve on my committee.

Finally, I'd like to thank my family for all of their support and encouragement throughout the course of my graduate studies.

May 2, 2003

To the Graduate School:

This thesis entitled “Automated Drusen Detection in a Retinal Image Using Multi-level Analysis” and written by Lee Brandon is presented to the Graduate School of Clemson University. I recommend that it be accepted in partial fulfillment of the requirements for the degree of Master of Science with a major in Department of Electrical Engineering.

Adam W. Hoover, Advisor

We have reviewed this thesis
and recommend its acceptance:

Robert J. Schalkoff

Ian D. Walker

Accepted for the Graduate School:

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 develop a wavelet-based feature vector for pixel-level classification which uses discretized versions of the Mexican hat wavelet. 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%.

Contents

TITLE PAGE	i
ABSTRACT	ii
1 Introduction	6
1.1 Overview of Approach	7
1.2 STARE Project	10
1.3 Related Work	11
2 Multi-level Analysis	14
2.1 Preprocessing	15
2.2 Pixel Level Classification	20
2.3 Region-level Classification	26
2.4 Area-level Classification	32
2.5 Image-level Classification	36
3 Results	38
4 Conclusions	48
A Wavelet Transform Overview	50
A.1 Continuous Wavelet Transforms and Introduction to Wavelet Theory .	51
A.2 Framed Wavelet Transforms	53

A.3 Orthonormal wavelet transforms 53

List of Figures

1.1	Retinal image examples of each drusen class	8
1.2	Example showing the need for multilevel analysis.	9
2.1	Overview of multi-level analysis	15
2.2	Intensity based equalization example	16
2.3	Wavelet-based equalization with various m values	18
2.4	Normalizing effect of wavelet-based equalization	19
2.5	Preliminary efforts at drusen detection	21
2.6	The Mexican hat wavelet	22
2.7	Drusen wavelet signatures	24
2.8	Non-drusen wavelet signatures	25
2.9	Pixel-level classification example	26
2.10	Region-level density rule	29
2.11	Region-level uniform gradient rule	30
2.12	Region-level border intensity rule	30
2.13	Region-level border gradient rule	31
2.14	Region-level classification example	31
2.15	Area-level background variance rule	33

	4
2.16 Area-level intensity difference rule	34
2.17 Area-level classification example	35
3.1 Large many example result	39
3.2 Fine many example result	40
3.3 Large few example result	41
3.4 Fine few example result	42
3.5 Normal image example result	43
3.6 Abnormal example result	44
3.7 Examples of debatable ground truths	46
3.8 Example of multiple drusen classifications in a single image	47
A.1 Orthonormal wavelet transform example	54

List of Tables

2.1	Dilation band average values for images in Figure 2.4	17
2.2	Rules for region-level classification	28
2.3	Rules for area-level filtering	33
2.4	Rules for area-level classification.	36
2.5	Rules for image-level classification	37
3.1	Image classification results summary	45
3.2	Image classification results summary with “reasonable” statistic	47

Chapter 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 also appear similar to normal features of the retina, such as the background pattern caused by the choroidal vessels.

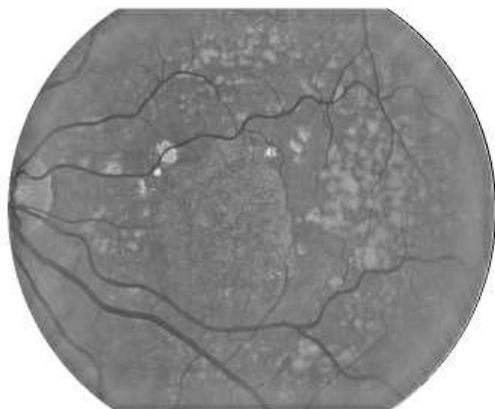
In this thesis we divide drusen into four classifications: large many, large few, fine many, or fine few. We test our methods on each of these four classes as well as images of normal retinas and of retinas containing various types of non-drusen abnormalities. Figure 1.1 shows examples¹ of each drusen class as well as a normal retina and a retina containing non-drusen lesions.

1.1 Overview of Approach

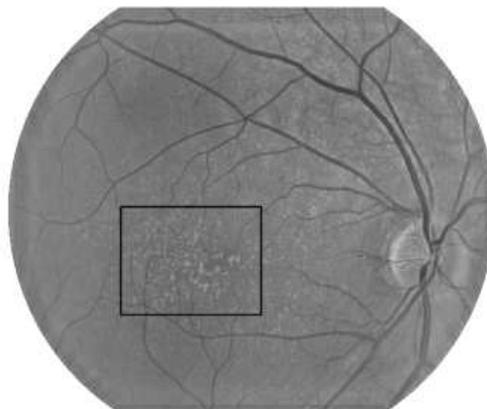
In this thesis, we introduce a technique to detect and segment (quantify) drusen in a full retinal image. We report our results on images containing a variety of types of drusen, as well as on images containing no drusen and on images containing lesions that could be confused with drusen. To our knowledge,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

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 classification steps further prune the segmentation, successively

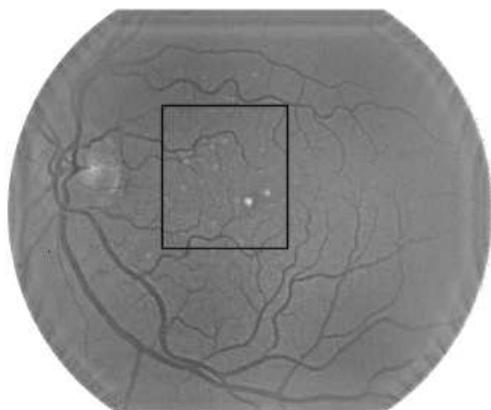
¹Contrast enhanced for better viewing



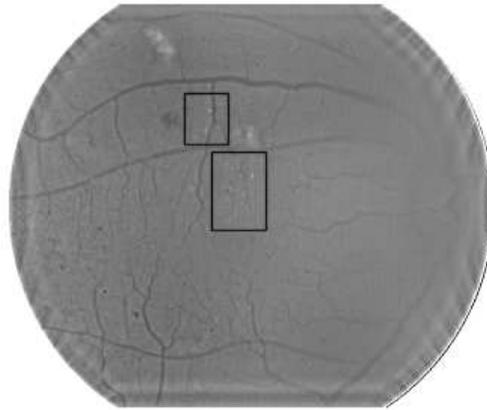
(a) Large many drusen



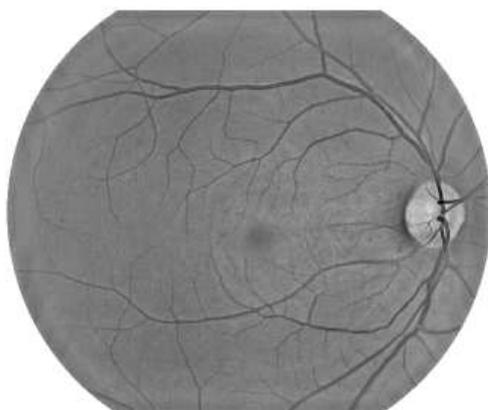
(b) Fine many drusen



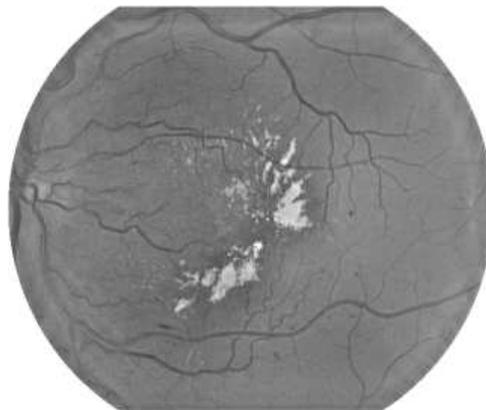
(c) Large few drusen



(d) Fine few drusen



(e) Normal retina



(f) Abnormal retina (lesions)

Figure 1.1: Retinal images showing each of the four drusen classes as well as a normal retina and a drusen free abnormal retina. Drusen in (a) is present throughout the image. Boxes have been superimposed in (b), (c), and (d) to indicate drusen locations.

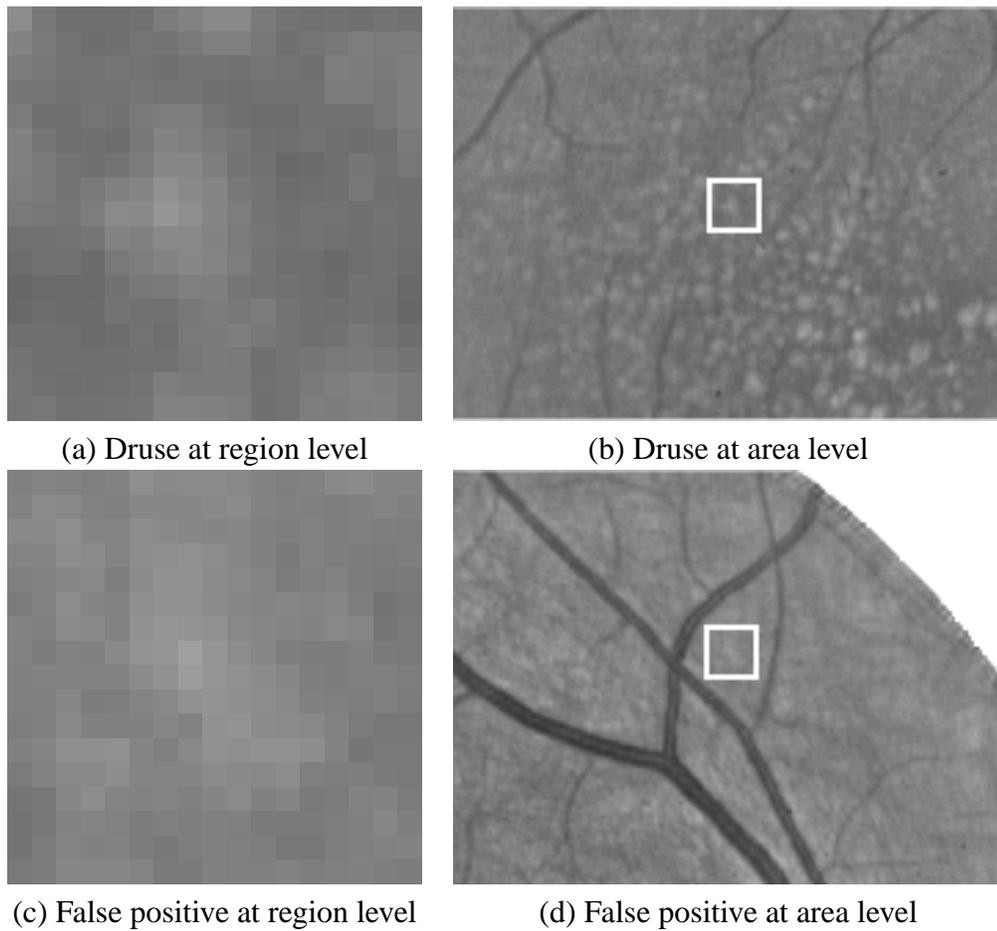


Figure 1.2: Example showing the need for multilevel analysis.

taking into account larger contextual properties. This approach 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. Figure 1.2 shows an example region of a druse and a region of a non-druse that at the region-level cannot be distinguished, but at the area level (when compared to similar surrounding regions) can be distinguished.

Pixel-level classification is the most complex processing stage and serves as the basis of our approach. In this stage, a wavelet-based feature vector is developed which describes

the shape of the area around the pixel at multiple resolutions. An artificial neural network is then used for classification. The region-level stage uses region growing and region properties to analyze groups of connected pixels. The area-level stage analyzes the properties of clusters of regions and produces a set of statistics. The image-wide statistics are then used to produce the overall image classification.

Our final result is 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.

1.2 STARE Project

This work is part of the STARE (STuctured Analysis of the REtina) project. The overall goal of STARE is to create a system which is capable of automated diagnosis of retinal images. Additional project goals include the ability to detect and measure key features, annotate the image contents, and to track manifestation change over time in individual subjects.

The STARE project was started by Micheal Goldbaum, M.D., in 1975 at the University of California, San Diego, and is currently funded by the National Institutes of Health (U.S.A.). Over thirty people have contributed to STARE, with backgrounds ranging from medicine to science to engineering. Images and clinical data have been provided by the Shirley Eye Center at the University of California, San Diego, and the Veterans Administration Medical Center in San Diego.

An ophthalmologist is a doctor who specializes in the structure, function, and diseases of the human eye. Ophthalmologists rely on visual examination of the retina to detect and

diagnose disorders. Retinal imaging provides a permanent record of a patient's eye that can be consulted at any time for diagnostic reasoning. A retinal image is acquired by imaging the inner surface of the eyeball through the patient's pupil using an optical camera. A retinal image generally shows the optic nerve, fovea, surrounding vessels, and the retinal layer.

Previous efforts in the STARE project have resulted in numerous publications². Some efforts concentrated on vessel segmentation [13], optic nerve localization [12] [14], image registration [9], the measurement of vessel tortuosity [11], and generic lesion detection [10]. This thesis presents the first effort under the STARE project directly targeted at detecting and classifying drusen.

1.3 Related Work

The varying characteristics of drusen prevent simple image processing methods from segmenting arbitrary drusen. Color and intensity based thresholding methods can not distinguish between drusen, other lesions, the optic nerve head and other features. The boundaries of drusen are vague and difficult to define, making edge detection ineffective. The size and shape of drusen vary greatly, making template matching methods impractical.

There are very few works published specifically on automated drusen detection and segmentation in retinal imagery. Sbeh et. al. [22] approach the problem using a new morphology operator to detect the brightest points (peaks) within individual drusen. Raptzikos and Zervakis [20] approach the problem using an adaptive local histogram to identify an appropriate local threshold for segmenting each druse. Thaibaoui et. al. [25] use a fuzzy logic approach to classify pixels as drusen or non-drusen based on intensity levels. The results shown in each of these papers are promising, but unlike our work, their

²Full publication list available at <http://www.parl.clemson.edu/stare>.

methods were applied only to image sub-areas that contained drusen. The specific areas in each image were selected manually prior to processing.

Other works in drusen processing have concentrated on image enhancement [24] or interactive approaches to segmentation [4] [8]. In [24], the effect of stereo viewing on the ability of an ophthalmologist to correctly classify drusen is explored. The effect of other standard image enhancements such as image contrast is also tested. Donohoe et. al. [8] present a system which allows an ophthalmologist to efficiently label drusen. This is done by the human user manually selecting a pixel location and then manually adjusting a region growing threshold so that the entire region is labelled. The system then has the capability to measure properties of the labelled regions in the image such as region sizes and total region count. While enhancement is not the focus of our work, we do present a preprocessing technique which can greatly aid the visualization of drusen. Our technique is not intended as an interactive tool, but it could be applied to refined human interaction.

More work has been done in the general field of retinal image segmentation. Some retinal image processing problems include blood vessel segmentation [13] [16] and detection of the optic nerve [12] [14]. Hoover et. al. [13] approach blood vessel segmentation via matched spatial filtering at different orientations. Leandro et. al. [16] use wavelet transforms for blood vessel segmentation. As in our work, they take advantage of the ability of the framed wavelet transform to describe shape at multiple scales. Hoover and Goldbaum [12] [14] develop a method to locate the optic nerve which takes advantage of the fact that all retinal blood vessels converge on the optic nerve. Work has also been done on identifying and segmenting disease manifestations other than drusen. Ballerini [3] introduces a method to measure the properties of the foveal avascular zone and detect irregularities which could indicate diabetic retinopathy. The method uses active contours to define the borders of the foveal avascular zone and detects irregularities using region properties such

as moments and gradients. Our work could benefit each of these efforts by recognizing drusen and eliminating them from consideration prior to processing.

The idea of multi-level, or multi-scale, segmentation has been approached in a number of different ways. Scale-space models [2] [15] [19] form the basis of the most common type of multi-scale segmentation. These models seek to represent the original image at varying scales. This allows image features to be segmented at scales appropriate to their size. For example, at scales where each pixel represents a large part of the original image, large image features are easier to segment. Zhang et. al. [29] introduce an interesting neural net approach to multi-level segmentation of medical images. A top level neural net provides the coarsest segmentation, and each top level segmented element is passed to lower level neural nets which perform successively finer segmentations. The idea of modelling the image at the pixel, region, and area level presented in this thesis is not commonly used. A similar method for analyzing mammograms is presented by Yu and Guan [27]. In their paper, wavelet-based feature vectors are used for pixel classification and properties of individual regions and groups of regions formed by the pixel-level segmentation are analyzed.

The ability of the wavelet transform to represent an image at multiple scales has also been used previously for segmentation. Wavelet based filtering has been applied successfully to problems such as blood vessel detection in retinal images [16]. Zayed et. al. [28] develop wavelet-based feature vectors to segment fetal ultrasound images. A number of variations on the wavelet-based feature vector technique have been used to detect breast cancer as well [5] [27] [30]. As in our work, Yu and Guan [27] use a feed-forward neural network to classify pixels based on wavelet feature vectors. This work differs from ours however in that an orthonormal transform is used. This produces a very small feature vector whose resolution would be insufficient for drusen segmentation.

Chapter 2

Multi-level Analysis

Our approach consists of five steps: preprocessing, pixel-level classification, region-level classification, area-level classification, and image-level classification. The preprocessing step equalizes the grey-level average intensity and contrast of features both within the image and between images. The pixel-level step uses wavelet based analysis to identify candidate drusen pixels. The region-level step analyzes each connected group of candidate pixels to determine whether it meets the criteria of an individual druse. The area-level step, which analyzes clusters of regions, proceeds in two phases. The first pass, hereafter referred to as the area-level filtering pass, is used to eliminate obvious false positives at the area level so that statistics gathered in the second pass are as accurate as possible. The second area-level pass attempts to give each region an exact classification. The image-level classification is based on area-level statistics. Figure 2.1 outlines the multi-level approach.

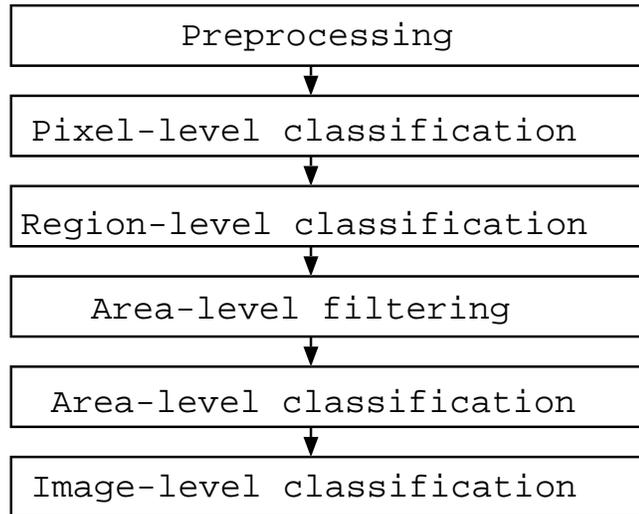


Figure 2.1: Overview of multi-level analysis

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 mid-grey (e.g. 128 in an 8 bit image) using a sliding window approach. This has the effect of making the image background uniform without sacrificing local detail. The intensity based equalization algorithm is fully described in [14] and is illustrated in Figure 2.2.

The wavelet-based equalization step was developed to compensate for the varying contrast in the test image set. This caused background variations in sharp images to have positive responses to the pixel-level classifier and actual drusen in the smoother images to be missed. Since the pixel-level classifier uses wavelet based analysis, we sought to somehow normalize the images in the wavelet space.

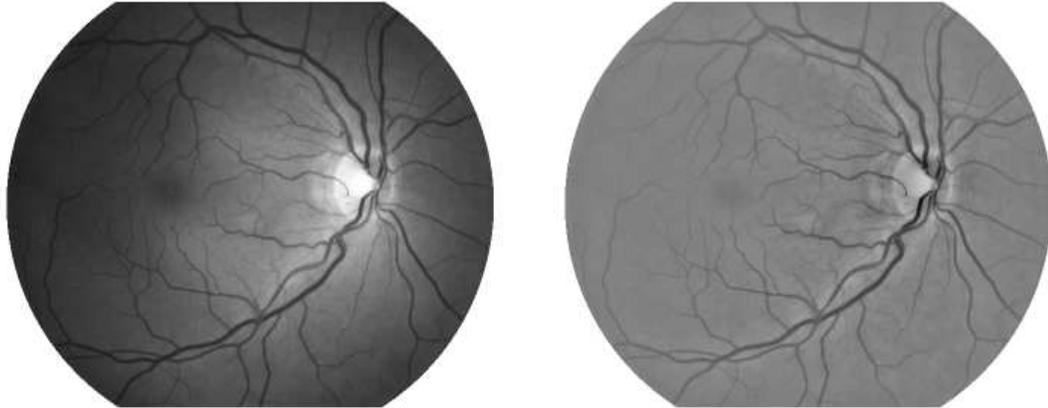


Figure 2.2: Example of intensity based equalization. Notice that the non-uniform illumination is corrected.

The algorithm works by transforming the image into a wavelet space by using an orthonormal wavelet transform. In our case we use a high order coiflet [7]. Next, each dilation 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 inverse transformed and the preprocessing phase is complete.

To make our images compatible with the orthonormal wavelet transform we pad them with zeros to produce a 1024 by 1024 matrix. When transformed this yields $\log_2 1024 - 1$, or 9, dilation bands. We number these bands from zero to eight with zero being the largest dilation, and therefore the lowest frequency, and eight being the highest frequency.

The value of the predetermined constant c_n for each dilation band is determined by

$$c_n = b * m^{-n} \quad (2.1)$$

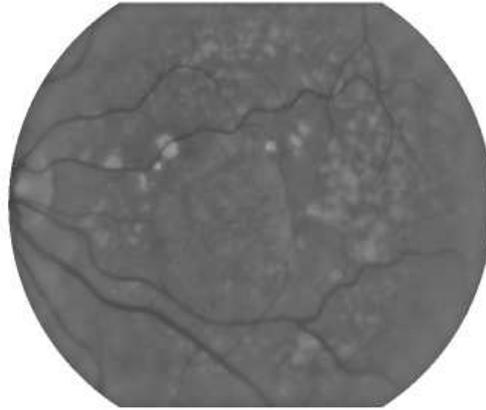
where b determines base dilation band energy, m determines the rate of change, and n is the dilation band number. The m value can be thought of as tuning low versus high frequency

n	Figure 2.4(a)	Figure 2.4(b)	Target
0	12179	11999	12000
1	3385	3383	3000
2	700	696	750
3	178	179	188
4	50	62	47
5	14.4	20.7	15.6
6	4.4	7.7	5.2
7	1.2	2.2	1.7
8	0.32	0.6	0.58

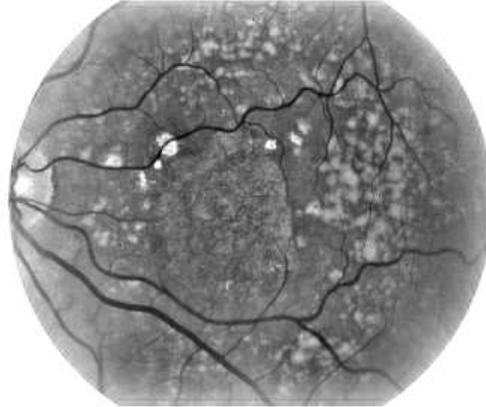
Table 2.1: Average dilation band values before wavelet-based equalization for images in Figure 2.4. The target values are average dilation band values after equalization. Notice that in the high frequency bands ($n = 5$ through $n = 8$) the low contrast image values are consistently under the target while the high contrast image values are consistently over the target.

ratios and changing the overall smoothness or sharpness of the output image. In our work we found that varying m could produce favorable results. For our image set, we used an m value of 4 for the low frequency half ($n=0$ through $n=4$) of the dilation band and an m value of 3 for the high frequency half ($n=5$ through $n=8$). This produces output images that are relatively smooth at low frequencies and sharp at high frequencies. Figure 2.3 illustrates these effects.

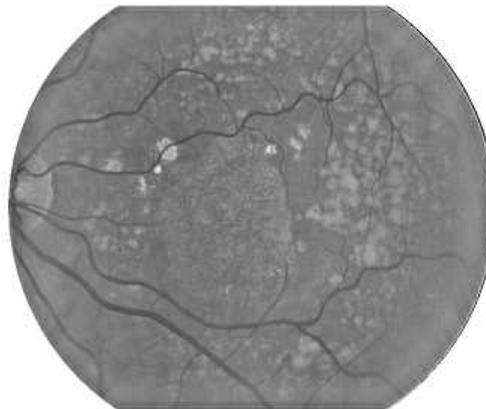
Figure 2.4 (a) and (b) shows two intensity equalized images with significantly different contrast levels. Table 2.1 shows the average dilation band values for the two images before equalization as well as the target values given by Equation 2.1. As expected, the high contrast image has relatively large values in the high frequency dilation bands while the low contrast image has relatively small values. Figure 2.4(c) and (d) shows the images after wavelet-based equalization. The contrast levels are similar in each image.



(a) $m = 4$ (all frequencies)



(b) $m = 3$ (all frequencies)



(c) $m = 4$ (low frequencies) and $m = 3$ (high frequencies)

Figure 2.3: 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.

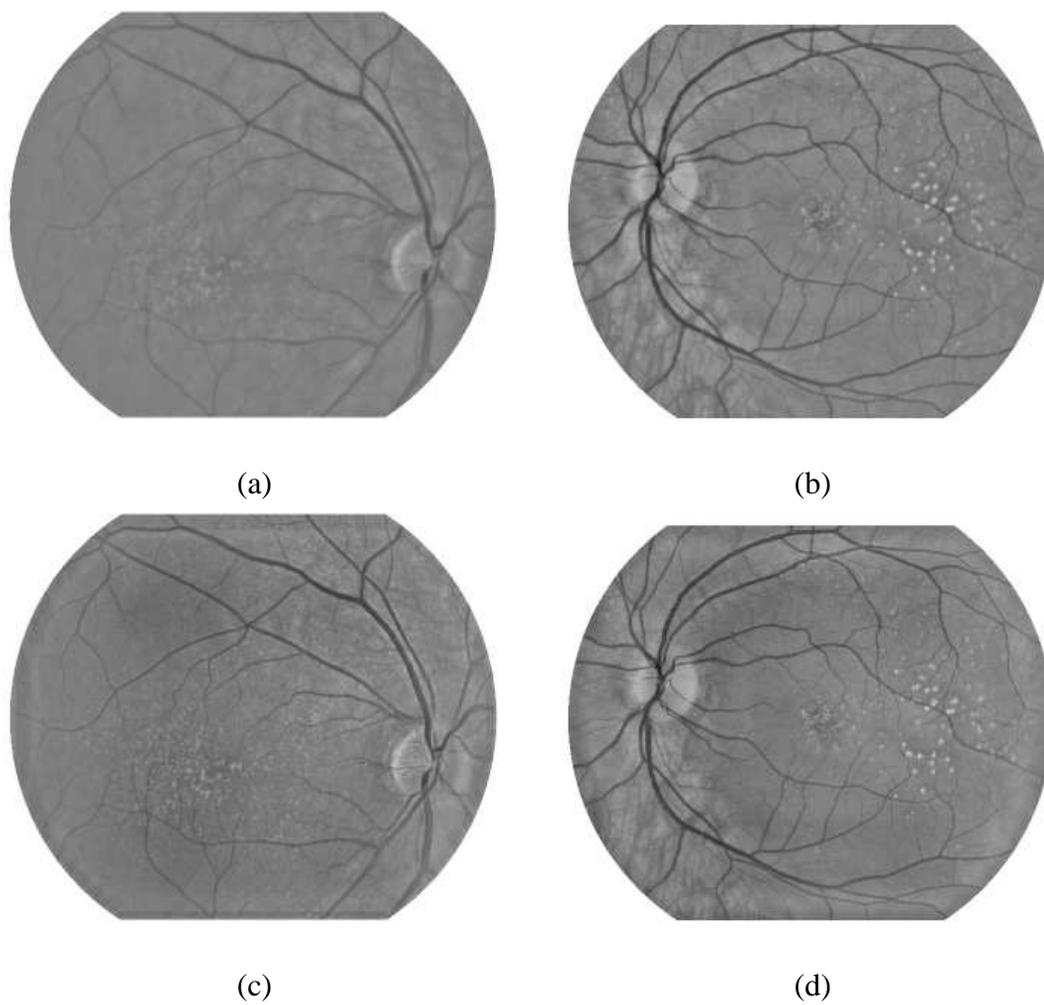


Figure 2.4: Images before and after wavelet-based equalization. The low contrast image (left) is sharpened while the high contrast image (right) is smoothed.

A similar algorithm has been used to enhance radiological images [26]. Our purpose, however, is different. We are trying to equalize detail levels between images 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 a wavelet transform regardless of the size, shape, or texture of the druse.

Our initial drusen detection work focused on the idea that drusen should produce a high energy peak in the Fourier space corresponding to their approximate average size. We used Fourier band-pass filters to attempt to extract drusen from the image. Fourier filtering met with limited success primarily due to the non-repeating and spatially localized nature of drusen. Drusen clusters are not continuous sine waves, but individual peaks of varying size and amplitude. The best result Fourier filtering could produce was a blur in the general area of a drusen patch, but of course, similar blurs were produced by other features as well. Figure 2.5(a) shows the effect of Fourier filtering drusen.

Next we tried the wavelet transform. The wavelet transform offered two big advantages over the Fourier transform: spatial localization and basis function flexibility. There were several basis functions for the orthonormal wavelet transform which were similar in shape to a typical druse. Using these basis functions we tried band-pass filtering the images. This technique was capable of enhancing drusen in the images but was not capable of producing a useful pixel level segmentation. Orthonormal wavelet transforms, like the discrete Fourier transform, produce data only at powers of two. Close examination of the wavelet space revealed that drusen did not stand out in any of the dilations represented by

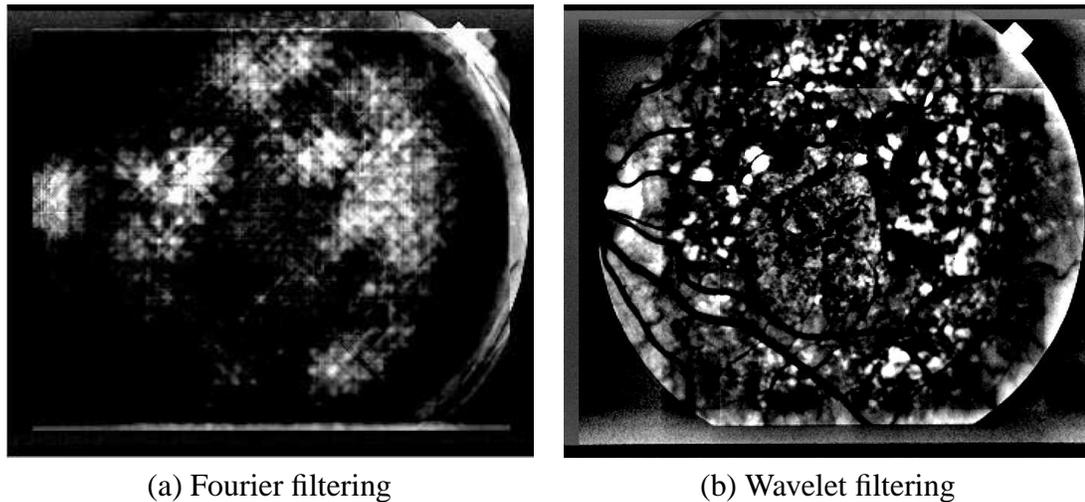


Figure 2.5: Preliminary work in drusen detection. (Original image is shown in Figure 1.1a.) (a) Fourier filtering (blood vessels are removed manually prior to processing) (b) Wavelet filtering using the Daubechies 4 basis function (blood vessels not removed).

the orthonormal wavelet transform. Figure 2.5(b) illustrates the effect of wavelet filtering drusen using an orthonormal wavelet transform.

In order to increase our resolution in scale-space and allow more flexibility for our basis wavelet we turned to the framed wavelet transform. The framed wavelet transform is simply a discretized version of the continuous wavelet transform (see Appendix A for an overview of wavelets). Our approach uses the Mexican hat wavelet, which has a basic shape similar to that of an individual druse [7]. Figure 2.6 shows the Mexican hat wavelet.

Framed wavelet transforms can provide a non-destructive representation of a signal if every integral dilation is represented. Since we don't need a complete representation of the image we can sample the wavelet space only in the dilation range we are interested in. We found empirically that a dilation range of 10 to 100 pixels works well for drusen detection, and sampling only multiple of ten dilations within that range provides adequate resolution.

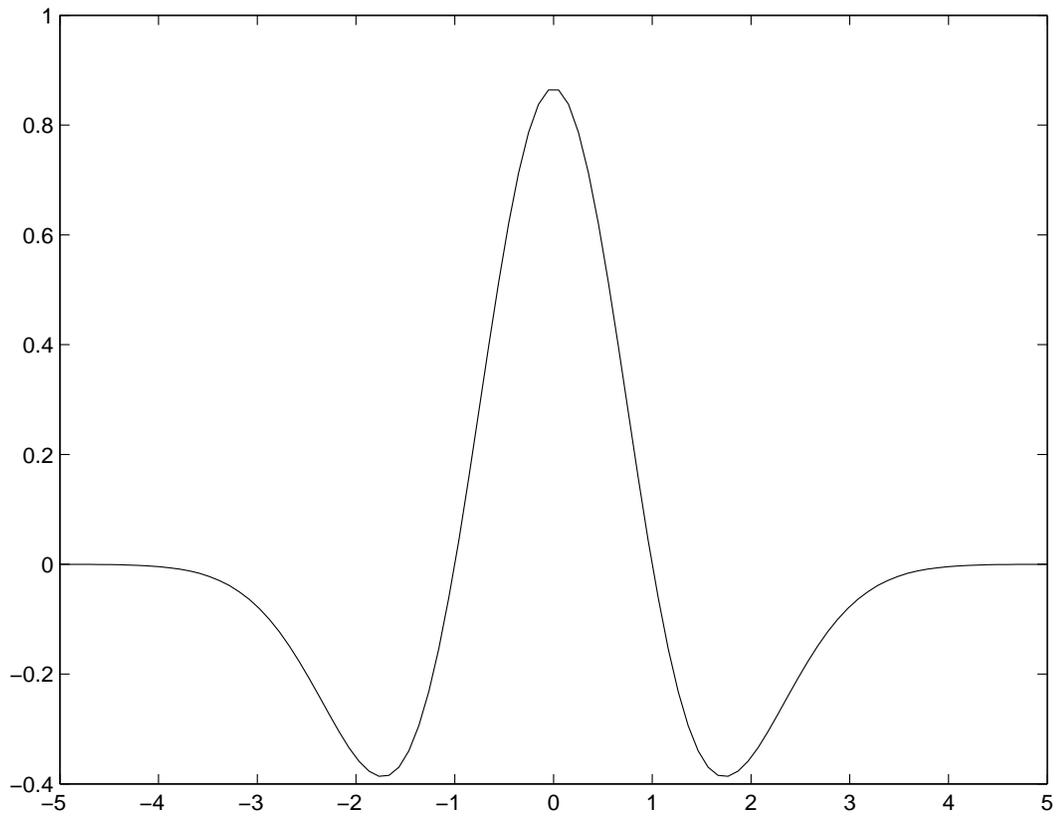
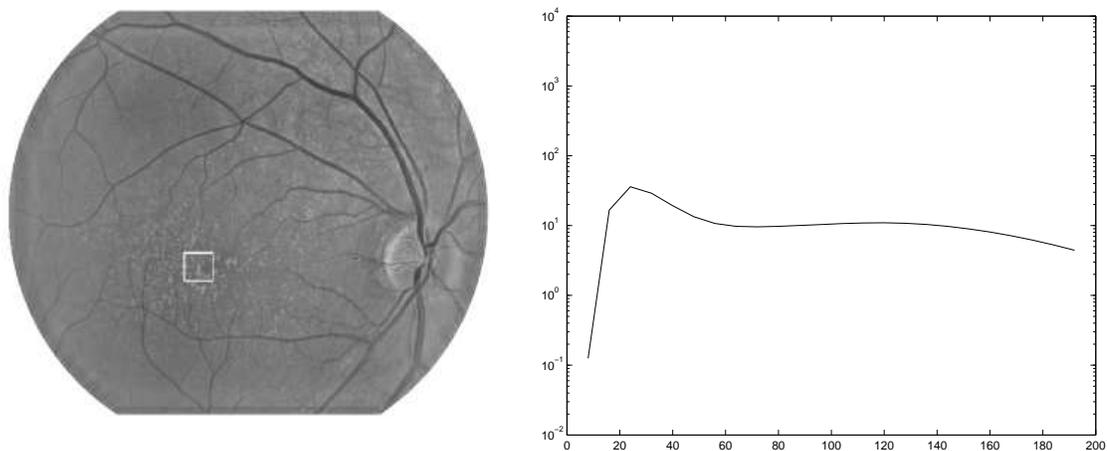


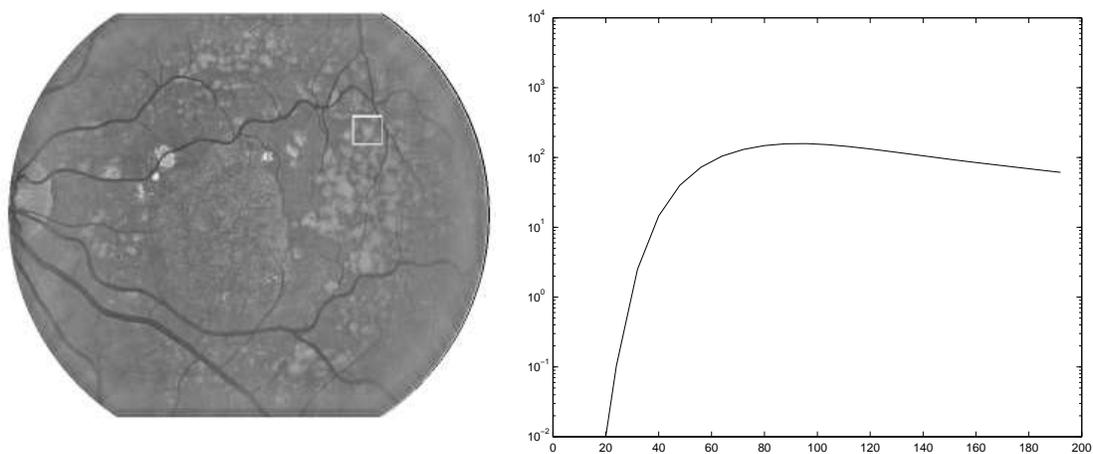
Figure 2.6: The Mexican hat wavelet.

We additionally simplify the computational complexity of our approach by sampling the wavelet space of one dimensional slices through the pixel we are interested in. To form the wavelet signature for a given pixel we compute four one dimensional wavelet responses centered around the pixel at each dilation: one horizontal, one vertical, and one for each diagonal. Each one dimensional response is computed by convolving the one dimensional slice of the image with the Mexican hat wavelet. Once the four one dimensional responses are computed, their product is then used as a measure of the similarity about the pixel and at the given dilation to the Mexican hat wavelet. The set of these similarity values for all dilations forms the pixel's wavelet signature. This wavelet signature vector, as expected, has a nice smooth peak for any pixel near the center of a druse or other blob-like feature. Non-druse pixels tend to produce a wavelet signature that either has an amplitude that is out of the expected range or peaks in the wrong dilation range. Figure 2.7 shows example wavelet signatures for fine and large drusen. Figure 2.8 shows several non-drusen features that can be confused with drusen.

Once we have a pixel's wavelet signature we classify it as drusen or non-drusen using a feed forward neural network with sigmoidal activation functions. The neural net has three layers structured as follows: 11 inputs (wavelet signature plus bias), 10 hidden units, and 2 output units. The two outputs are trained to represent the degree of certainty the pixel is druse and the degree of certainty the pixel is not druse. The neural network was trained over approximately 1000 sample pixels of which about half were drusen and half were non-drusen. The neural network is then used to classify each pixel in a given image. The final output is a labelling on the preprocessed image. If the degree of certainty that a pixel is druse is greater than the degree of certainty the pixel is not druse then the pixel is labelled. We refer to pixels labelled by this step as "candidate pixels" and the labelling is passed to

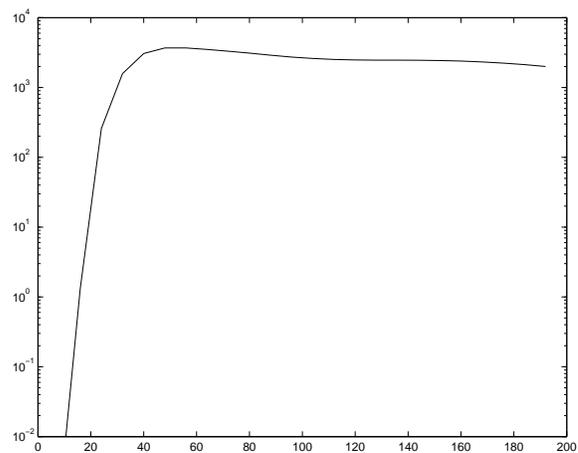
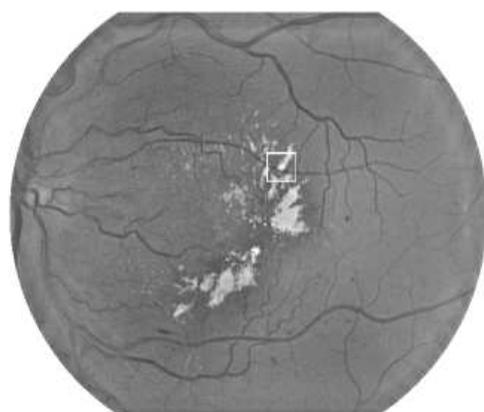


(a) Fine Drusen

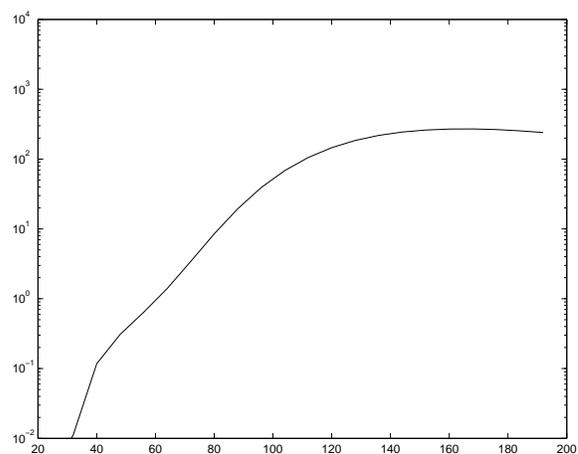
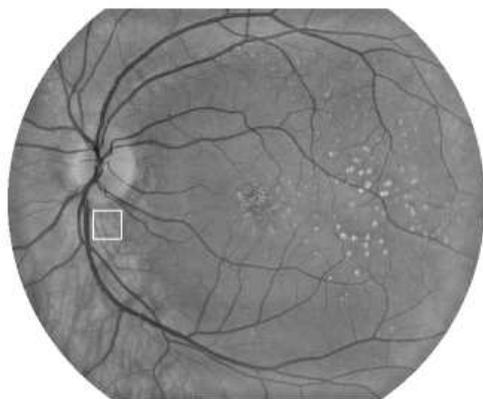


(b) Large Drusen

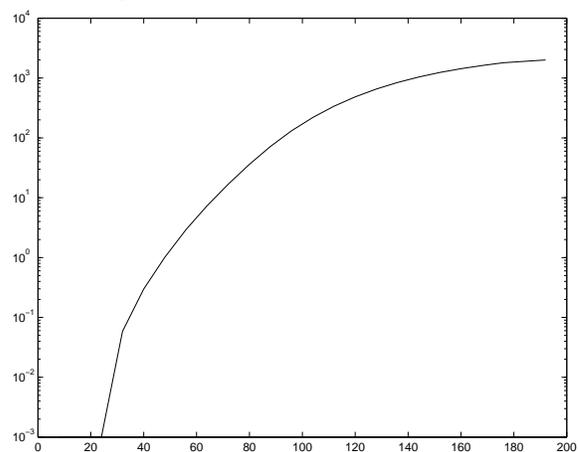
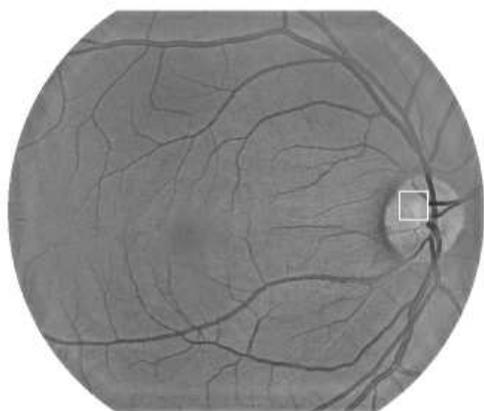
Figure 2.7: Typical wavelet signatures for large and fine drusen.



(a) Lesion (amplitude out of tolerance)



(b) Choroidal Vessels (dilation range out of tolerance)



(c) Optic Nerve (both amplitude and dilation range out of tolerance)

Figure 2.8: Wavelet signatures for various image features which can be confused with drusen.

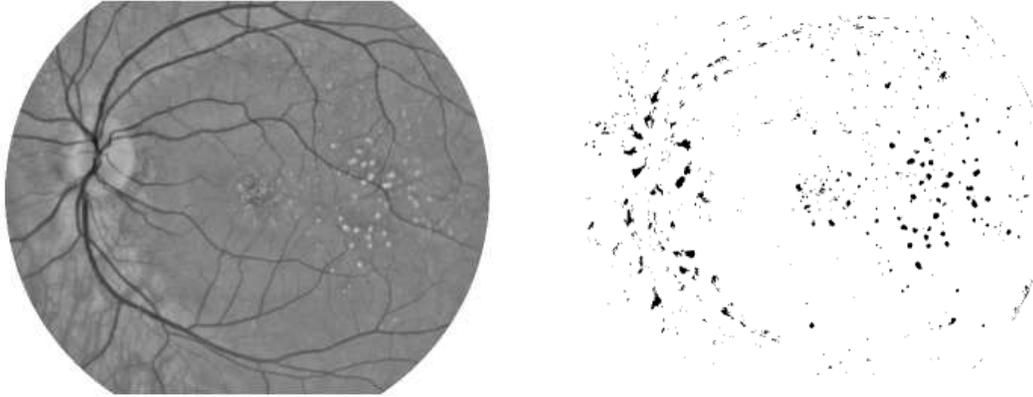


Figure 2.9: A preprocessed image and its pixel classification.

the region-level classifier to be further refined. Figure 2.9 shows an example result from this step.

2.3 Region-level Classification

For region-level classification, we use the output of the pixel-level classifier as a labelling on the preprocessed image. Groups of 4-connected, labelled pixels form regions. We measure properties of each region in the preprocessed image to see if the region 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}
- border average intensity: b_{ave}

- border maximum intensity: b_{max}
- border minimum intensity: b_{min}
- border average gradient: g_{ave}
- border gradient standard deviation: g_{σ}
- horizontal density: μ_h
- 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. The density measures are calculated as follows:

$$\mu_h = \frac{\sum_{k=0}^a |r_k - x_c| p_k}{i_{ave} a} \quad \mu_v = \frac{\sum_{k=0}^a |c_k - y_c| p_k}{i_{ave} a} \quad (2.2)$$

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.

Simple rules are now developed based on the region properties to eliminate observed false positives from the pixel-level classifier. The classification rules are listed in Table 2.2. If a region does not pass all of these rules it is eliminated. Regions passing all tests are referred to as candidate regions and are passed on to the area-level classifier. The specific values used for each rule were selected through repeated experimentation.

The first region-level rule eliminates stray pixel responses. We define a stray response as less than four connected candidate pixels. Stray pixel responses tend to occur in areas where the image has a grainy texture.

	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 bright or dark 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

Table 2.2: Rules for region-level classification. A region will be eliminated as a potential druse if it does not pass all of these tests.

The second region-level rule uses the moment based density measures to eliminate elongated, skewed, and non-compact regions. We expect a druse to be relatively round and compact with its intensity concentrated near its center and uniformly balanced around its centroid. Elongated regions often occur parallel to blood vessels or along the edge of the optic nerve. Skewed regions usually occur near lesions or other high energy non-drusen features. The intensity of the region tends to slope away from the high energy feature, causing the measures in Equation 2.2 to increase. Non-compact regions can be induced by choroidal vessels or other irregularly shaped bright patches in the image. Figure 2.10 shows examples of false positive regions from the pixel-level classifier which are eliminated by this rule.

The third region-level rule eliminates regions having an average border gradient that is not uniform. The rule targets false positives induced by nearby strong gradients such as

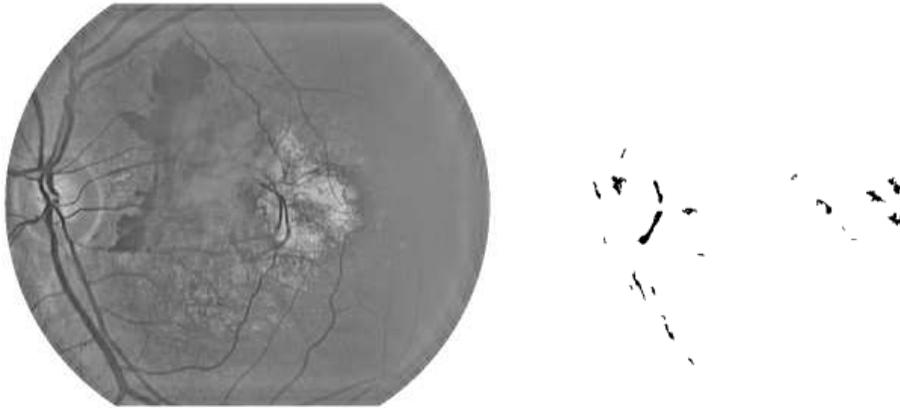


Figure 2.10: False positives eliminated by region-level rule two. The image on the right shows false positives produced by the pixel-level classifier which are eliminated at the region level due to elongation, skew or lack of compactness.

blood vessel edges as well as regions which are part of a larger feature such as a lesion. Rules two and three are complimentary, and can overlap in the false regions eliminated.

The fourth region-level rule takes advantage of the fact that drusen do not normally occur adjacent to very bright or very dark retinal features, but false positives with drusen-like appearance often do appear adjacent to very bright or very dark retinal features. Figure 2.12 shows the effect of this rule. The fifth region-level rule constrains the gradient range of a candidate drusen. This rule eliminates bright compact false positives which are only discernable from drusen by their excessively high gradients as well as false positives induced by the cloudy appearance of choroidal vessels. Figure 2.13 shows the effect of this rule. The sixth region-level rule further helps to eliminate small false positives induced by background image textures. A final region-level classification result, after all rules are applied, is shown in Figure 2.14.

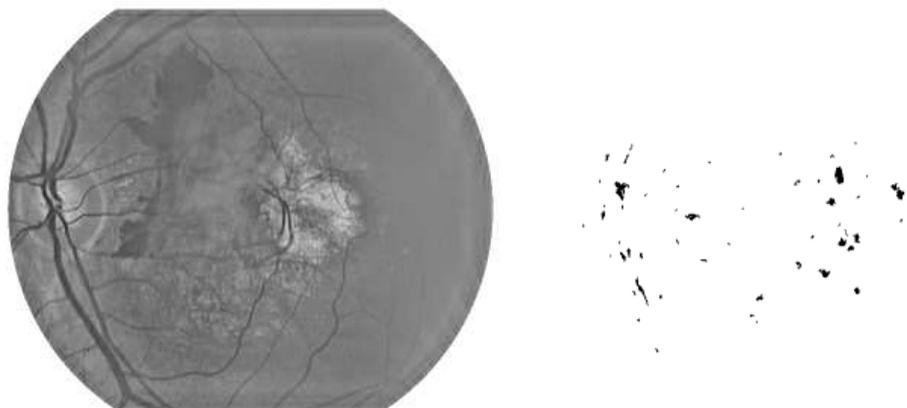


Figure 2.11: False positives eliminated by region-level rule three. The image on the right shows false positives produced by the pixel-level classifier which are eliminated at the region level due gradient variation.

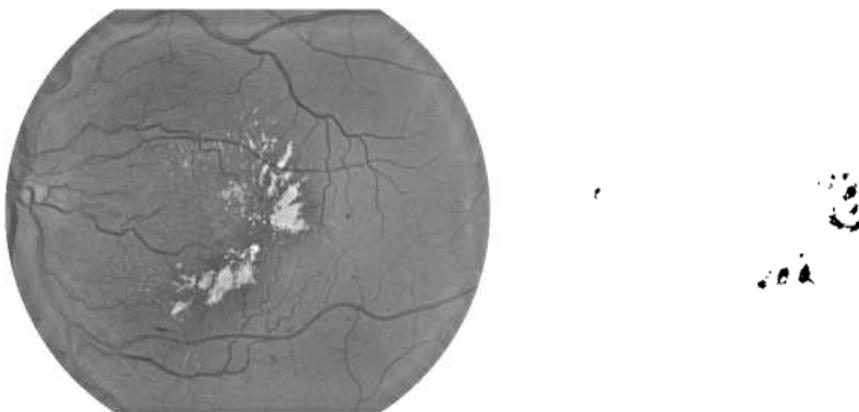


Figure 2.12: False positives eliminated by region-level rule four. The image on the right shows false positives produced by the pixel-level classifier which are eliminated at the region level due to adjacency to very bright or dark features.

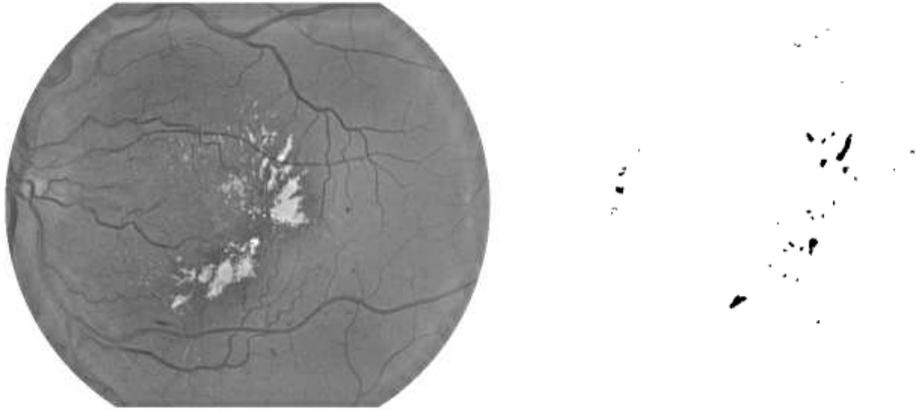


Figure 2.13: False positives eliminated by region-level rule five. The image on the right shows false positives produced by the pixel-level classifier which are eliminated at the region level due to out of tolerance gradient strength.



Figure 2.14: Pixel-level classification results from Figure 2.9 and the refined segmentation after region-level classification.

2.4 Area-level Classification

Area-level classification is done on the preprocessed image using the region level output as a labelling. Area-level classification uses a 100×100 pixel window to define an individual area. The window is centered on the centroid of each region in the image rather than convolved through the entire image. Each area-level classification label is only applied to the region in the center of that 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 best pixel-wise drusen labelling comes out of this step, however, the second area-level pass and the image-level classification steps are needed to provide a global classification.

The classification rules for area-level filtering are listed in Table 2.3. 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 look at their averages and standard deviations over all regions in the area. The 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-level classification rules, specific values used in each test for both area-level classification passes were selected through repeated experimentation.

The first area-level filtering rule is based on the assumption that drusen should be the only feature present in a drusen cluster. If there is a large amount of variance in the unlabelled pixels then non-drusen features are probably inducing false positives. Figure 2.15 shows an example of the effect of this rule.

	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	$\bar{i}_{max} - \bar{b}_{max} > 6$	Most regions in the area are significantly brighter than their surroundings
3	$\sigma_a < 1.5 * \bar{a}$	We expect clustered drusen to have similar size
4	$\bar{g}_\sigma < 0.5 * \bar{g}_{ave}$	A stricter version of the region-level gradient variation rule applied to the averages across an area rather than individual regions
5	$\bar{\mu}_h < 2.6, \bar{\mu}_v < 2.6$	A stricter version of the region-level density rule applied to averages for the area
6	$150 > \bar{g}_{ave} > 35$	A stricter version of the region-level average gradient rule applied to averages for the area

Table 2.3: 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.

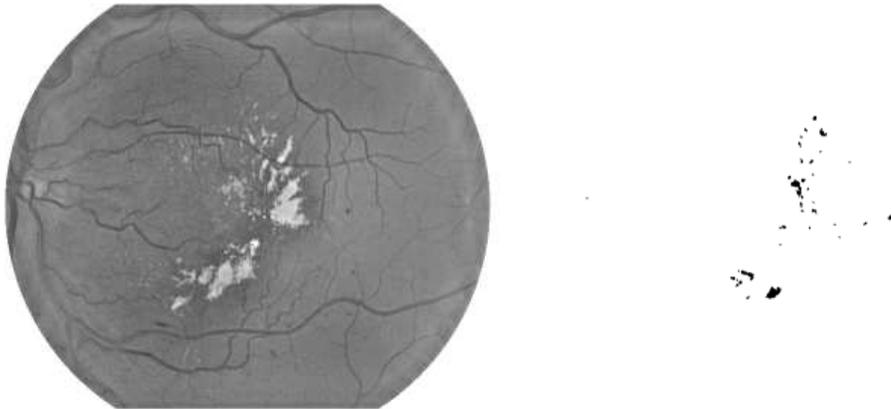


Figure 2.15: False positives eliminated by area-level filtering rule one. The image on the right shows false positives produced by the region-level classifier which are eliminated at the area level due to high background variance.

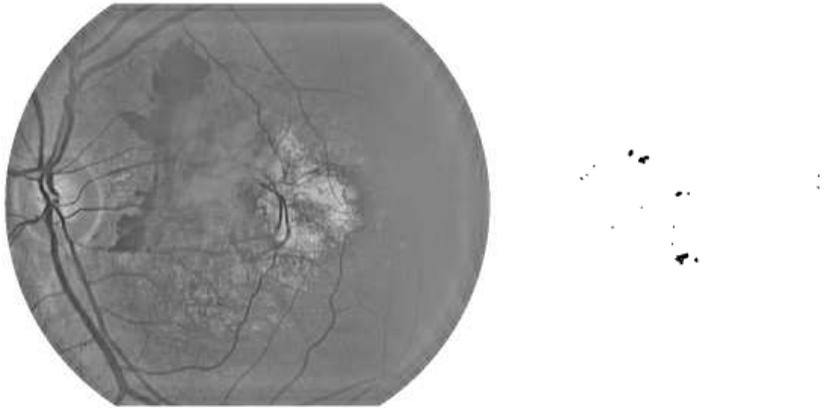


Figure 2.16: False positives eliminated by area-level filtering rule two. The image on the right shows false positives produced by the region-level classifier which are eliminated at the area level due to lack of difference from background intensity.

The second area-level filtering rule eliminates regions in drusen clusters where the average region intensity does not stand out from the background. This rule primarily targets false positives induced by faint background fluctuations such as those caused by choroidal vessels. Figure 2.16 shows an example of the effect of this rule.

The third area-level filtering rule relies on the assumption that clustered drusen should be relatively similar in size. This rule also focuses on eliminating non-drusen abnormalities. The remainder of the area-level filtering rules are simply stricter versions of region-level rules applied to averages of region properties for all regions in an area. A sample area-level filtering result is shown in Figure 2.17.

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. In this pass we take advantage of the different properties of each of these five classifications.



Figure 2.17: Region-level classification results from Figure 2.14 and the refined segmentation after area-level filtering.

This is the first classification step that is not tuned to err on the side of letting all drusen pass since its goal is to provide statistics for image-wide classification. In this pass we have a different rule set for each class. If none of the rule sets are passed, the region is classified as a “none”. The rules for the classification pass are listed in Table 2.4.

A candidate region is first classified as a potential large many, fine many, large few, or fine few based on the number and average size of the regions in its area. The region is then subjected to tests specific to its class. Both of the many classes must pass only one class specific test, a limitation on background standard deviation and average gradient. A background standard deviation test is also applied in area-level filtering, but here we use a stricter limit since additional false positives have already been eliminated in the filtering pass.

The few classes can be subjected to stricter rule sets than the many classes because we assume the individual drusen are well spaced and do not overlap. A stricter version of the background standard deviation and average gradient rule used for the many classes

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.

Table 2.4: Rules for area-level classification.

is applied to the few classes. The few classes must pass a stricter rule limiting border gradient variance as well as a rule which places a minimum requirement on the difference between region average intensity and border average intensity. For the fine few class only, a minimum is set for the average border gradient.

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 separately as none, few, or many. We then give an overall image classification as none, fine few, fine many, large few, or large many. The rules for image-level classification are listed in Table 2.5. As can be seen from the table, a hierarchy is used to resolve situations where requirements are met for more than one image class.

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 fine 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.
None	All others.

Table 2.5: Rules for image-level classification. The rules are evaluated from the top down and the first rule passed gives the overall image classification.

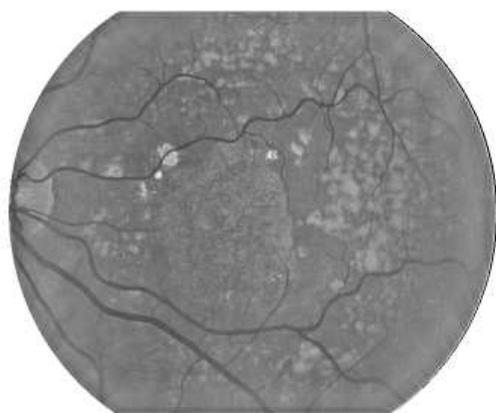
Chapter 3

Results

The algorithm was evaluated on a set of 119 images, 62 of which contained drusen, 38 of which were of healthy retinas, and 19 of which contained non-drusen abnormalities. Of the drusen images, 20 were large many, 13 were large few, 10 were fine many, and 19 were fine few. Sample results for each image class are given in Figures 3.1- 3.6¹. The results were evaluated based on two criteria: percent correct (drusen versus no drusen) and percent correct (exact classification). The percent correct (drusen versus no drusen) statistic is also referred to as the drusen detection rate. Both statistics are based on comparison to a ground truth provided by an ophthalmologist. A summary of the image classification results is shown in Table 3.1.

As expected, drusen detection percentages for the many classes were very high. Only one many drusen image in the set was misclassified as non-drusen. This failure occurred in a fine many image at the image-level classification step. The fine many regions in the image were labelled correctly by the area-level classifier, but there were an insufficient number of fine many regions to label the image fine many. This occurred because a single highly

¹Area-level classification results not shown. Full results available at www.parl.clemson.edu/stare/drusen/



(a) Preprocessed image



(b) Pixel-level classification

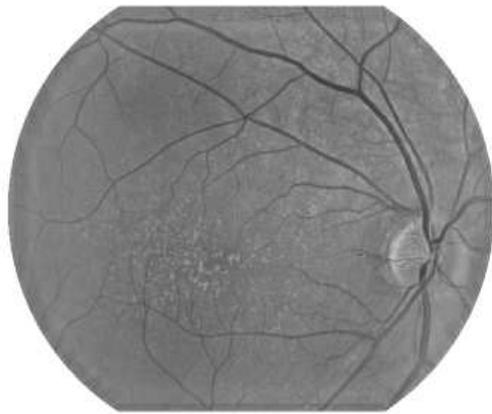


(c) Region-level classification

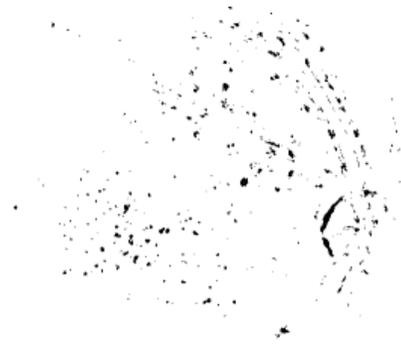


(d) Area-level filtering

Figure 3.1: Large many drusen image at each stage of processing. Image was correctly classified as large many at the image level.



(a) Preprocessed image



(b) Pixel-level classification



(c) Region-level classification



(d) Area-level filtering

Figure 3.2: Fine many drusen image at each stage of processing. Image was correctly classified as fine many at the image level.



(a) Preprocessed image



(b) Pixel-level classification

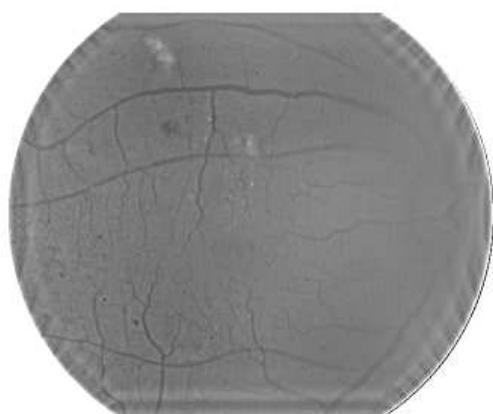


(c) Region-level classification



(d) Area-level filtering

Figure 3.3: Large few drusen image at each stage of processing. Image was correctly classified as large few at the image level.



(a) Preprocessed image



(b) Pixel-level classification

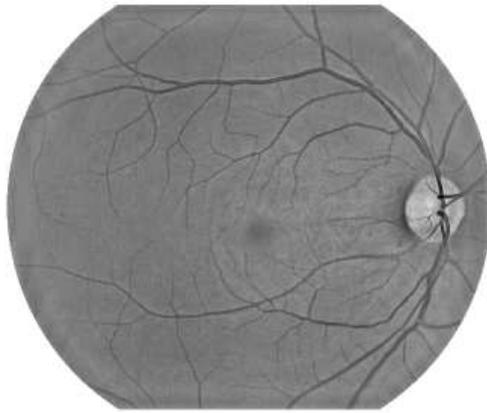


(c) Region-level classification



(d) Area-level filtering

Figure 3.4: Fine few drusen image at each stage of processing. Image was correctly classified as fine few at the image level.



(a) Preprocessed image



(b) Pixel-level classification

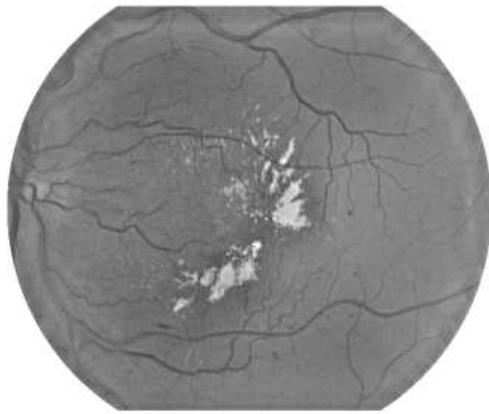


(c) Region-level classification



(d) Area-level filtering

Figure 3.5: Normal image at each stage of processing. Image was correctly classified as no drusen at the image level.



(a) Preprocessed image



(b) Pixel-level classification



(c) Region-level classification



(d) Area-level filtering

Figure 3.6: Abnormal image at each stage of processing. Image was incorrectly classified as fine many at the image level.

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%

Table 3.1: Results (percentage of 119 images correctly classified)

compact cluster was present which appeared to be many drusen at the area level, but at the image level not enough drusen were present to allow a many classification.

For normal images, the drusen detection rate was 89%, or 34 out of 38 images. Each of the failures was due to misclassification as few drusen. Of these 4 misclassifications, 3 appeared to possibly contain drusen. One of these ambiguous images is shown in Figure 3.7a.

Our method misclassified 2 out of 13 large few images and 4 out of 19 fine few images as none. A fair number of these misclassifications also had debatable ground truths. One of these debatable images is shown in Figure 3.7b. Our method failed most frequently in the abnormal catagory, missing 5 of 19 images. All five of these failures were induced by lesions. It should be noted however that there were five other lesion-containing images in the abnormal set which were correctly classified.

Percentages for exact classification were significantly lower than drusen detection for each of the drusen classes. Most of this performance decrease was due to ambiguity and difficulty in the labelling of the training set. It is difficult for a human to be consistent when using such qualitative descriptions as large, fine, many, and few. It is also possible for an image to have multiple classifications, for example, the image shown Figure 3.8 has relatively equal amounts of fine and large drusen. In fact, two images of this retina are

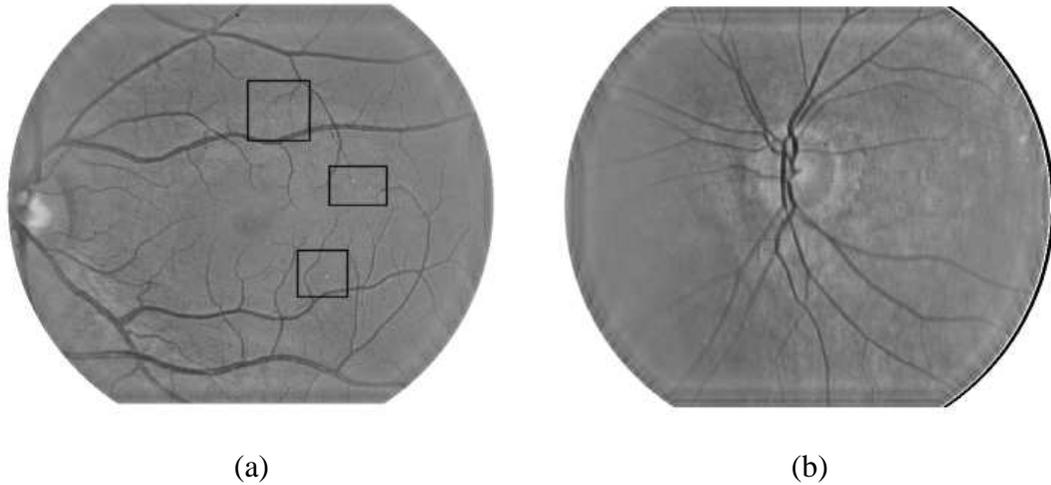


Figure 3.7: Examples of debatable ground truth classifications. The image on the left was labelled normal in the ground truth but the system classified the image as fine few drusen. Boxes indicate regions which appear to be drusen. The image on the right was labelled large few in the ground truth but the system classified the image as none.

present in the ground truth; one is labelled large many drusen and the other is labelled fine many drusen.

To fairly evaluate these results, we created a subjective third performance measure called “reasonable”. In order to pass this measure, an image’s area-level classification and image-level classification be consistent with how a human expert might classify the image. Table 3.2 restates the results summary with the reasonable statistic included.

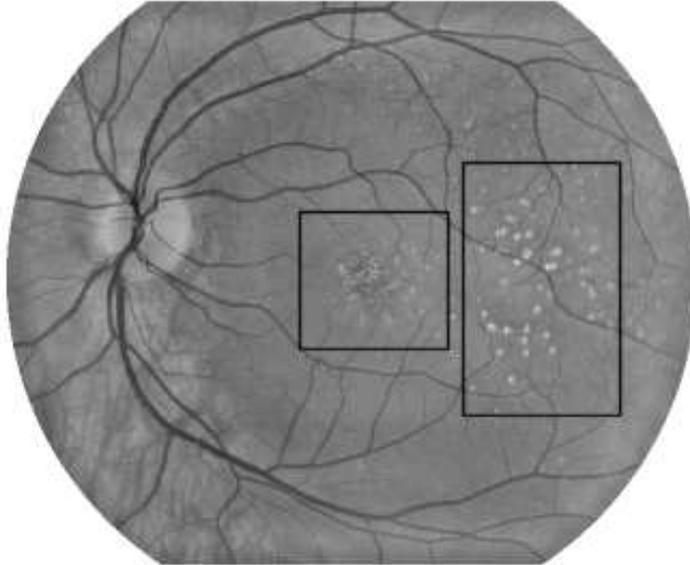


Figure 3.8: Example of an image containing multiple classes of drusen. The box on the left indicates the presence of fine drusen and the box on the right indicates the presence of large drusen.

Catagory	Correct (D vs. ND)	Correct (Exact Classification)	Reasonable
Normal	89%	89%	92%
Large Many	100%	85%	100%
Large Few	85%	54%	100%
Fine Many	90%	40%	90%
Fine Few	79%	47%	84%
Abnormal	74%	74%	79%
Overall	87%	71%	91%

Table 3.2: Results including the “reasonable” measure

Chapter 4

Conclusions

This paper presented 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 developed a wavelet based feature vector for pixel-level classification which uses discretized versions of the Mexican hat wavelet. 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% compared to a ground truth provided by an ophthalmologist.

Most failures occur in images that have non-drusen abnormalities such as lesions which induce drusen-like patterns. A human analyzing retinal images can recognize these areas as false positives by first recognizing the abnormality, but our algorithm has no such contextual knowledge (it can only recognize drusen). The algorithm is also susceptible to normal variations near the optic nerve and blood vessels. Normal intensity variations in this region

can sometimes have a drusen-like appearance. The algorithm also tends to fail when drusen are very close to the perimeter of the image. Without area context, perimeter drusen are often missed.

Combining this method with other retinal segmentation algorithms could greatly improve performance. Knowledge of the location of certain non-drusen abnormalities, the optic nerve, and blood vessels could be used to exclude most of the false positives that currently cause our method to fail.

Appendix A

Wavelet Transform Overview

Ingrid Daubechies, a pioneer in wavelet theory, summarizes the wavelet transform as follows: “The wavelet transform is a tool that cuts up data or functions or operators into different frequency components, then studies each component with a resolution matched to its scale” [7]. The wavelet transform is usually described as an extension of the Fourier transform which adds the following benefits: space-frequency localization and flexible basis functions. These two properties of the wavelet transform make it an ideal tool for analyzing the frequency content of natural or non-periodic data.

The most common use of the wavelet transform is as a data compression tool. The wavelet transform’s ability to describe natural and non-repeating signals makes it ideal for processing images, audio, and video. The jpeg2000 image compression standard replaces the cosine transform with a discrete wavelet transform, attaining significantly higher compression without noticeable image degradation [6]. Wavelets have also shown promise as signal enhancement tools. The wavelet-based equalization algorithm described in this thesis is an example of a wavelet based enhancement tool. Another area wavelets have only recently been applied to is pattern recognition problems such as image segmentation. Our

wavelet based approach to pixel-level classification is an example of wavelet based pattern recognition.

There are two major types of wavelet transforms: continuous wavelet transforms (CWT), and discrete wavelet transforms (DWT). Within the discrete wavelet transform we have two major subclassifications: framed wavelet transforms and orthonormal wavelet transforms. In this appendix we describe each of these classes of wavelet transform individually, starting with the continuous transform and an overview of the theory behind wavelet analysis.

A.1 Continuous Wavelet Transforms and Introduction to Wavelet Theory

The continuous wavelet transform (CWT) is the most fundamental wavelet transform and forms the basis for the two discrete types. In this section we will show how the CWT is an extension of the Fourier transform and explain how it works.

The continuous Fourier transform may be written as

$$X(j\omega) = \int_{-\infty}^{\infty} x(t) * e^{-j\omega t} dt \quad (\text{A.1})$$

and the continuous wavelet transform (CWT) may be written as

$$\Psi(\tau, s) = \frac{1}{\sqrt{s}} \int_{-\infty}^{\infty} x(t) * \psi\left(\frac{t-\tau}{s}\right) dt \quad (\text{A.2})$$

In the CWT, the Fourier transform's $e^{-j\omega t}$ term is replaced by the function $\psi\left(\frac{t-\tau}{s}\right)$. The function $\psi(t)$ is called the mother wavelet. The mother wavelet can be any one of a number of eligible functions. One example is the Mexican hat wavelet described in this thesis.

The variables s and τ denote scale and translation, so $\psi(\frac{t-\tau}{s})$ is simply the mother wavelet translated and scaled (dilated) to the s and τ for which the CWT is being calculated. In summary, the CWT of a function $x(t)$ at scale s and translation τ is the convolution of $x(t)$ with the scaled and translated mother wavelet, normalized by $\frac{1}{\sqrt{s}}$.

We can see from this comparison that the wavelet transform extends the Fourier transform in two ways. We replace the Fourier transform's sinusoidal basis function with a basis function of our choosing, and we retain space-time information in the transformed signal. The Fourier transform takes a signal in n dimensions of space and time and transforms it into n dimensions of frequency. The CWT transforms the same signal into n dimensions of scale (analogous to frequency) at n dimensions of translation, thus providing full frequency information at every space-time location in the original signal.

Although an infinite number of mother wavelets are possible, there are certain properties the mother wavelet must satisfy. For the mother wavelet to act as a valid basis function, it must integrate to zero over all real values:

$$\int_{-\infty}^{\infty} \psi(t) dt = 0 \quad (\text{A.3})$$

Another desirable mother wavelet property is compact support. This simply means that the mother wavelet integrates to zero over some range less than $-\infty$ to ∞ . For example, the smallest range over which the Mexican hat wavelet integrates to zero is -5 to 5 , we therefore say that it has a compact support range of -5 to 5 . Mathematically, compact support is not required, but practically, the CWT is difficult to implement without it.

A.2 Framed Wavelet Transforms

The framed wavelet transform is a discretized version of the continuous wavelet transform:

$$\Psi(\tau, s) = \frac{1}{\sqrt{s}} \int_{-\infty}^{\infty} x(n) * \psi\left(\frac{n - \tau}{s}\right) dn \quad (\text{A.4})$$

The basic idea behind the framed wavelet transform is that one can approximate a non-destructive transform of a discrete signal by computing the CWT with discrete approximations of the mother wavelet at every integral value of scale and translation. The framed wavelet transform is rarely used as a non-destructive transform however, due to the enormous computational cost of computing an accurate signal representation. The power of the framed wavelet transform lies more in its ability to extract patterns from a signal.

When sampled at reasonable intervals of scale and translation, the framed wavelet transform acts like a multi-scale matched spatial filter. Features which are similar in shape to the mother wavelet at a given scale and translation produce a high output at that scale and translation. This is the idea behind the pixel-level segmentation technique presented in this thesis. In our method, we sample the image's wavelet space at every integral translation (each pixel) and at scales (dilations) from 10 pixels to 100 pixels.

A.3 Orthonormal wavelet transforms

Orthonormal discrete wavelet transforms are by far the most commonly used type of wavelet transform. Orthonormal transforms provide fast, non-destructive, and compact image transformations which are useful for data compression and enhancement. Like the fast Fourier transform, the orthonormal wavelet transform is a one-to-one transformation that breaks the original signal into scale components at powers of 2 up to the length of the signal. For

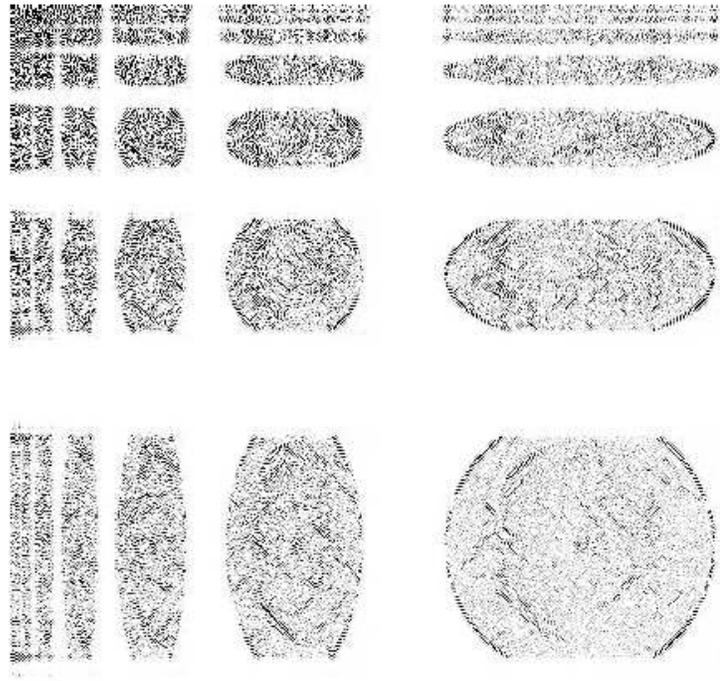


Figure A.1: Wavelet space of a retinal image transformed via an order 10 coiflet.

example, if a discrete one dimensional signal contains 1024 data points, its orthonormal wavelet transform will be composed of 512 coefficients representing the signal's responses at scale 2, 256 coefficients representing the signal's responses at scale 4, etc. When a second dimension is added, as in an image, a quadtree representation is used to store the transformed signal. Figure A.1 shows an orthonormal wavelet transform of a retinal image. The lower right corner of the image contains the wavelet coefficients for horizontal scale 2 and vertical scale 2. As one moves from right to left across the image, horizontal scale increases by powers of 2. Similarly, as one moves from bottom to top, vertical scale increases by powers of 2. The transform shown in this figure is the same one used for the wavelet-based equalization technique presented in the preprocessing section of this thesis.

Only mother wavelets which form orthonormal basis functions can be used for orthonormal wavelet transforms. More specifically, the set of all functions satisfying

$$\Psi_{m,n}(x) = 2^{-m/2}\psi(2^{-m}x - n) \quad (\text{A.5})$$

for integral values of m and n must together form an orthonormal basis. Equation A.5 defines the set of all possible functions formed by shifting the mother wavelet by n and scaling it by 2^m . A number of useful wavelet families have been derived which satisfy this condition. The wavelet families are usually chosen such that desirable properties are optimized in addition to meeting the fundamental requirement of forming an orthonormal basis. The coiflet, for example, seeks to maximize symmetry and minimize support width.

The order of an orthonormal wavelet basis determines the degree of precision with which the mother wavelet is represented. As the order increases the mother wavelet representation becomes smoother and usually meets its secondary requirements more optimally. The downside of using high order wavelet bases is that computational complexity is significantly higher. High order wavelet bases produce transforms in which frequency separation is cleaner, which is useful for enhancement. High order transforms also tend to cause more coefficients in the transformed signal to approach zero, which is useful for compression.

Bibliography

- [1] A. Abdelsalam, L. Del Priore and M. Zarbin, “Drusen in Age-Related Macular Degeneration: Pathogenesis, Natural Course, and Laser Photocoagulation-Induced Regression” in *Survey of Ophthalmology*, vol. 44, no. 1, Aug. 1999, pp. 1-29.
- [2] S. Acton and D. Mukherjee, “Scale space classification using area morphology” in *IEEE Trans. on Image Processing*, vol. 9, no. 4, Apr. 2000, pp. 623-635.
- [3] L. Ballerini, “Detection and quantification of diabetic retinopathy” in *SPIE Applications of Digital Image Processing XXII*, vol. 3808, Oct. 1999, pp. 213-223.
- [4] J. Berger, “Quantitative spatiotemporal image analysis of fluorescein angiography in age-related macular degeneration” in *SPIE Ophthalmic Technologies VIII*, vol. 3246, June 1998, pp. 48-53.
- [5] L. Bruce, R. Adhami, “Classifying mammographic mass shapes using the wavelet transform modulus-maxima method ” in *IEEE Trans. on Medical Imaging*, vol. 18, no. 12, Dec. 1999, pp. 1-8.
- [6] C. Christopoulos, A. Skodras, T. Ebrahimi, “The JPEG2000 still image coding system: an overview” in *IEEE Trans. on Consumer Electronics*, vol. 46, no. 4, Nov. 2000, pp. 1103-1127.

- [7] I. Daubechies, Ten Lectures on Wavelets, SIAM, Philadelphia, 1992.
- [8] G. Donohoe, S. Nemeth and P. Soliz, “An interactive system for computer-aided retinal image analysis” in the proceedings of *12th IEEE Symposium on Computer-Based Medical Systems*, Stamford, CT, June 18-20, 1999.
- [9] M. Goldbaum, V. Kouznetsova, B. Cote, M. Nelson, W. Hart, “Automated registration of fundus images for comparison of lesions”, in Parel J-M, Qiushi R (eds.): *Proceedings of Ophthalmic Technologies III* , SPIE Vol 1877:94-99, 1993. (research article)
- [10] M. Goldbaum, B. Cote, R. Garcia, W. Hart, P. Kube, M. Nelson, “Computer detection and classification of lesions found in diabetic retinopathy”, *Invest Ophthalmol Vis Sci* , 33:1082, 1992. (abstract)
- [11] W. Hart, M. Goldbaum, B. Cote, P. Kube, M. Nelson, “Automated measurement of retinal vascular tortuosity”, proceedings *AMIA Annual Fall Symposium* , pp. 459-463, 1997. (research article)
- [12] A. Hoover and M. Goldbaum, “Fuzzy convergence”, in the proceedings of *IEEE Computer Vision and Pattern Recognition*, 1998, pp. 716-721.
- [13] A. Hoover, V. Kouznetsova, M. Goldbaum, “Locating Blood Vessels in Retinal Images by Piece-wise Threhsold Probing of a Matched Filter Response” , in *IEEE Transactions on Medical Imaging*, vol. 19 no. 3, March 2000, pp. 203-210.
- [14] 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.

- [15] A. Kuijper and L. Florack, "Hierarchical pre-segmentation without prior knowledge" in the proceedings of *Eighth IEEE International Conference on Computer Vision*, vol. 2, Vancouver, Canada, July 9-12, 2001, pp. 487-493.
- [16] J. Leandro, R. Cesar and H. Jelinek, "Blood Vessel Segmentation in Retina: Preliminary Assessment of the Mathematical Morphology and of the Wavelet Transform Techniques" in the proceedings of *XIV Brazilian Symposium on Computer Graphics and Image Processing*, Oct. 2001, pp. 84-91.
- [17] B. McCane and T. Caelli, "Multi-scale adaptive segmentation using edge and region based attributes" in the proceedings of *1997 First International Conference on Knowledge-Based Intelligent Electronic Systems*, vol. 1, Adelaide, South Australia, May 1997, pp. 72-80.
- [18] B. McCormick and M. Goldbaum, "STARE = Structured Analysis of the Retina: Image processing of TV fundus image", presented at the *USA-Japan Workshop on Image Processing*, Jet Propulsion Laboratory, Pasadena, CA, October 31, 1975.
- [19] N. Mudigonda, R. Rangayyan, J. Desautels, "Segmentation and classification of mammographic masses" in *SPIE Medical Imaging 2000: Image Processing*, vol. 3979, June 2000, pp. 55-67.
- [20] K. Rapantzikos and M. Zervakis, "Nonlinear enhancement and segmentation algorithm for the detection of age-related macular degeneration (AMD) in human eye's retina" in the proceedings of *2001 International Conference on Image Processing*, vol. 3, 2001, pp. 1055-1058.

- [21] K. Rapantzikos, "Detection of Age-related Macular Degeneration (AMD) in Human Eye's Retina Using Image Processing and Analysis Methods" Thesis Presented to Technical University of Crete, April 2000.
- [22] Z. Sbeh and L. Cohen, "A New Approach of Geodesic Reconstruction for Drusen Segmentation in Eye Fundus Images" in *IEEE Trans. on Medical Imaging*, vol. 20, no. 12, Dec. 2001, pp. 1321-1333.
- [23] Z. Sbeh, L. Cohen, G. Mimoun, G. Coscas and G. Soubrane, "An adaptive contrast method for segmentation of drusen" in the proceedings of *International Conference on Image Processing*, vol. 1, Oct. 1997, pp. 255 -258.
- [24] P. Soliz, S. Nemeth, M. Swift, A. Edwards, S. Meuer, J. Berger, "Improving the visualization of drusen in age-related macular degeneration through maximum entropy digitization and stereo viewing" in *SPIE Medical Imaging 2000: Image Perception and Performance*, vol. 3981, April 2000, pp. 271-281.
- [25] A. Thaibaoui, A. Raji and P. Bunel, "A Fuzzy Logic Approach to Drusen Detection in Retinal Angiographic Images" in the proceedings of *International Conference on Pattern Recognition (ICPR'00)*, vol. 4, Sept. 2000, pp. 748-751.
- [26] M. Vidaurrezaga, L.A. Diago and A. Cruz, "Contrast Enhancement with Wavelet Transform in Radiological Images" in the proceedings of *EMBS International Conference*, Chicago, Illinois, July 23-28, 2000, pp. 1760-1763.
- [27] S. Yu and L. Guan, "A CAD System for the Automatic Detection of Clustered Microcalcifications in Digitized Mammogram Films" in *IEEE Transactions on Medical Imaging*, vol. 19, no. 2, Feb. 2000, pp. 115-126.

- [28] N. Zayed, A. Badwi, A. Elsayad, M. Elsherif, A. Youssef, "Wavelet segmentation for fetal ultrasound images" in the proceedings of *4th IEEE 2001 Midwest Symposium on Circuits and Systems*, vol. 1, 2001.
- [29] D. Zhang and D. Valentino, "Segmentation of anatomical structures in x-ray computed tomography images using artificial neural networks" in *SPIE Medical Imaging 2002: Image Processing*, vol. 4684, May 2002, pp. 1640-1652.
- [30] L. Zheng, A. Chan, "An artificial intelligent algorithm for tumor detection in screening mammogram " in *IEEE Transactions on Medical Imaging*, vol. 20, no. 7, July 2001, pp. 559-567.